Monitor and troubleshoot

StorageGRID 11.8

NetApp
May 10, 2024

Table of Contents

Monitor and troubleshoot a StorageGRID system

Monitor StorageGRID system

Monitor a StorageGRID system: Overview

Monitor your StorageGRID system regularly to ensure it is performing as expected.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

> To change units for the storage values displayed in the Grid Manager, select the user drop-down in the upper right of the Grid Manager, then select **User preferences**.

About this task

These instructions describe how to:

- View and manage the dashboard
- View the Nodes page
- Monitor these aspects of the system regularly:
  - System health
  - Storage capacity
  - Information lifecycle management
  - Networking and system resources
  - Tenant activity
  - Load balancing operations
  - Grid federation connections
  - Archival capacity
- Manage alerts and legacy alarms
- View log files
- Configure audit messages and log destinations
- Use an external syslog server to collect audit information
- Use SNMP for monitoring
- Obtain additional StorageGRID data, including metrics and diagnostics

View and manage the dashboard

You can use the dashboard to monitor system activities at a glance. You can create custom dashboards to monitor your implementation of StorageGRID.

To change units for the storage values displayed in the Grid Manager, select the user drop-down in the upper right of the Grid Manager, then select **User preferences**.

Your dashboard might be different based on system configuration.

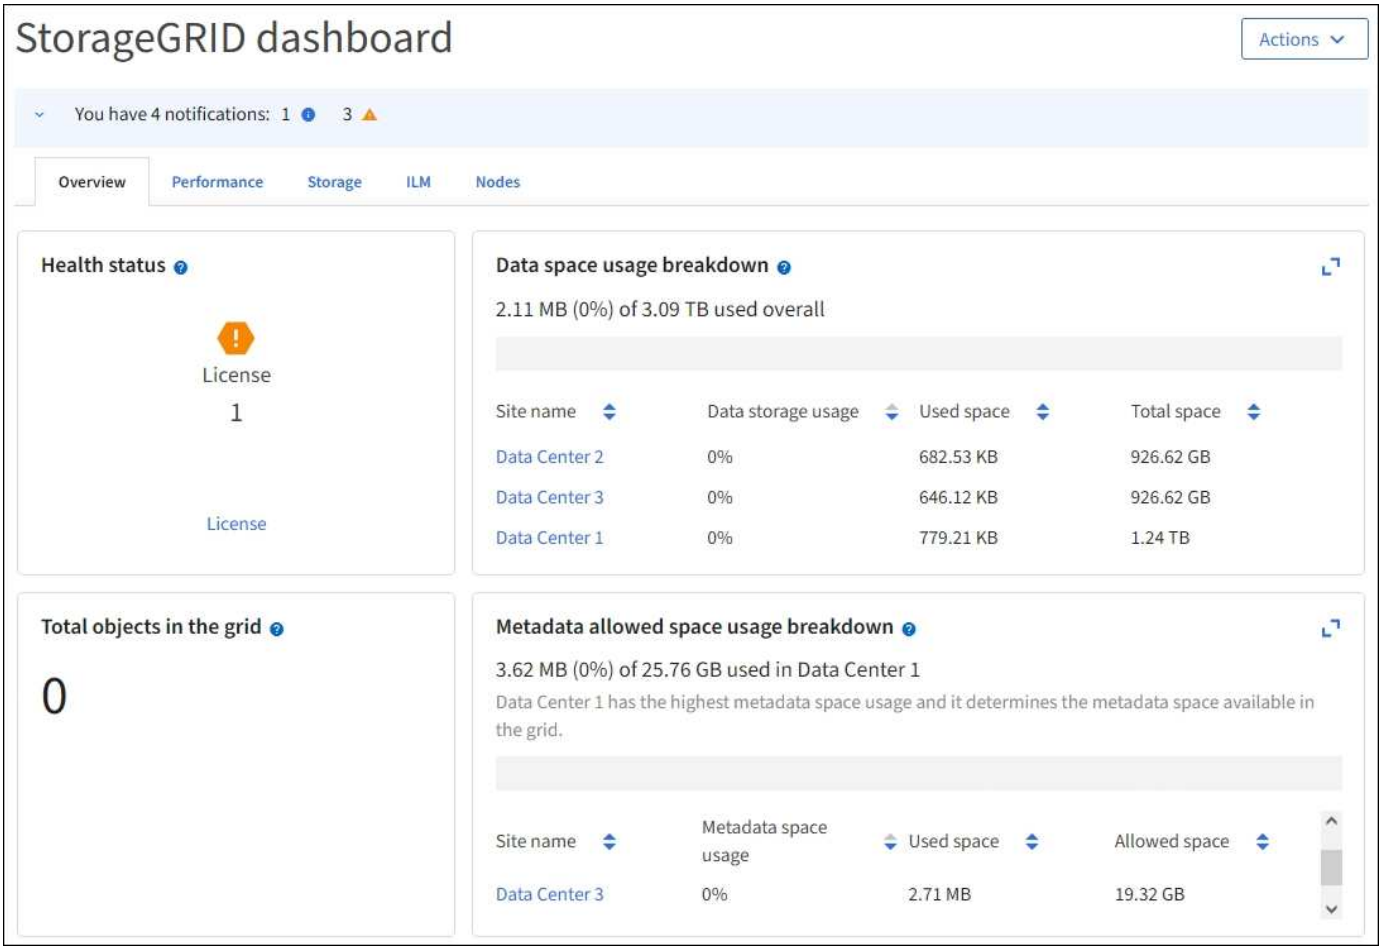

View the dashboard

The dashboard consists of tabs that contain specific information about the StorageGRID system. Each tab contains categories of information displayed on cards.

You can use the system-provided dashboard as is. Additionally, you can create custom dashboards that contain only the tabs and cards that are relevant to monitoring your implementation of StorageGRID.

The system-provided dashboard tabs contain cards with the following types of information:

Tab on system-provided dashboard	Contains
Overview	General information about the grid, such as active alerts, space usage, and total objects in the grid.
Performance	Space usage, storage used over time, S3 or Swift operations, request duration, error rate.

Tab on system-provided dashboard	Contains
Storage	Tenant quota usage and logical space usage. Forecasts of space usage for user data and metadata.
ILM	Information lifecycle management queue and evaluation rate.
Nodes	CPU, data, and memory usage by node. S3 or Swift operations by node. Node to site distribution.

Some of the cards can be maximized for easier viewing. Select the maximize icon in the upper right corner of the card. To close a maximized card, select the minimize icon or select **Close**.

Manage dashboards

If you have Root access (see Admin group permissions), you can perform the following management tasks for dashboards:

- Create a custom dashboard from scratch. You can use custom dashboards to control which StorageGRID information is displayed and how that information is organized.
- Clone a dashboard to create custom dashboards.
- Set an active dashboard for a user. The active dashboard can be the system-provided dashboard or a custom dashboard.
- Set a default dashboard, which is what all users see unless they activate their own dashboard.
- Edit a dashboard name.
- Edit a dashboard to add or remove tabs and cards. You can have a minimum of 1 and a maximum of 20 tabs.
- Remove a dashboard.

If you have any other permission besides Root access, you can only set an active dashboard.

To manage dashboards, select **Actions** > **Manage dashboards**.

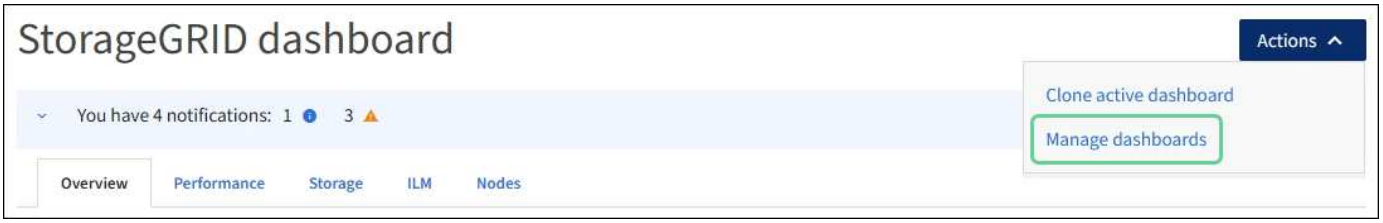

Configure dashboards

To create a new dashboard by cloning the active dashboard, select **Actions** > **Clone active dashboard**.

To edit or clone an existing dashboard, select **Actions** > **Manage dashboards**.

The system-provided dashboard can't be edited or removed.

When configuring a dashboard, you can:

- Add or remove tabs
- Rename tabs and give new tabs unique names
- Add, remove, or rearrange (drag) cards for each tab
- Select the size for individual cards by selecting **S**, **M**, **L** or **XL** at the top of the card

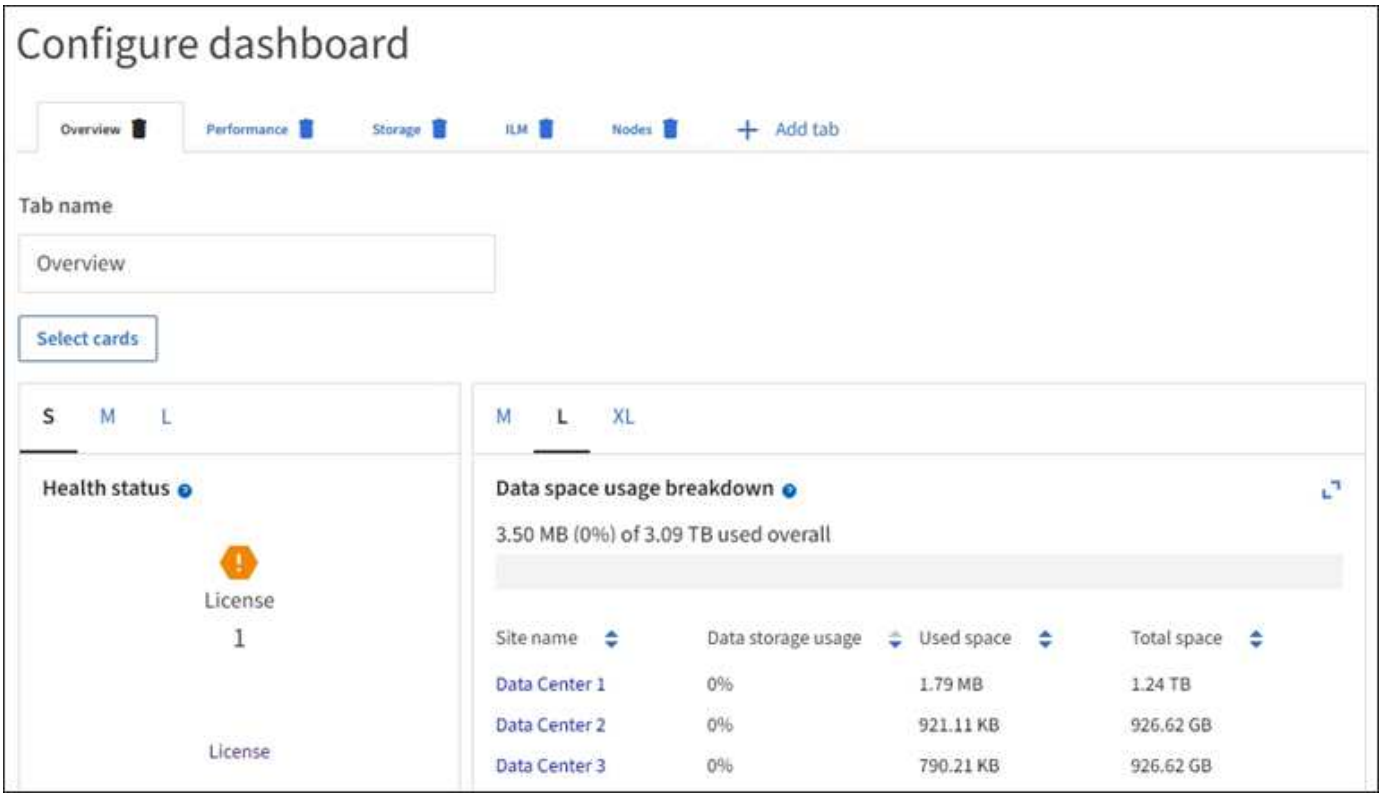

View the Nodes page

View the Nodes page: Overview

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 summary information for the entire grid, each site, and each node. If a node is disconnected or 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.

> When a node is not connected to the grid, such as during upgrade or a disconnected state, certain metrics might be unavailable or excluded from site and grid totals. After a node reconnects to the grid, wait several minutes for the values to stabilize.

> To change units for the storage values displayed in the Grid Manager, select the user drop-down in the upper right of the Grid Manager, then select **User preferences**.

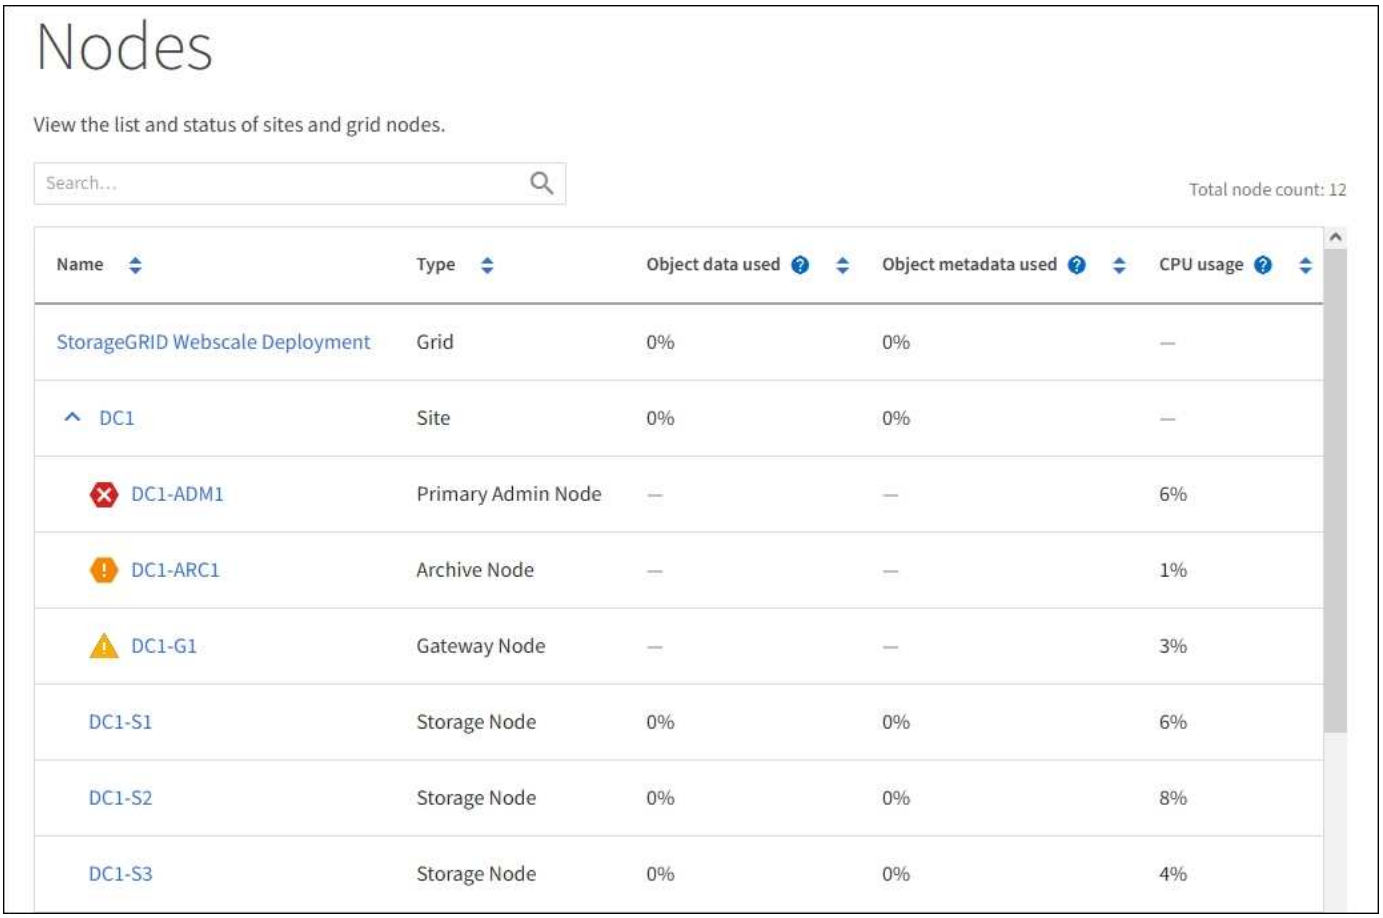

Connection state icons

If a node is disconnected from the grid, either of the following icons appears next to the node name.

Icon	Description	Action required
	Not connected - Unknown For an unknown reason, a node is disconnected or services on the node are unexpectedly down. For example, a service on the node might be stopped, or the node might have lost its network connection because of a power failure or unexpected outage. The **Unable to communicate with node** alert might also be triggered. Other alerts might also be active.	Requires immediate attention. Select each alert and follow the recommended actions. For example, you might need to restart a service that has stopped or restart the host for the node. **Note**: A node might appear as Unknown during managed shutdown operations. You can ignore the Unknown state in these cases.

Icon	Description	Action required
	Not connected - Administratively down For an expected reason, node is not connected to grid. 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. Based on the underlying issue, these nodes often go back online with no intervention.	Determine if any alerts are affecting this node. If one or more alerts are active, Select each alert and follow the recommended actions.

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 don't clear on their own to ensure they don't result in a more serious problem.

View details for a system, site, or node

To filter the information shown in the Nodes table, enter a search string in the **Search** field. You can search by system name, display name, or type (for example, enter **gat** to quickly locate all Gateway Nodes).

To view the information for the grid, site, or node:

- Select the grid name to see an aggregate summary of the statistics for your entire StorageGRID system.
- 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.

View the Overview tab

The Overview tab provides basic information about each node. It also shows any alerts currently affecting the node.

The Overview tab is shown for all nodes.

Node Information

The Node Information section of the Overview tab lists basic information about the node.

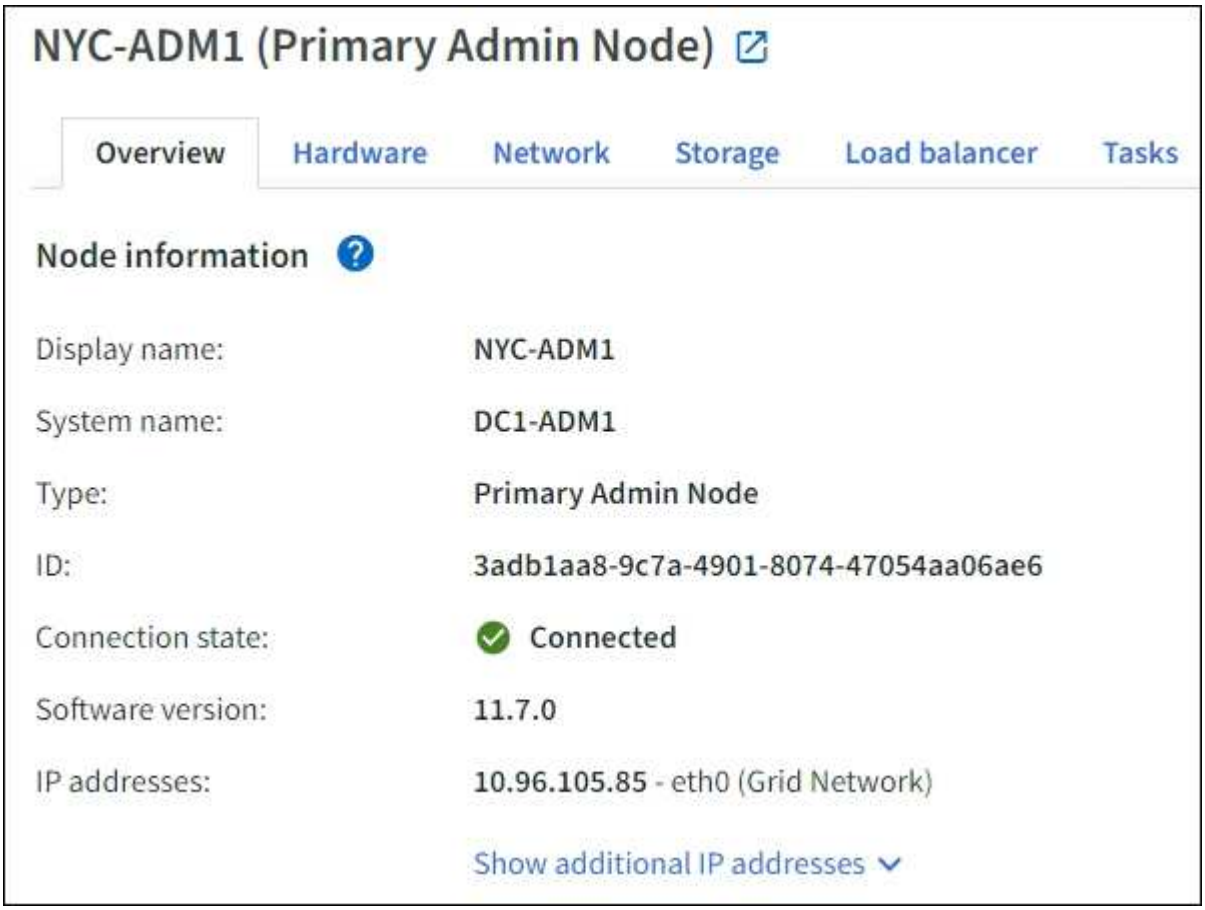

The overview information for a node includes the following:

- **Display name** (shown only if the node has been renamed): The current display name for the node. Use the Rename grid, sites, and nodes procedure to update this value.
- **System name**: The name you entered for the node during installation. System names are used for internal StorageGRID operations and can't be changed.
- **Type**: The type of node — Admin Node, primary Admin Node, Storage Node, Gateway Node, or Archive Node.

> Support for Archive Nodes is deprecated and will be removed in a future release. Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by ILM Cloud Storage Pools, which offer more functionality.

- **ID**: The unique identifier for the node, which is also referred to as the UUID.
- **Connection state**: One of three states. The icon for the most severe state is shown.
  - **Unknown** : For an unknown reason, the node is not connected to the grid, or one or more services are unexpectedly down. For example, the network connection between nodes has been lost, the power is down, or a service 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.

  - **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.
  - **Connected**: The node is connected to the grid.
- **Storage used**: For Storage Nodes only.
  - **Object data**: The percentage of the total usable space for object data that has been used on the Storage Node.
  - **Object metadata**: The percentage of the total allowed space for object metadata that has been used on the Storage Node.
- **Software version**: The version of StorageGRID that is installed on the node.
- **HA groups**: For Admin Node and Gateway Nodes only. Shown if a network interface on the node is included in a high availability group and whether that interface is the Primary interface.
- **IP addresses**: The node's IP addresses. Click **Show additional IP addresses** to view the node's IPv4 and IPv6 addresses and interface mappings.

Alerts

The Alerts section of the Overview tab lists any alerts currently affecting this node that have not been silenced. Select the alert name to view additional details and recommended actions.

Alerts are also included for node connection states.

View the Hardware tab

The Hardware tab displays CPU utilization and memory usage for each node, and additional hardware information about appliances.

> The Grid Manager is updated with each release and might not match the example screenshots on this page.

The Hardware tab is shown for all nodes.

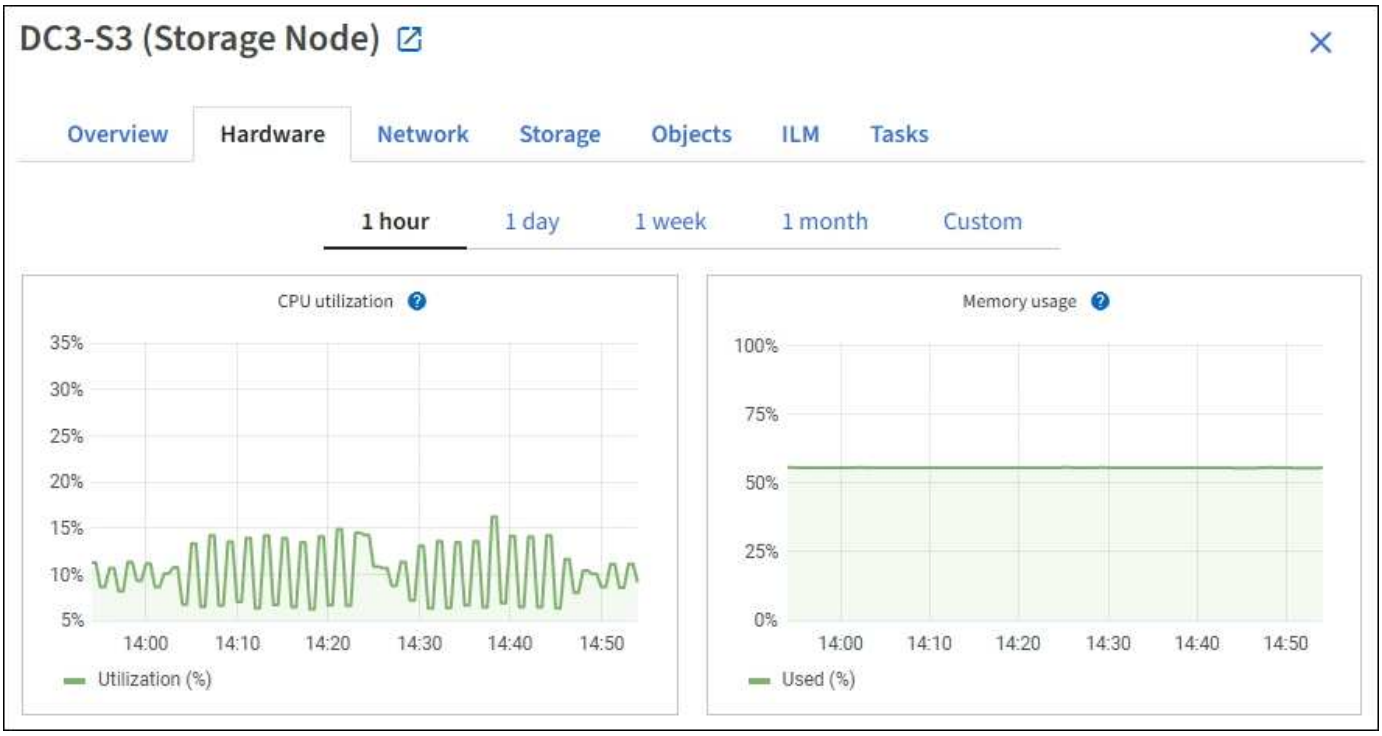

To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.

To see details for CPU utilization and memory usage, position your cursor over each graph.

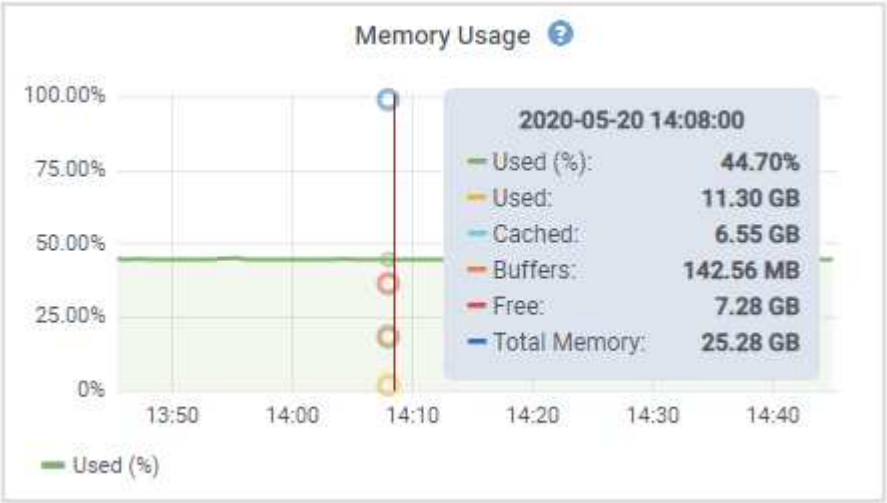

If the node is an appliance node, this tab also includes a section with more information about the appliance hardware.

View information about appliance Storage Nodes

The Nodes page lists information about service health and all computational, disk device, and network resources for each appliance Storage Node. You can also see memory, storage hardware, controller firmware version, network resources, network interfaces, network addresses, and receive and transmit data.

Steps

1. From the Nodes page, select an appliance Storage Node.
2. Select **Overview**.

   The Node information section of the Overview tab displays summary information for the node, such as the node's name, type, ID, and connection state. The list of IP addresses includes the name of the interface for each address, as follows:

   - **eth**: The Grid Network, Admin Network, or Client Network.
   - **hic**: One of the physical 10, 25, or 100 GbE ports on the appliance. These ports can be bonded together and connected to the StorageGRID Grid Network (eth0) and Client Network (eth2).
   - **mtc**: One of the physical 1 GbE ports on the appliance. One or more mtc interfaces are bonded to form the StorageGRID Admin Network interface (eth1). You can leave other mtc interfaces available for temporary local connectivity for a technician in the data center.

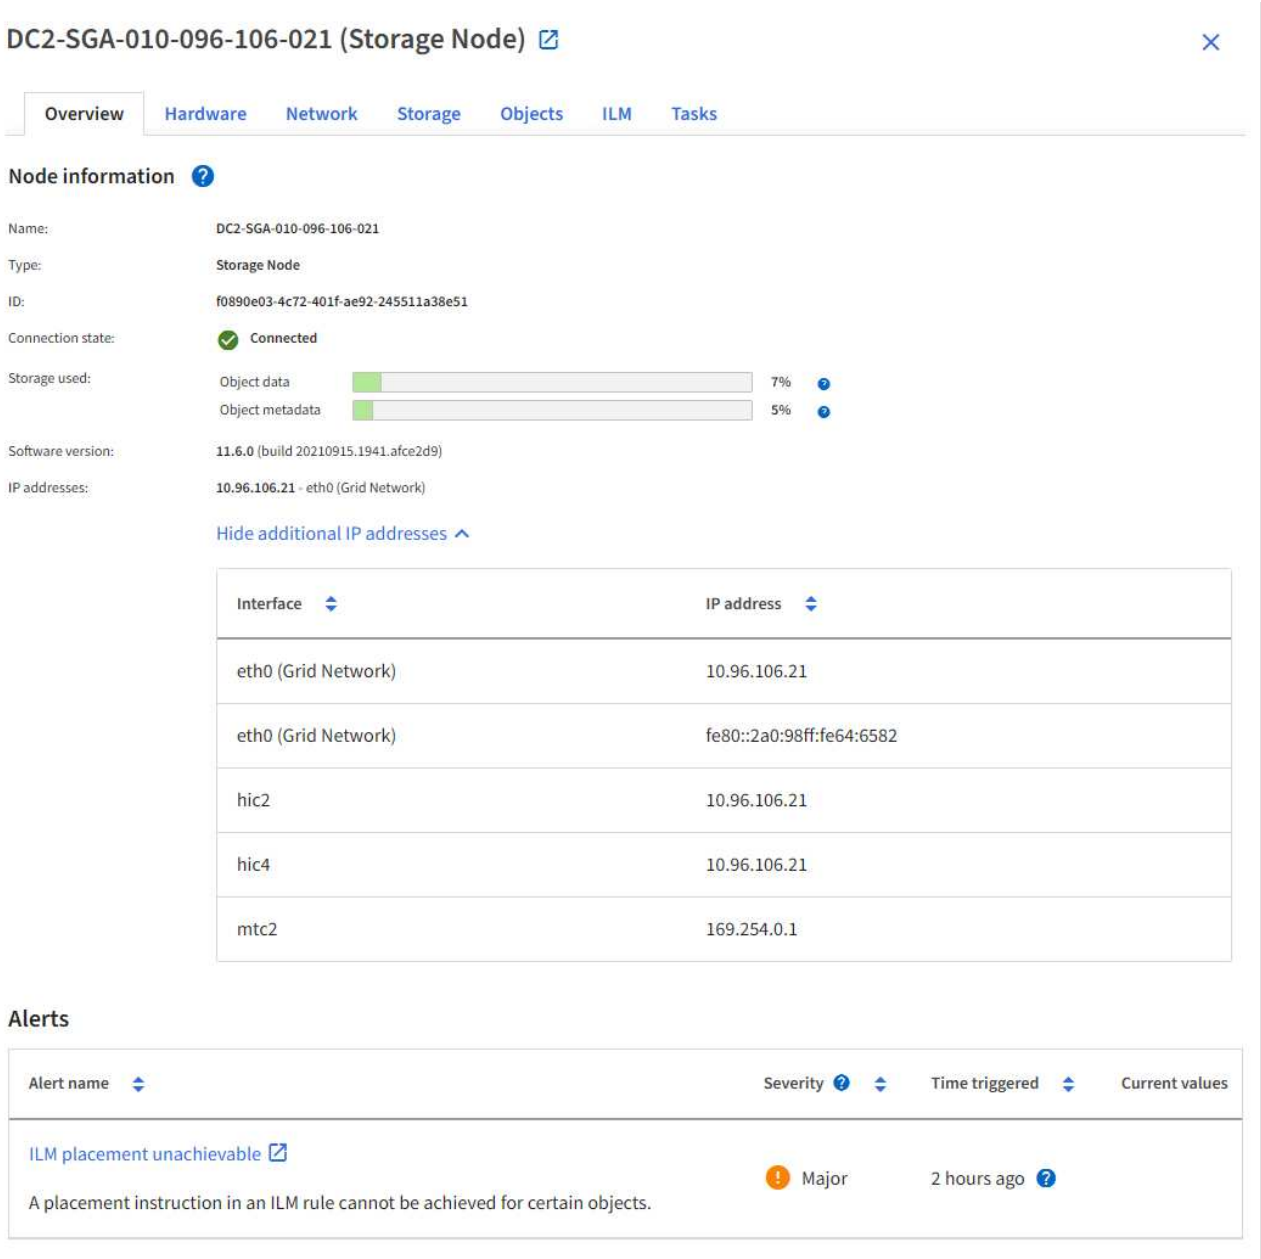

   The Alerts section of the Overview tab displays any active alerts for the node.

3. Select **Hardware** to see more information about the appliance.
   a. View the CPU Utilization and Memory graphs to determine the percentages of CPU and memory usage over time. To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.

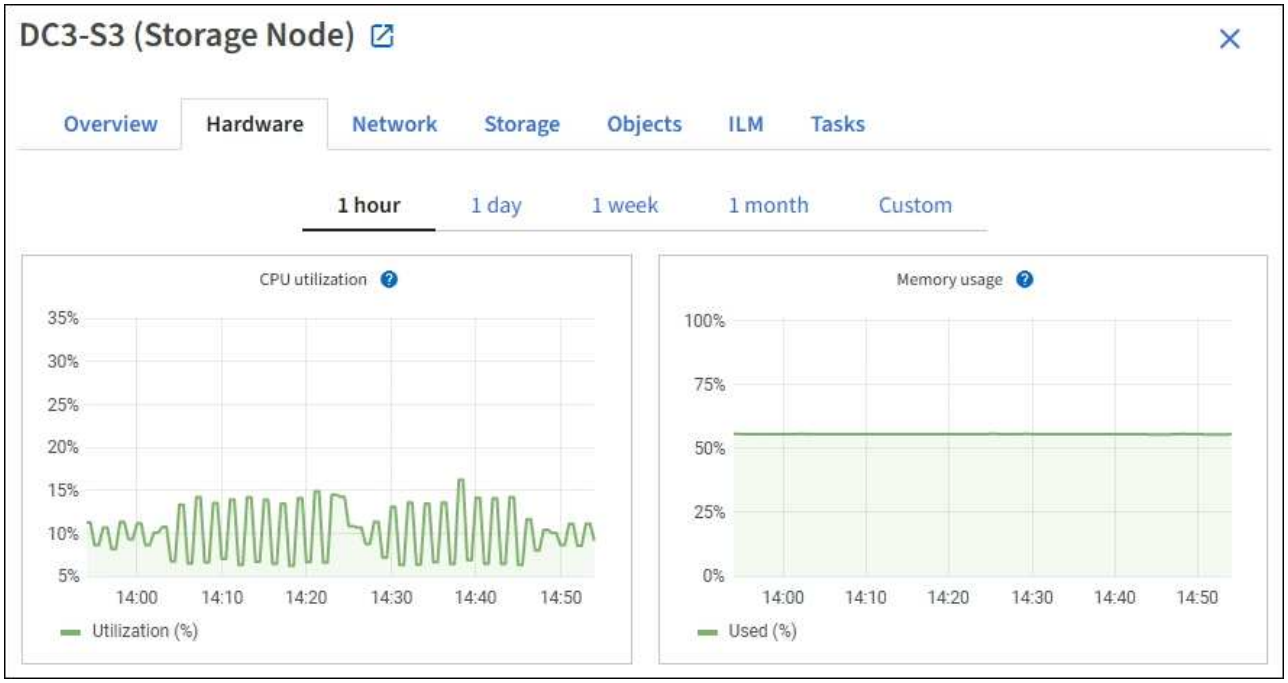

   b. Scroll down to view the table of components for the appliance. This table contains information such as the model name of the appliance; controller names, serial numbers, and IP addresses; and the status of each component.

> Some fields, such as Compute controller BMC IP and Compute hardware, appear only for appliances with that feature.

   Components for the storage shelves, and expansion shelves if they are part of the installation, appear in a separate table below the appliance table.

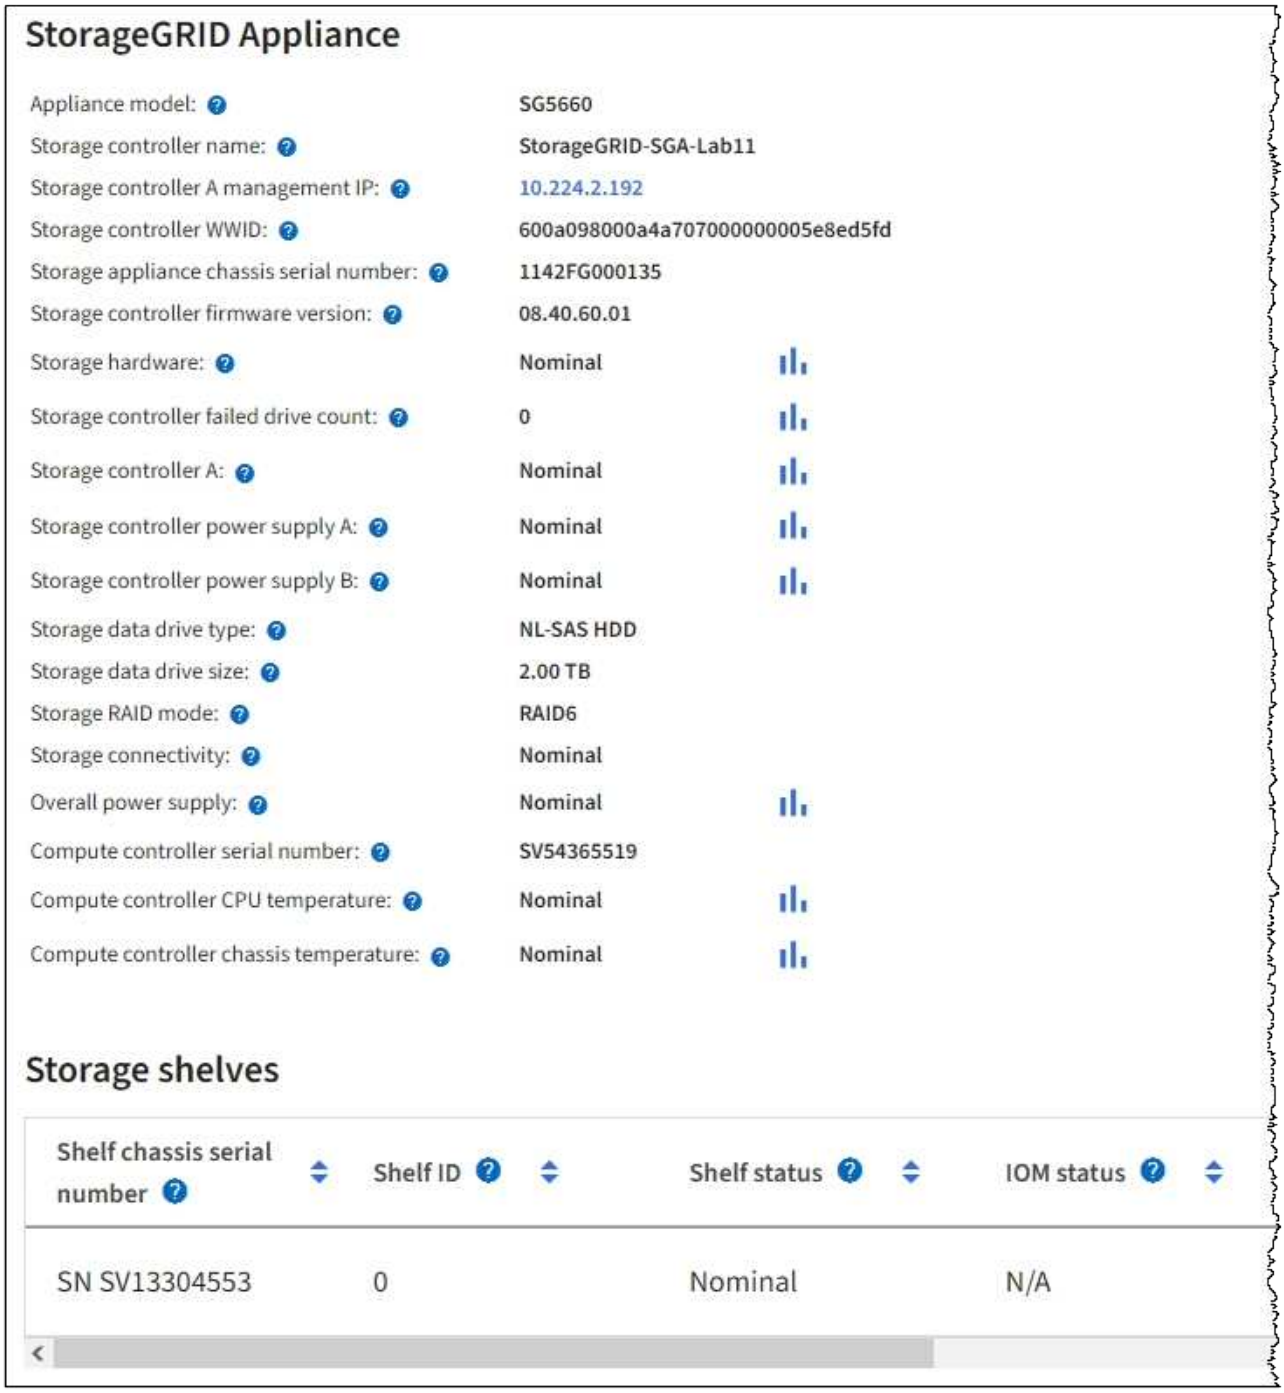

Field in the Appliance table	Description
Appliance model	The model number for this StorageGRID appliance shown in SANtricity OS.
Storage controller name	The name for this StorageGRID appliance shown in SANtricity OS.
Storage controller A management IP	IP address for management port 1 on storage controller A. You use this IP to access SANtricity OS to troubleshoot storage issues.

Field in the Appliance table	Description
Storage controller B management IP	IP address for management port 1 on storage controller B. You use this IP to access SANtricity OS to troubleshoot storage issues. Some appliance models don't have a storage controller B.
Storage controller WWID	The worldwide identifier of the storage controller shown in SANtricity OS.
Storage appliance chassis serial number	The chassis serial number of the appliance.
Storage controller firmware version	The version of the firmware on the storage controller for this appliance.
Storage hardware	The overall status of the storage controller hardware. If SANtricity System Manager reports a status of Needs Attention for the storage hardware, the StorageGRID system also reports this value. If the status is "needs attention," first check the storage controller using SANtricity OS. Then, ensure that no other alarms exist that apply to the compute controller.
Storage controller failed drive count	The number of drives that aren't optimal.
Storage controller A	The status of storage controller A.
Storage controller B	The status of storage controller B. Some appliance models don't have a storage controller B.
Storage controller power supply A	The status of power supply A for the storage controller.
Storage controller power supply B	The status of power supply B for the storage controller.
Storage data drive type	The type of drives in the appliance, such as HDD (hard drive) or SSD (solid state drive).
Storage data drive size	The effective size of one data drive. **Note**: For nodes with expansion shelves, use the Data drive size for each shelf instead. Effective drive size might differ by shelf.
Storage RAID mode	The RAID mode configured for the appliance.

Field in the Appliance table	Description
Storage connectivity	The storage connectivity state.
Overall power supply	The status of all power supplies for the appliance.
Compute controller BMC IP	The IP address of the baseboard management controller (BMC) port in the compute controller. You use this IP to connect to the BMC interface to monitor and diagnose the appliance hardware. This field is not displayed for appliance models that don't contain a BMC.
Compute controller serial number	The serial number of the compute controller.
Compute hardware	The status of the compute controller hardware. This field is not displayed for appliance models that don't have separate compute hardware and storage hardware.
Compute controller CPU temperature	The temperature status of the compute controller's CPU.
Compute controller chassis temperature	The temperature status of the compute controller.

Column in the Storage shelves table	Description
Shelf chassis serial number	The serial number for the storage shelf chassis.
Shelf ID	The numeric identifier for the storage shelf. • 99: Storage controller shelf • 0: First expansion shelf • 1: Second expansion shelf **Note:** Expansion shelves apply to the SG6060 only.
Shelf status	The overall status of the storage shelf.
IOM status	The status of the input/output modules (IOMs) in any expansion shelves. N/A if this is not an expansion shelf.
Power supply status	The overall status of the power supplies for the storage shelf.

Column in the Storage shelves table	Description
Drawer status	The status of the drawers in the storage shelf. N/A if the shelf does not contain drawers.
Fan status	The overall status of the cooling fans in the storage shelf.
Drive slots	The total number of drive slots in the storage shelf.
Data drives	The number of drives in the storage shelf that are used for data storage.
Data drive size	The effective size of one data drive in the storage shelf.
Cache drives	The number of drives in the storage shelf that are used as cache.
Cache drive size	The size of the smallest cache drive in the storage shelf. Normally, cache drives are all the same size.
Configuration status	The configuration status of the storage shelf.

c. Confirm that all statuses are "Nominal."

   If a status is not "Nominal," review any current alerts. You can also use SANtricity System Manager to learn more about some of these hardware values. See the instructions for installing and maintaining your appliance.

4. Select **Network** to view information for each network.

   The Network Traffic graph provides a summary of overall network traffic.

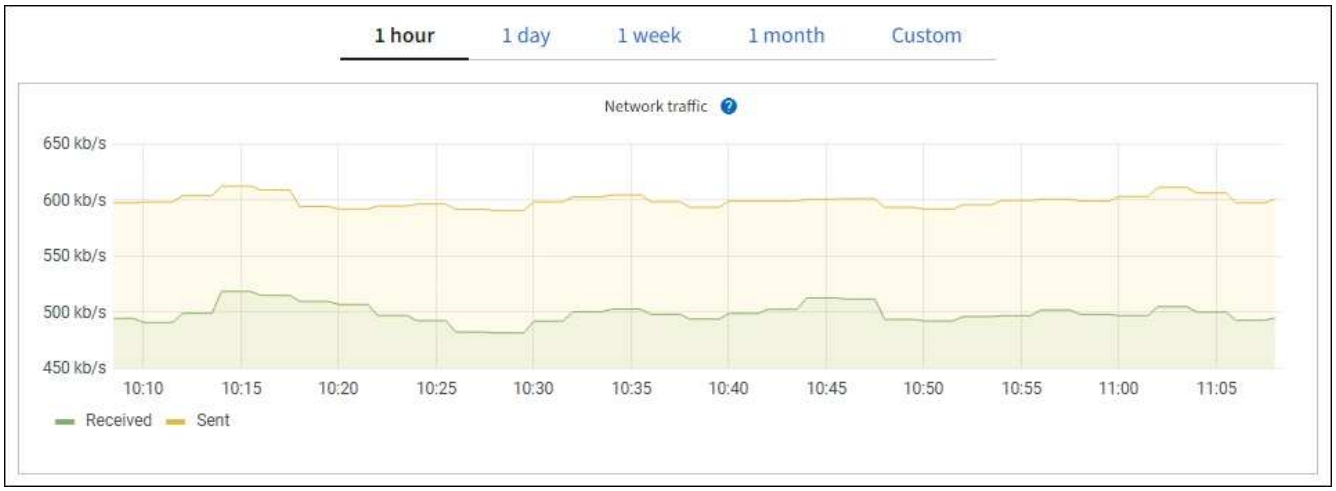

   a. Review the Network Interfaces section.

Network interfaces

Name	Hardware address	Speed	Duplex	Auto-negotiation	Link status
eth0	00:50:56:A7:66:75	10 Gigabit	Full	Off	Up

Use the following table with the values in the **Speed** column in the Network Interfaces table to determine whether the 10/25-GbE network ports on the appliance were configured to use active/backup mode or LACP mode.

The values shown in the table assume all four links are used.

Link mode	Bond mode	Individual HIC link speed (hic1, hic2, hic3, hic4)	Expected Grid/Client Network speed (eth0,eth2)
Aggregate	LACP	25	100
Fixed	LACP	25	50
Fixed	Active/Backup	25	25
Aggregate	LACP	10	40
Fixed	LACP	10	20
Fixed	Active/Backup	10	10

See Configure network links for more information about configuring the 10/25-GbE ports.

b. Review the Network Communication section.

The Receive and Transmit tables show how many bytes and packets have been received and sent across each network as well as other receive and transmit metrics.

Network communication

Receive

Interface	Data	Packets	Errors	Dropped	Frame overruns	Frames
eth0	2.89 GB	19,421,503	0	24,032	0	0

Transmit

Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	3.64 GB	18,494,381	0	0	0	0

5. Select **Storage** to view graphs that show the percentages of storage used over time for object data and object metadata, as well as information about disk devices, volumes, and object stores.

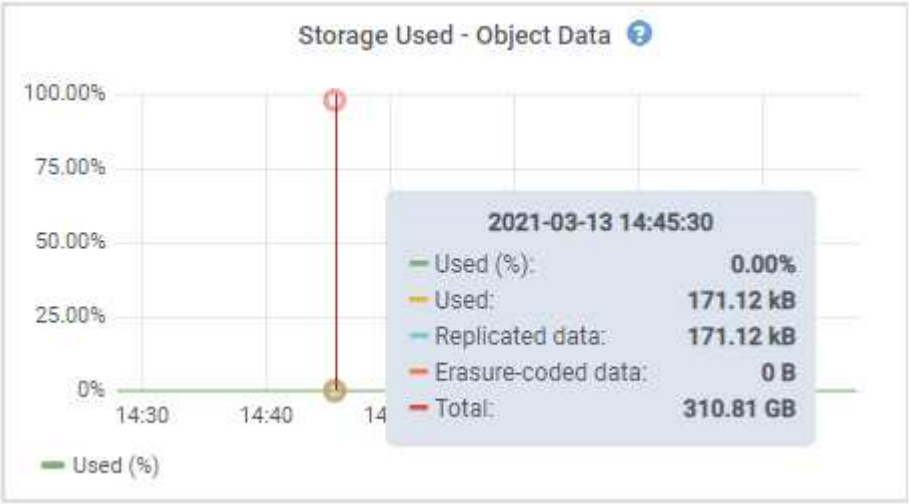

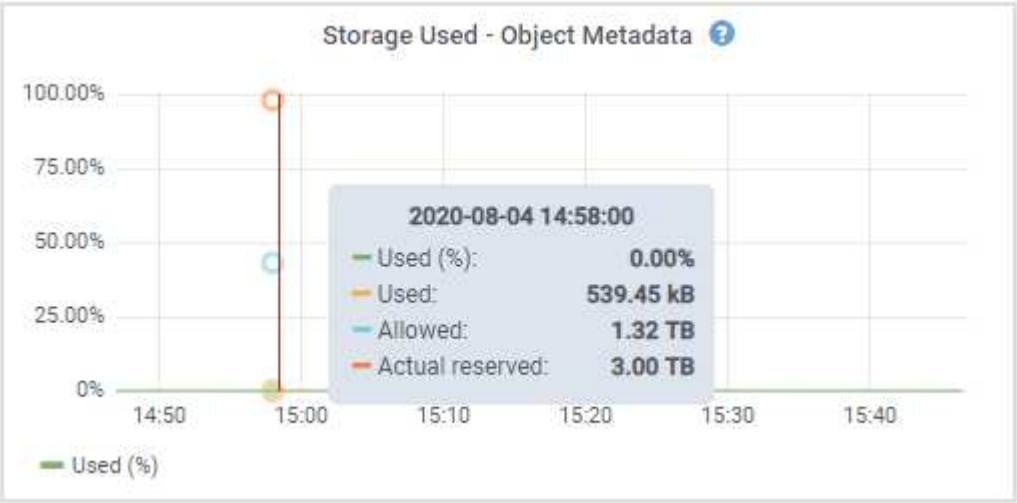

a. Scroll down to view the amounts of available storage for each volume and object store.

   The Worldwide Name for each disk matches the volume world-wide identifier (WWID) that appears

when you view standard volume properties in SANtricity OS (the management software connected to the appliance's storage controller).

To help you interpret disk read and write statistics related to volume mount points, the first portion of the name shown in the **Name** column of the Disk Devices table (that is, *sdc*, *sdd*, *sde*, and so on) matches the value shown in the **Device** column of the Volumes table.

Disk devices

Name	World Wide Name	I/O load	Read rate	Write rate
croot(8:1,sda1)	N/A	0.04%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.67%	0 bytes/s	50 KB/s
sdc(8:16,sdb)	N/A	0.03%	0 bytes/s	4 KB/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s

Volumes

Mount point	Device	Status	Size	Available	Write cache status
/	croot	Online	21.00 GB	14.75 GB	Unknown
/var/local	cvloc	Online	85.86 GB	84.05 GB	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.17 GB	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/2	sde	Online	107.32 GB	107.18 GB	Enabled

Object stores

ID	Size	Available	Replicated data	EC data	Object data (%)	Health
0000	107.32 GB	96.44 GB	124.60 KB	0 bytes	0.00%	No Errors
0001	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0002	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors

View information about appliance Admin Nodes and Gateway Nodes

The Nodes page lists information about service health and all computational, disk device, and network resources for each services appliance that is used as an Admin Node or a Gateway Node. You can also see memory, storage hardware, network resources, network interfaces, network addresses, and receive and transmit data.

Steps

1. From the Nodes page, select an appliance Admin Node or an appliance Gateway Node.
2. Select **Overview**.

   The Node information section of the Overview tab displays summary information for the node, such as the node's name, type, ID, and connection state. The list of IP addresses includes the name of the interface for each address, as follows:

   - **adllb** and **adlli**: Shown if active/backup bonding is used for the Admin Network interface
   - **eth**: The Grid Network, Admin Network, or Client Network.
   - **hic**: One of the physical 10, 25, or 100 GbE ports on the appliance. These ports can be bonded together and connected to the StorageGRID Grid Network (eth0) and Client Network (eth2).
   - **mtc**: One of the physical 1-GbE ports on the appliance. One or more mtc interfaces are bonded to form the Admin Network interface (eth1). You can leave other mtc interfaces available for temporary local connectivity for a technician in the data center.

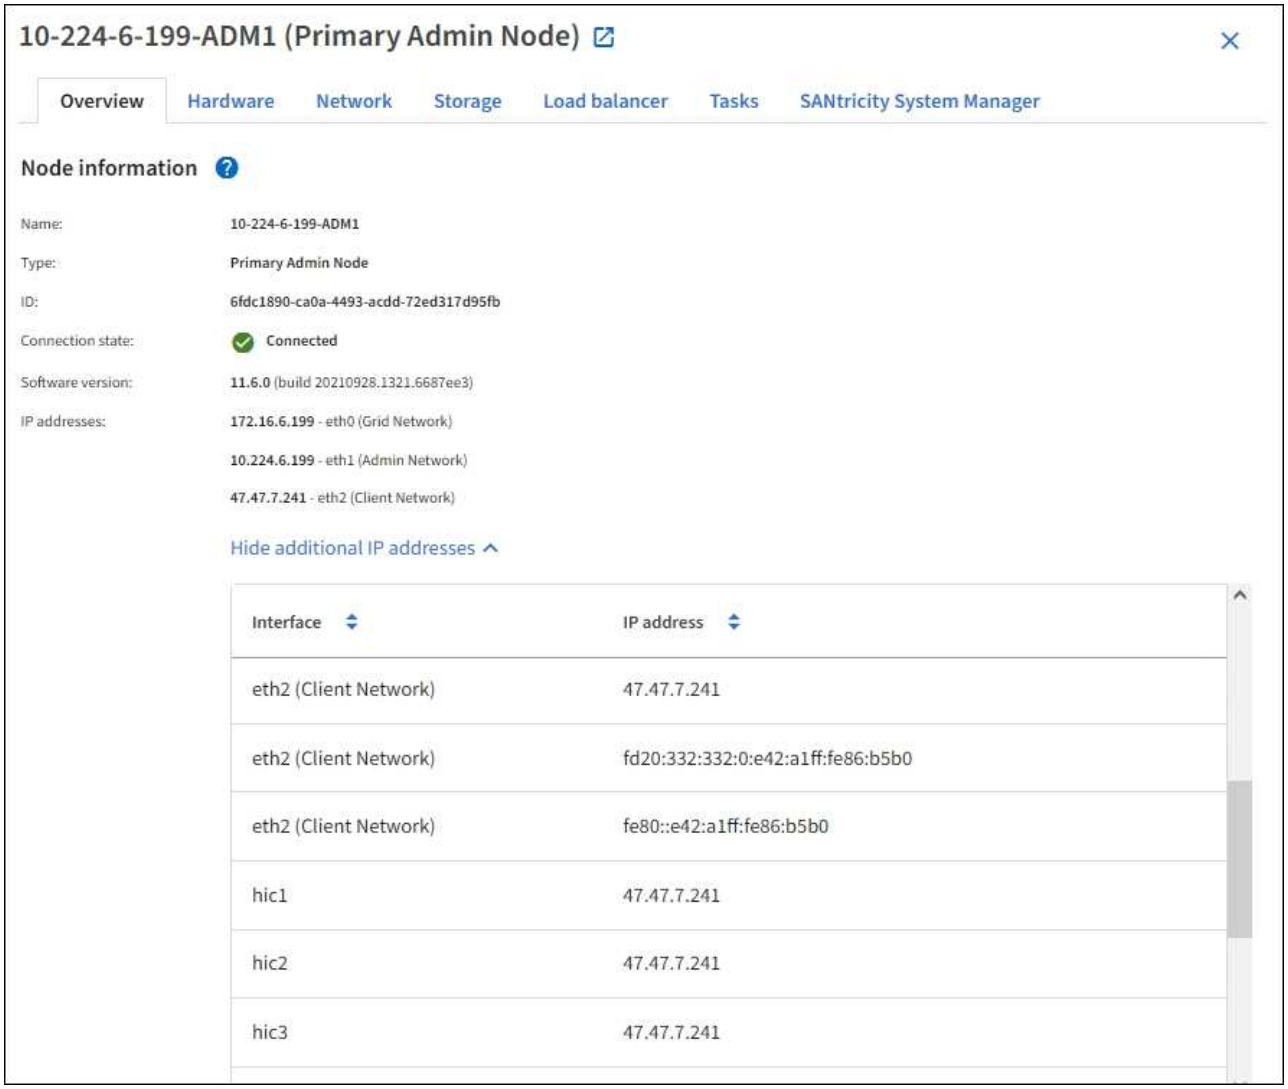

   The Alerts section of the Overview tab displays any active alerts for the node.

3. Select **Hardware** to see more information about the appliance.
   a. View the CPU Utilization and Memory graphs to determine the percentages of CPU and memory usage over time. To display a different time interval, select one of the controls above the chart or graph. You

can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.

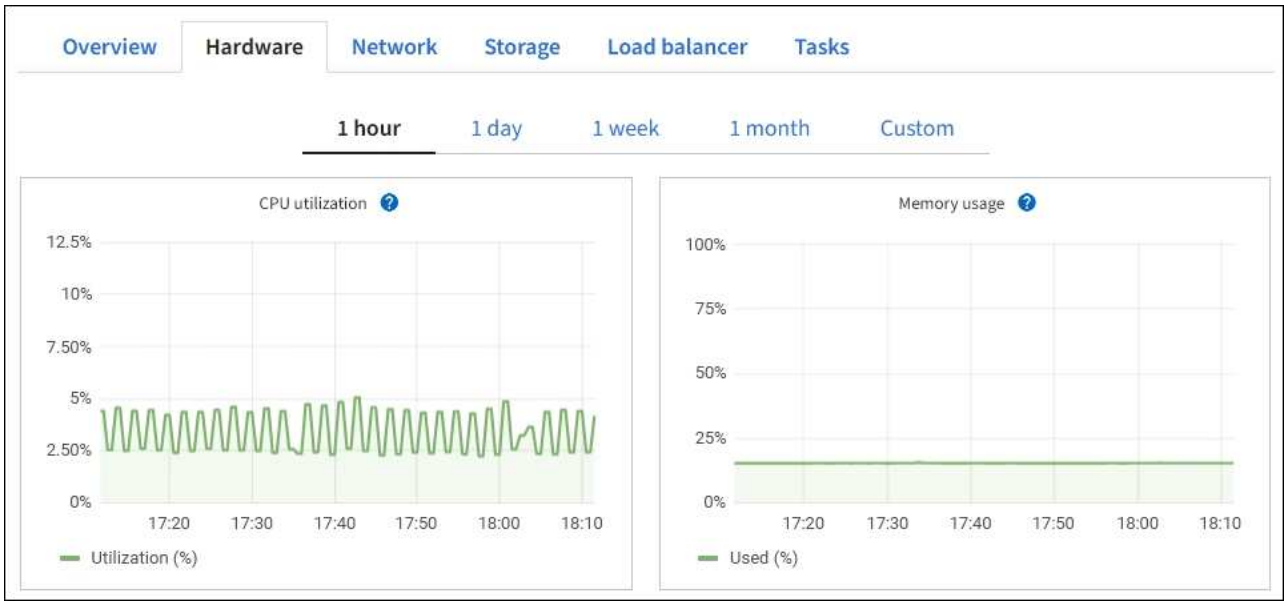

b. Scroll down to view the table of components for the appliance. This table contains information such as the model name, serial number, controller firmware version, and the status of each component.

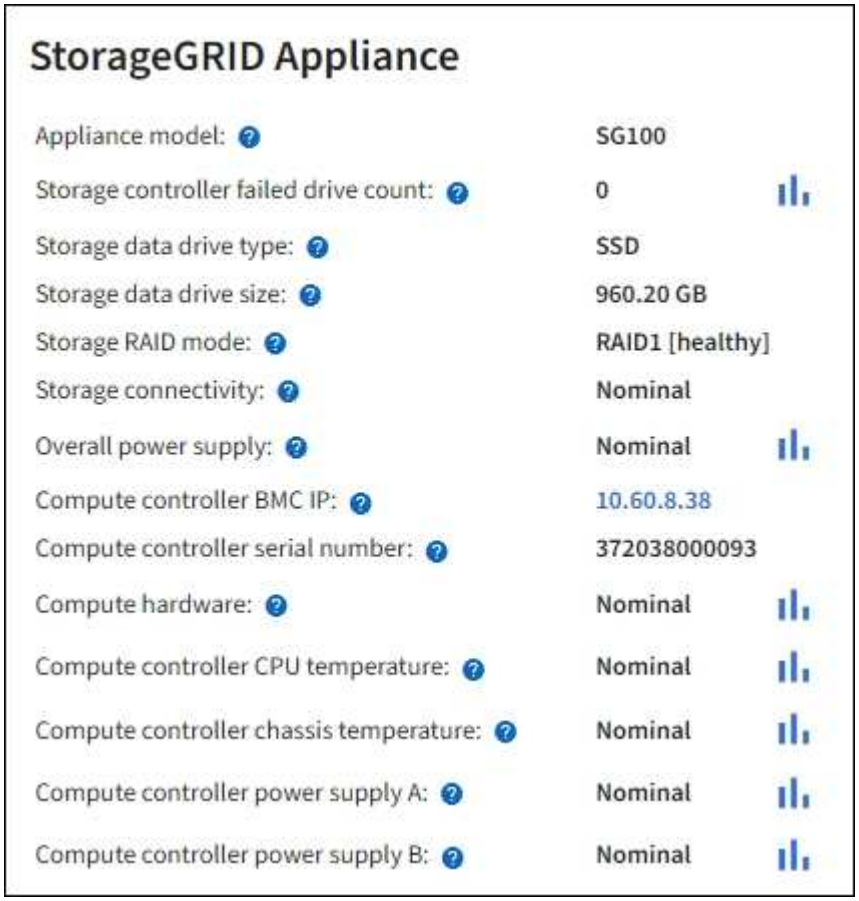

Field in the Appliance table	Description
Appliance model	The model number for this StorageGRID appliance.

Field in the Appliance table	Description
Storage controller failed drive count	The number of drives that aren't optimal.
Storage data drive type	The type of drives in the appliance, such as HDD (hard drive) or SSD (solid state drive).
Storage data drive size	The effective size of one data drive.
Storage RAID mode	The RAID mode for the appliance.
Overall power supply	The status of all power supplies in the appliance.
Compute controller BMC IP	The IP address of the baseboard management controller (BMC) port in the compute controller. You can use this IP to connect to the BMC interface to monitor and diagnose the appliance hardware. This field is not displayed for appliance models that don't contain a BMC.
Compute controller serial number	The serial number of the compute controller.
Compute hardware	The status of the compute controller hardware.
Compute controller CPU temperature	The temperature status of the compute controller's CPU.
Compute controller chassis temperature	The temperature status of the compute controller.

c. Confirm that all statuses are "Nominal."

   If a status is not "Nominal," review any current alerts.

4. Select **Network** to view information for each network.

   The Network Traffic graph provides a summary of overall network traffic.

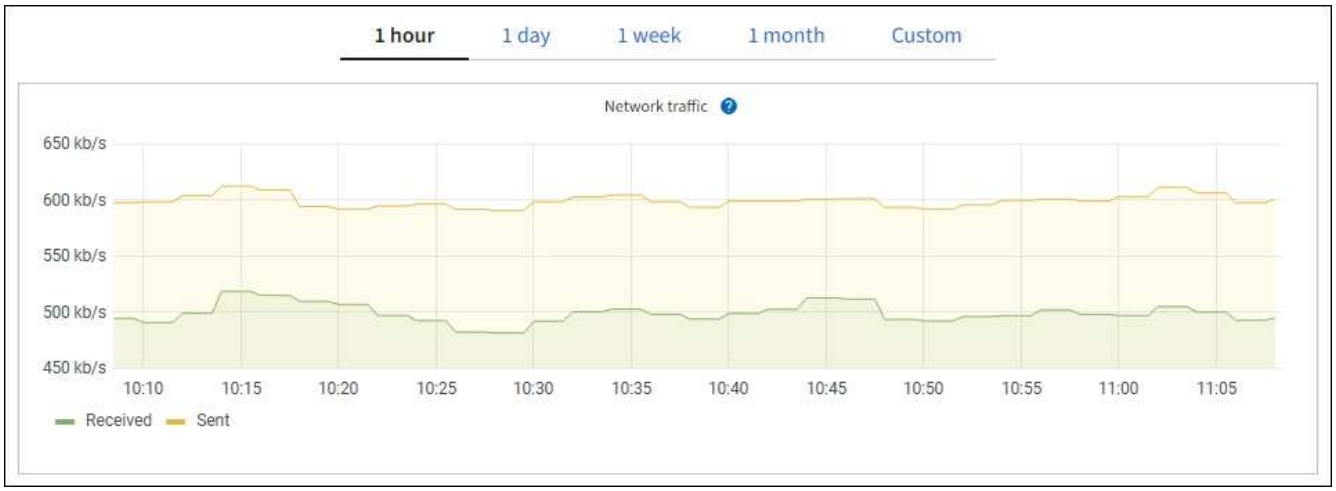

a. Review the Network Interfaces section.

Network interfaces

Name	Hardware address	Speed	Duplex	Auto-negotiation	Link status
eth0	0C:42:A1:86:B5:B0	100 Gigabit	Full	Off	Up
eth1	B4:A9:FC:71:68:36	Gigabit	Full	Off	Up
eth2	0C:42:A1:86:B5:B0	100 Gigabit	Full	Off	Up
hic1	0C:42:A1:86:B5:B0	25 Gigabit	Full	On	Up
hic2	0C:42:A1:86:B5:B0	25 Gigabit	Full	On	Up
hic3	0C:42:A1:86:B5:B0	25 Gigabit	Full	On	Up
hic4	0C:42:A1:86:B5:B0	25 Gigabit	Full	On	Up
mtc1	B4:A9:FC:71:68:36	Gigabit	Full	On	Up
mtc2	B4:A9:FC:71:68:35	Gigabit	Full	On	Up

Use the following table with the values in the **Speed** column in the Network Interfaces table to determine whether the four 40/100-GbE network ports on the appliance were configured to use active/backup mode or LACP mode.

The values shown in the table assume all four links are used.

Link mode	Bond mode	Individual HIC link speed (hic1, hic2, hic3, hic4)	Expected Grid/Client Network speed (eth0, eth2)
Aggregate	LACP	100	400
Fixed	LACP	100	200
Fixed	Active/Backup	100	100
Aggregate	LACP	40	160
Fixed	LACP	40	80
Fixed	Active/Backup	40	40

b. Review the Network Communication section.

The Receive and Transmit tables show how many bytes and packets have been received and sent across each network as well as other receive and transmission metrics.

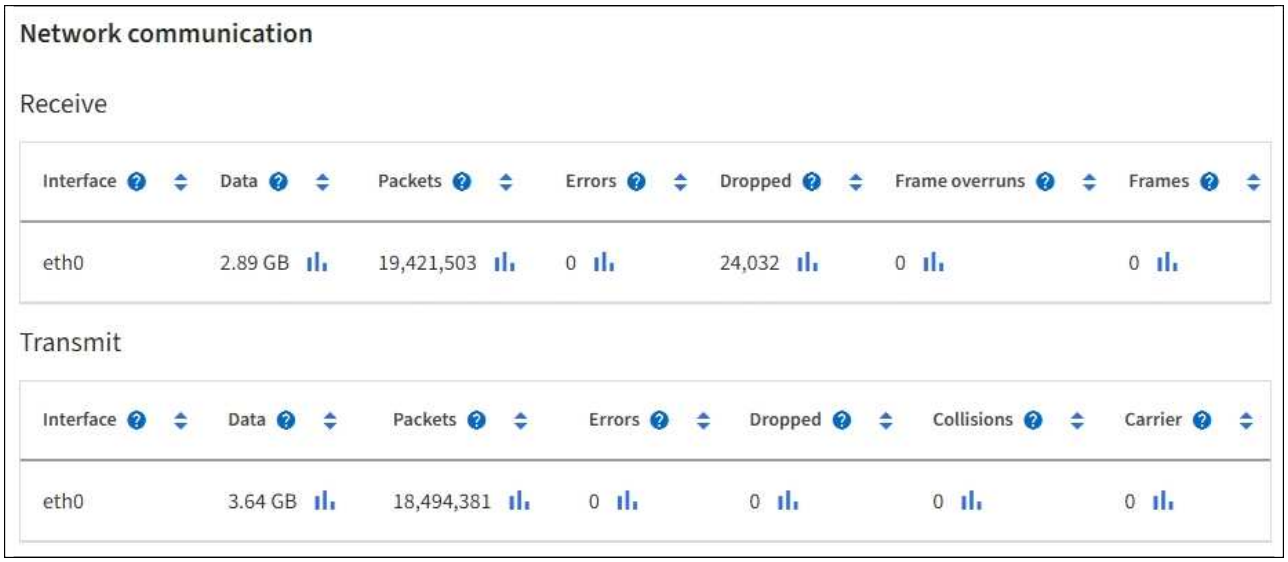

Network communication

Receive

Interface	Data	Packets	Errors	Dropped	Frame overruns	Frames
eth0	2.89 GB	19,421,503	0	24,032	0	0

Transmit

Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	3.64 GB	18,494,381	0	0	0	0

5. Select **Storage** to view information about the disk devices and volumes on the services appliance.

DO-REF-DC1-GW1 (Gateway Node)

Overview | Hardware | Network | Storage | Load balancer | Tasks

Disk devices

Name	World Wide Name	I/O load	Read rate	Write rate
croot(8:1,sda1)	N/A	0.02%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.03%	0 bytes/s	6 KB/s

Volumes

Mount point	Device	Status	Size	Available	Write cache status
/	croot	Online	21.00 GB	14.73 GB	Unknown
/var/local	cvloc	Online	85.86 GB	84.63 GB	Unknown

View the Network tab

The Network tab displays a graph showing the network traffic received and sent across all of the network interfaces on the node, site, or grid.

The Network tab is shown for all nodes, each site, and the entire grid.

To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.

For nodes, the Network interfaces table provides information about each node's physical network ports. The Network communications table provides details about each node's receive and transmit operations and any driver reported fault counters.

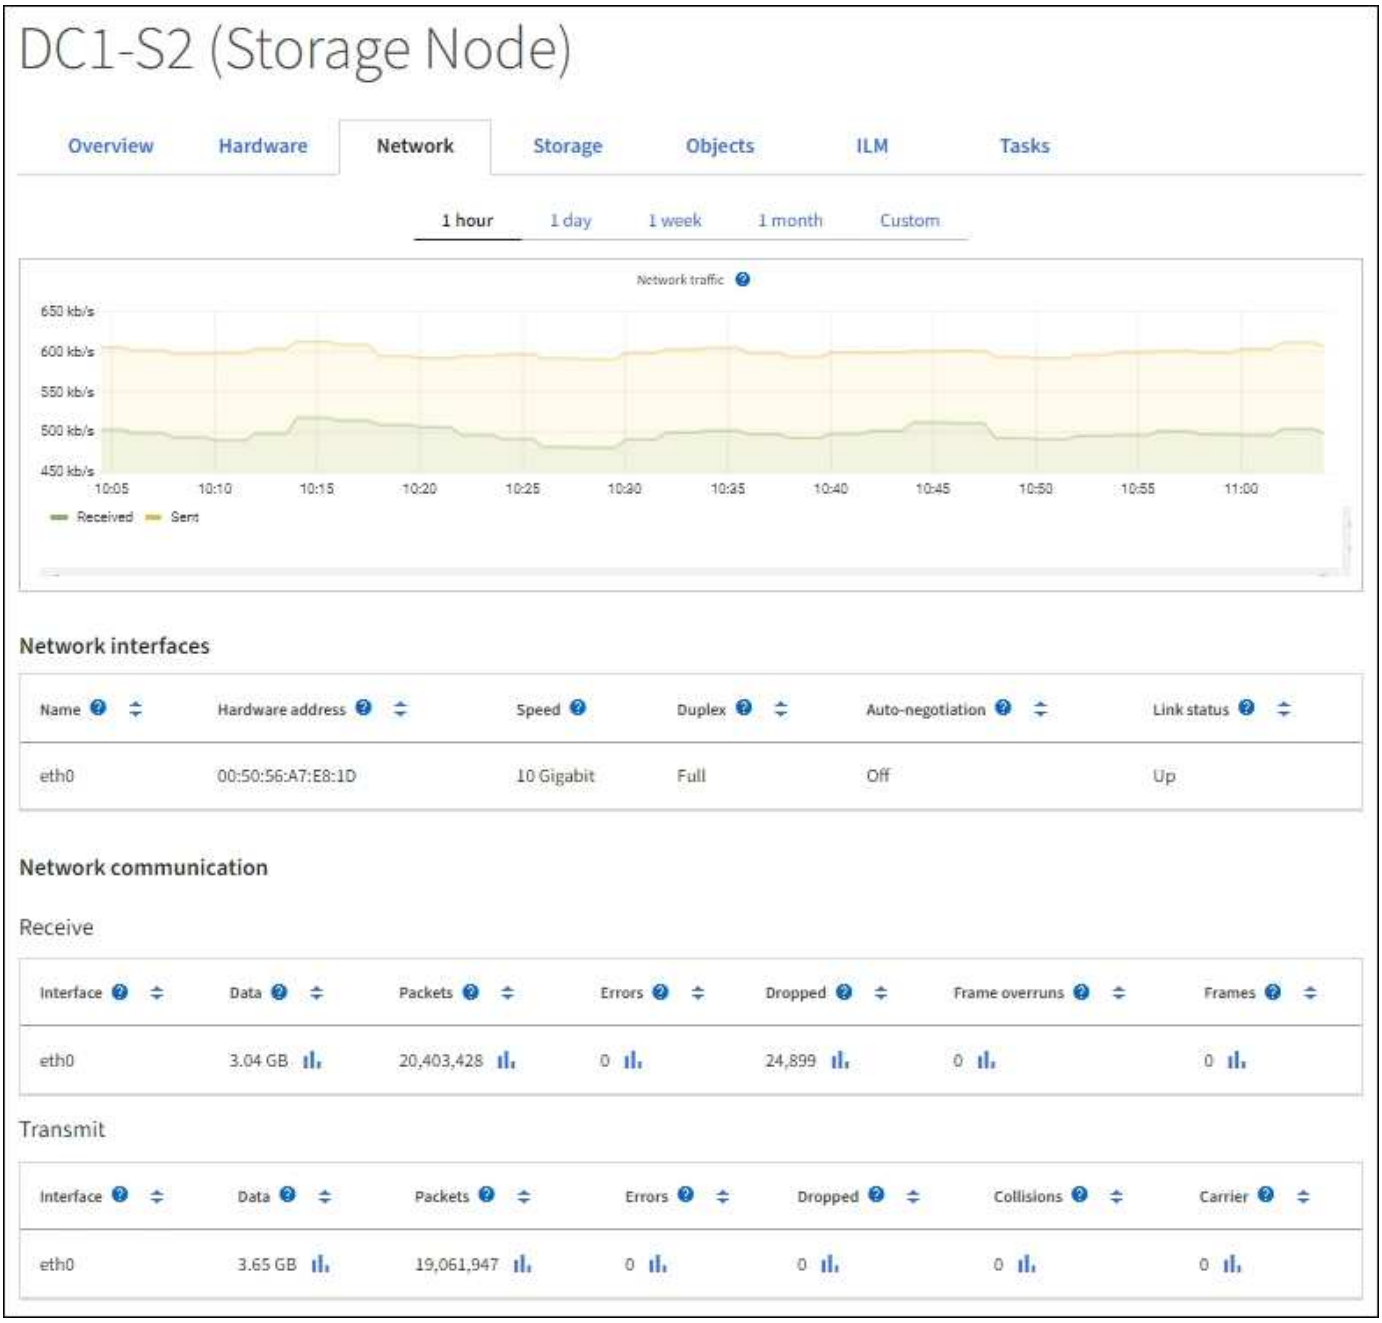

Related information

Monitor network connections and performance

View the Storage tab

The Storage tab summarizes storage availability and other storage metrics.

The Storage tab is shown for all nodes, each site, and the entire grid.

Storage used graphs

For Storage Nodes, each site, and the entire grid, the Storage tab includes graphs showing how much storage has been used by object data and object metadata over time.

When a node is not connected to the grid, such as during upgrade or a disconnected state, certain metrics might be unavailable or excluded from site and grid totals. After a node reconnects to the grid, wait several minutes for the values to stabilize.

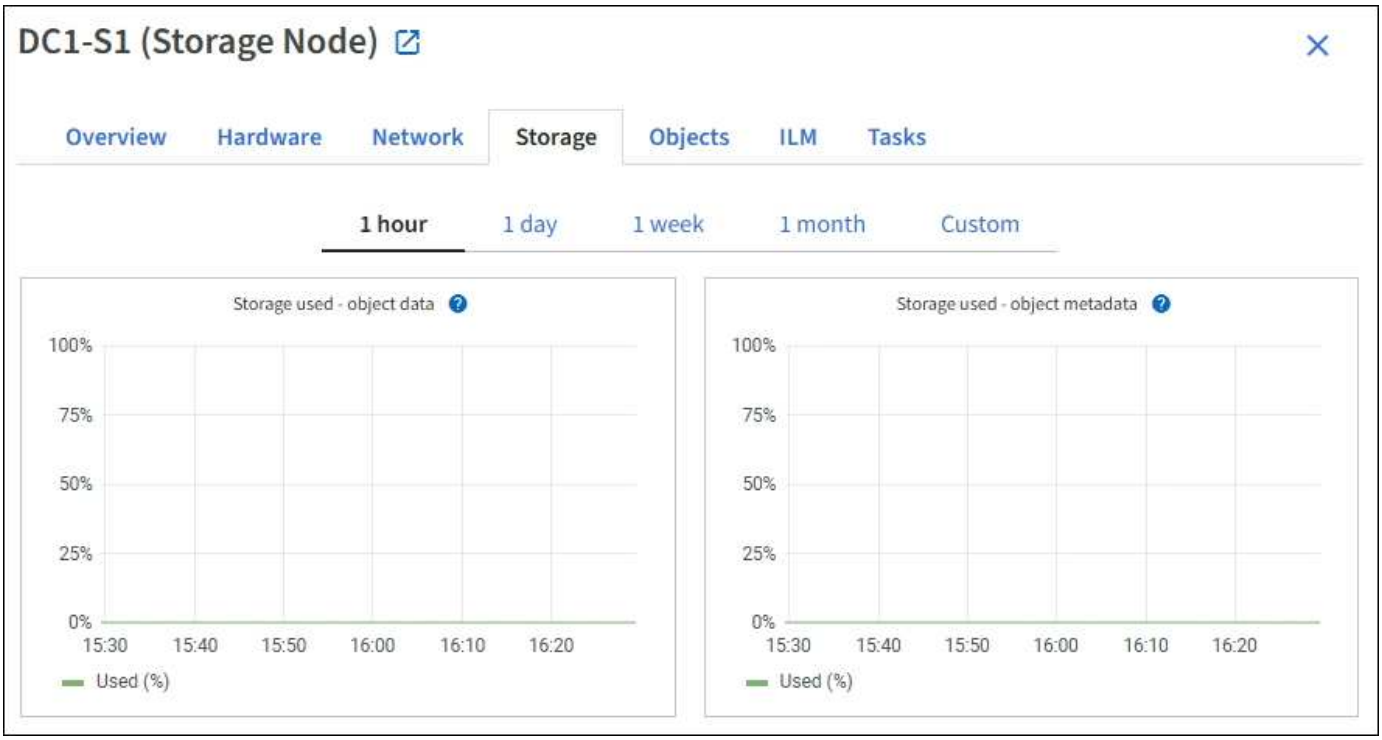

Disk devices, Volumes, and Object stores tables

For all nodes, the Storage tab contains details for the disk devices and volumes on the node. For Storage Nodes, the Object Stores table provides information about each storage volume.

Disk devices

Name	World Wide Name	I/O load	Read rate	Write rate
croot(8:1,sda1)	N/A	0.04%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.67%	0 bytes/s	50 KB/s
sdc(8:16,sdb)	N/A	0.03%	0 bytes/s	4 KB/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s

Volumes

Mount point	Device	Status	Size	Available	Write cache status
/	croot	Online	21.00 GB	14.75 GB	Unknown
/var/local	cvloc	Online	85.86 GB	84.05 GB	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.17 GB	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/2	sde	Online	107.32 GB	107.18 GB	Enabled

Object stores

ID	Size	Available	Replicated data	EC data	Object data (%)	Health
0000	107.32 GB	96.44 GB	124.60 KB	0 bytes	0.00%	No Errors
0001	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0002	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors

Related information

Monitor storage capacity

View the Objects tab

The Objects tab provides information about S3 and Swift ingest and retrieve rates.

The Objects tab is shown for each Storage Node, each site, and the entire grid. For Storage Nodes, the Objects tab also provides object counts and information about metadata queries and background verification.

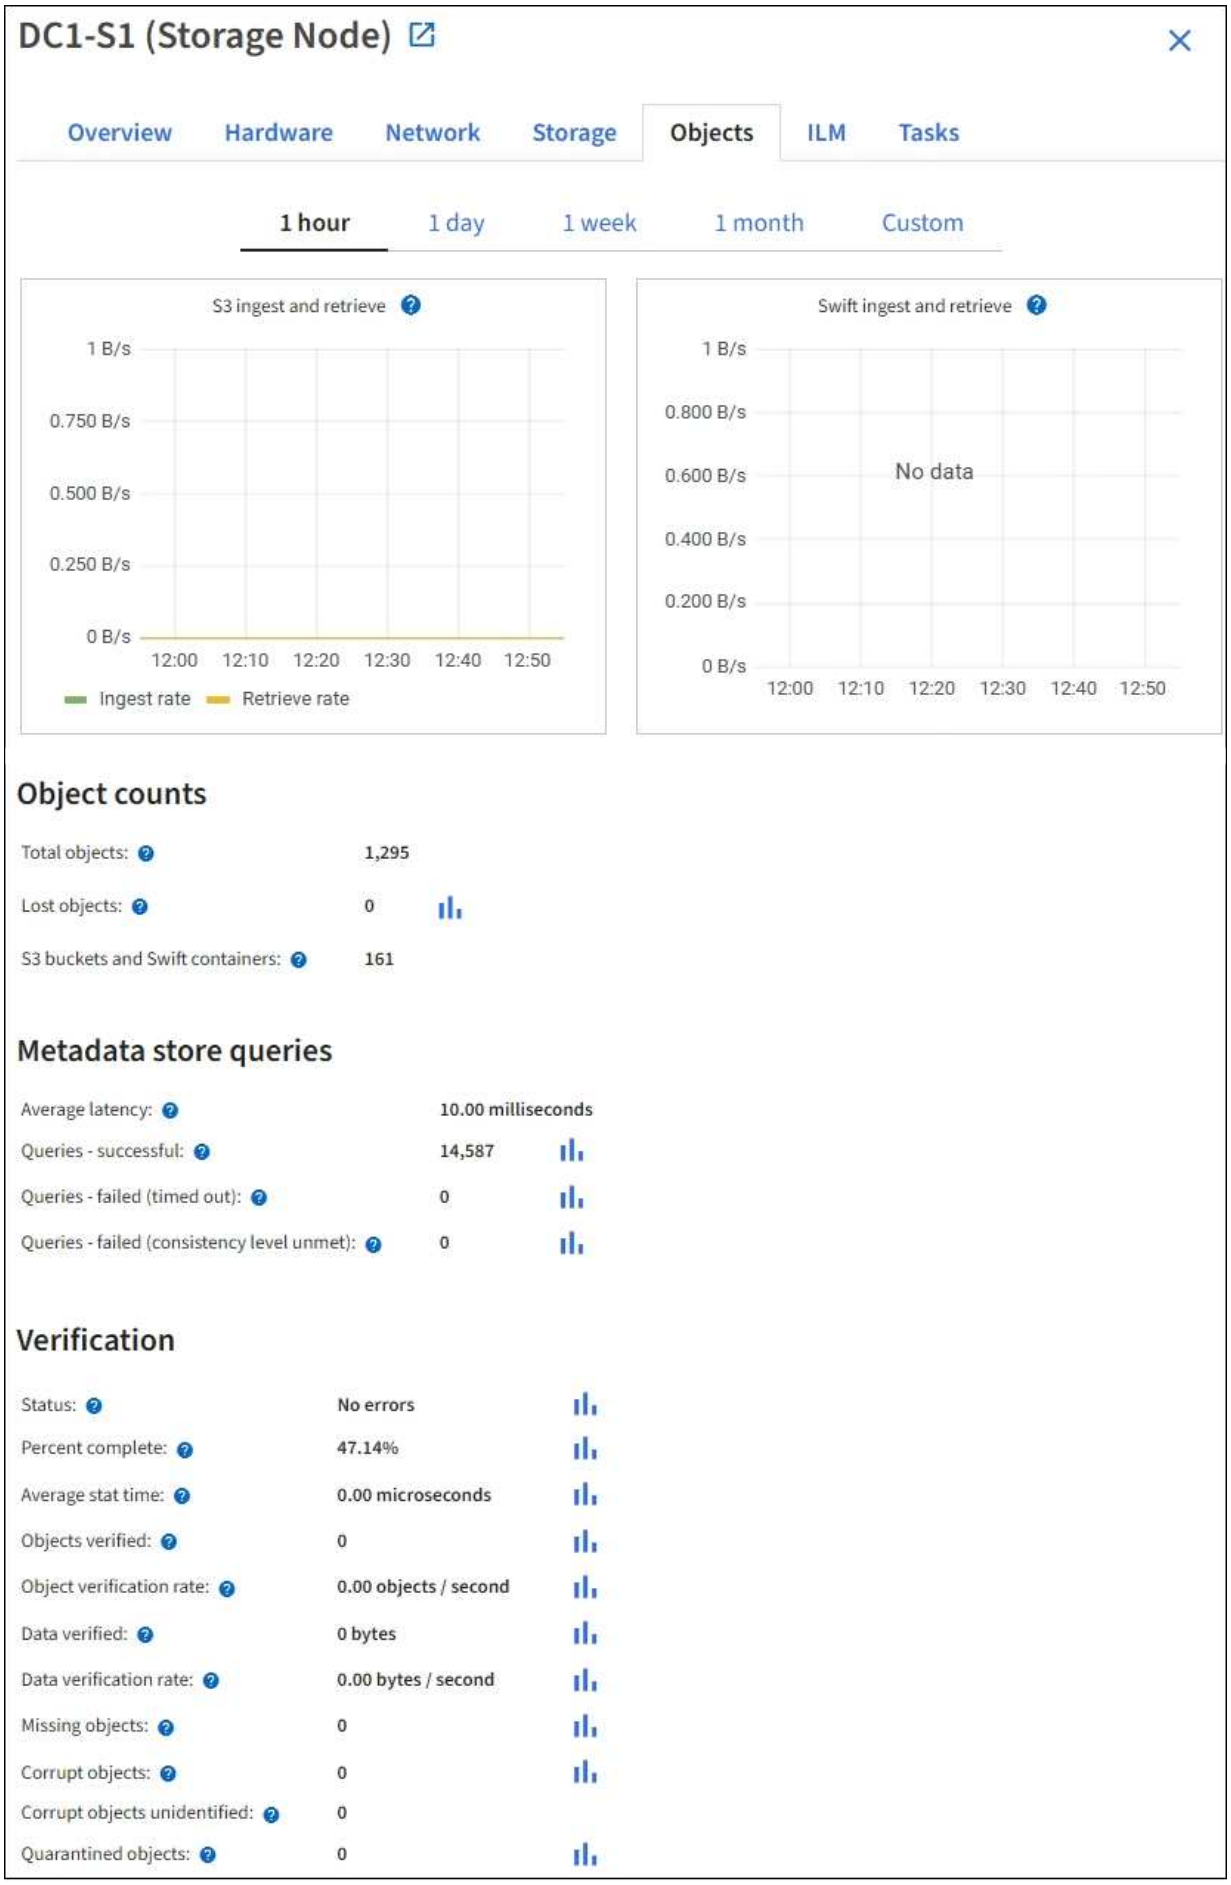
DC1-S1 (Storage Node)
Overview
Hardware
Network
Storage
Objects
ILM
Tasks
1 hour
1 day
1 week
1 month
Custom
S3 ingest and retrieve
1 B/s
0.750 B/s
0.500 B/s
0.250 B/s
0 B/s
12:00
12:10
12:20
12:30
12:40
12:50
Ingest rate
Retrieve rate
Swift ingest and retrieve
1 B/s
0.800 B/s
0.600 B/s
0.400 B/s
0.200 B/s
0 B/s
No data
12:00
12:10
12:20
12:30
12:40
12:50
Object counts
Total objects: 1,295
Lost objects: 0
S3 buckets and Swift containers: 161
Metadata store queries
Average latency: 10.00 milliseconds
Queries - successful: 14,587
Queries - failed (timed out): 0
Queries - failed (consistency level unmet): 0
Verification
Status: No errors
Percent complete: 47.14%
Average stat time: 0.00 microseconds
Objects verified: 0
Object verification rate: 0.00 objects / second
Data verified: 0 bytes
Data verification rate: 0.00 bytes / second
Missing objects: 0
Corrupt objects: 0
Corrupt objects unidentified: 0
Quarantined objects: 0

View the ILM tab

The ILM tab provides information about information lifecycle management (ILM) operations.

The ILM tab is shown for each Storage Node, each site, and the entire grid. For each site and the grid, the ILM tab shows a graph of the ILM queue over time. For the grid, this tab also provides the estimated time to complete a full ILM scan of all objects.

For Storage Nodes, the ILM tab provides details about ILM evaluation and background verification for erasure-coded objects.

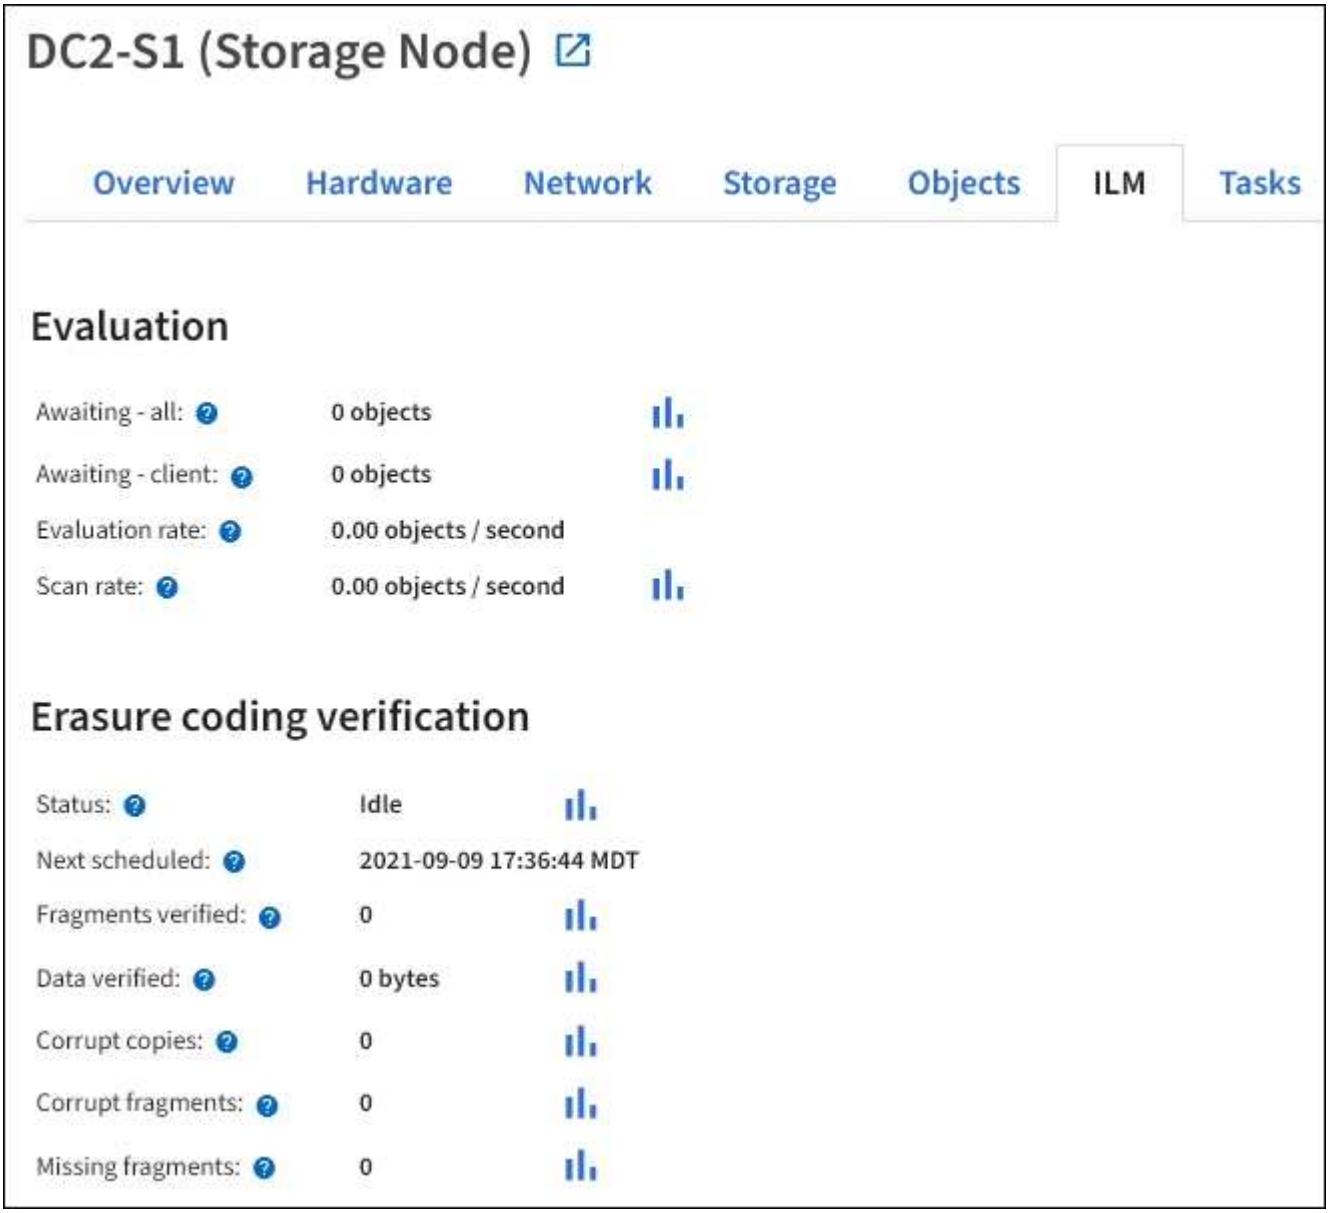

Related information

Monitor information lifecycle management

Administer StorageGRID

Use the Tasks tab

The Tasks tab is shown for all nodes. You can use this tab to rename or reboot a node or to put an appliance node into maintenance mode.

For the complete requirements and instructions for each option on this tab, see the following:

- Rename grid, sites, and nodes
- Reboot grid node
- Place appliance into maintenance mode

View the Load balancer tab

The Load Balancer tab includes performance and diagnostic graphs related to the operation of the Load Balancer service.

The Load Balancer tab is shown for Admin Nodes and Gateway Nodes, each site, and the entire grid. For each site, the Load Balancer tab provides an aggregate summary of the statistics for all nodes at that site. For the entire grid, the Load Balancer tab provides an aggregate summary of the statistics for all sites.

If there is no I/O being run through the Load Balancer service, or there is no load balancer configured, the graphs display "No data."

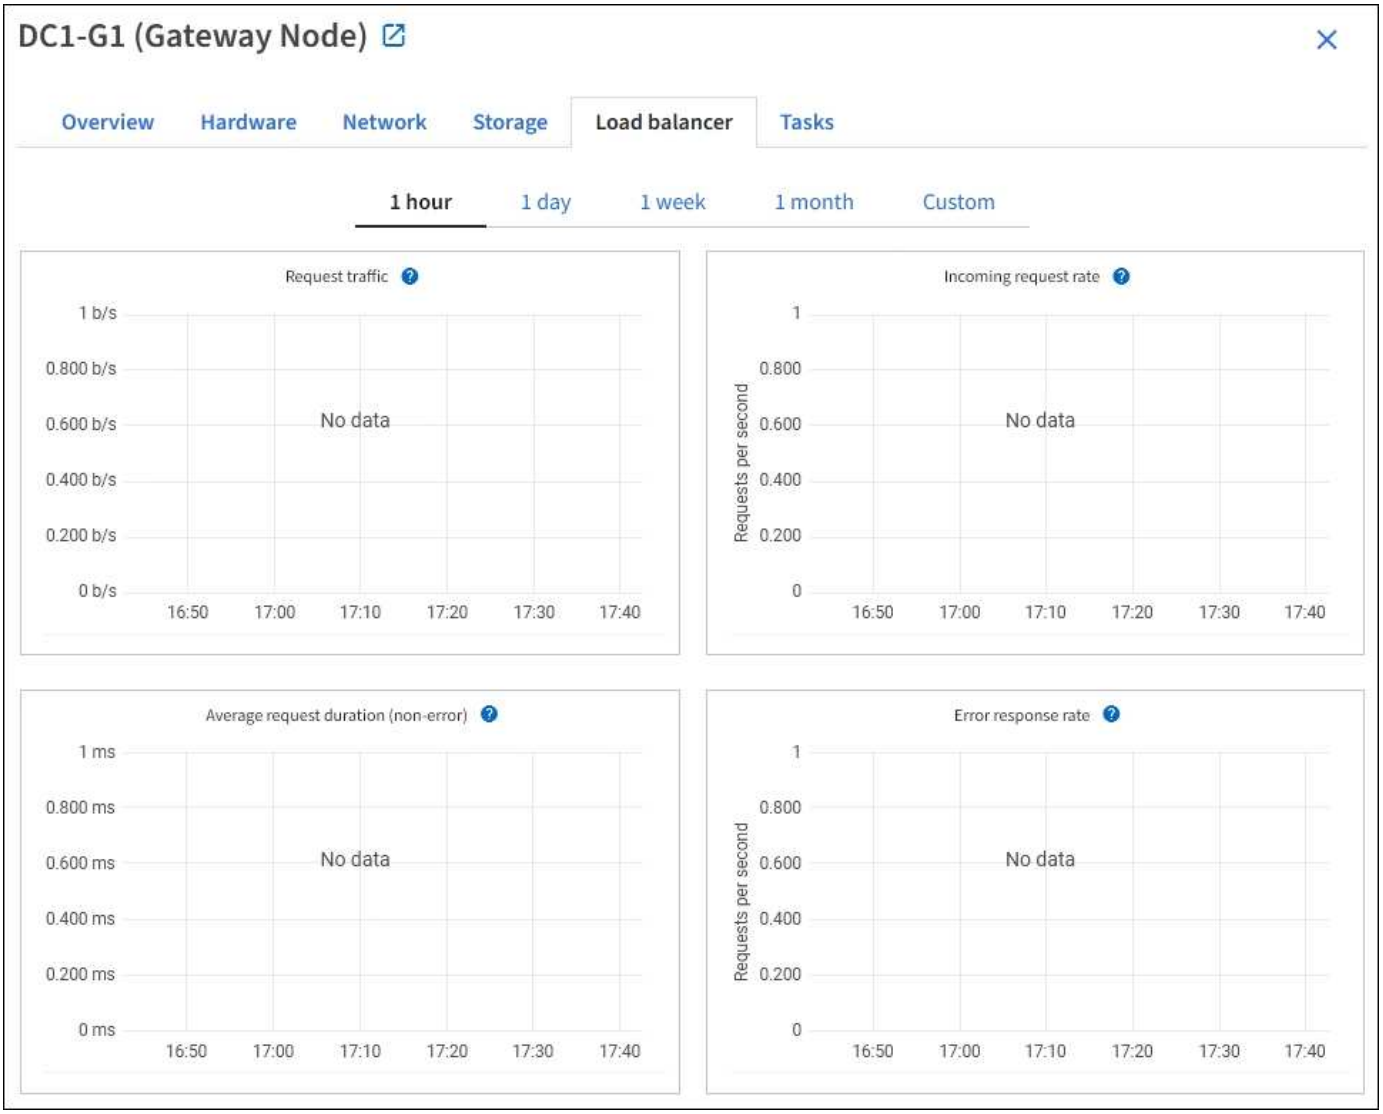

Request traffic

This graph provides a 3-minute moving average of the throughput of data transmitted between load balancer endpoints and the clients making the requests, in bits per second.

> This value is updated at the completion of each request. As a result, this value might differ from the real-time throughput at low request rates or for very long-lived requests. You can look at the Network tab to get a more realistic view of the current network behavior.

Incoming request rate

This graph provides a 3-minute moving average of the number of new requests per second, broken down by request type (GET, PUT, HEAD, and DELETE). This value is updated when the headers of a new request have been validated.

Average request duration (non-error)

This graph provides a 3-minute moving average of request durations, broken down by request type (GET, PUT, HEAD, and DELETE). Each request duration starts when a request header is parsed by the Load Balancer service and ends when the complete response body is returned to the client.

Error response rate

This graph provides a 3-minute moving average of the number of error responses returned to clients per second, broken down by the error response code.

Related information

Monitor load balancing operations

Administer StorageGRID

View the Platform services tab

The Platform services tab provides information about any S3 platform service operations at a site.

The Platform services tab is shown for each site. This tab provides information about S3 platform services, such as CloudMirror replication and the search integration service. Graphs on this tab display metrics such as the number of pending requests, request completion rate, and request failure rate.

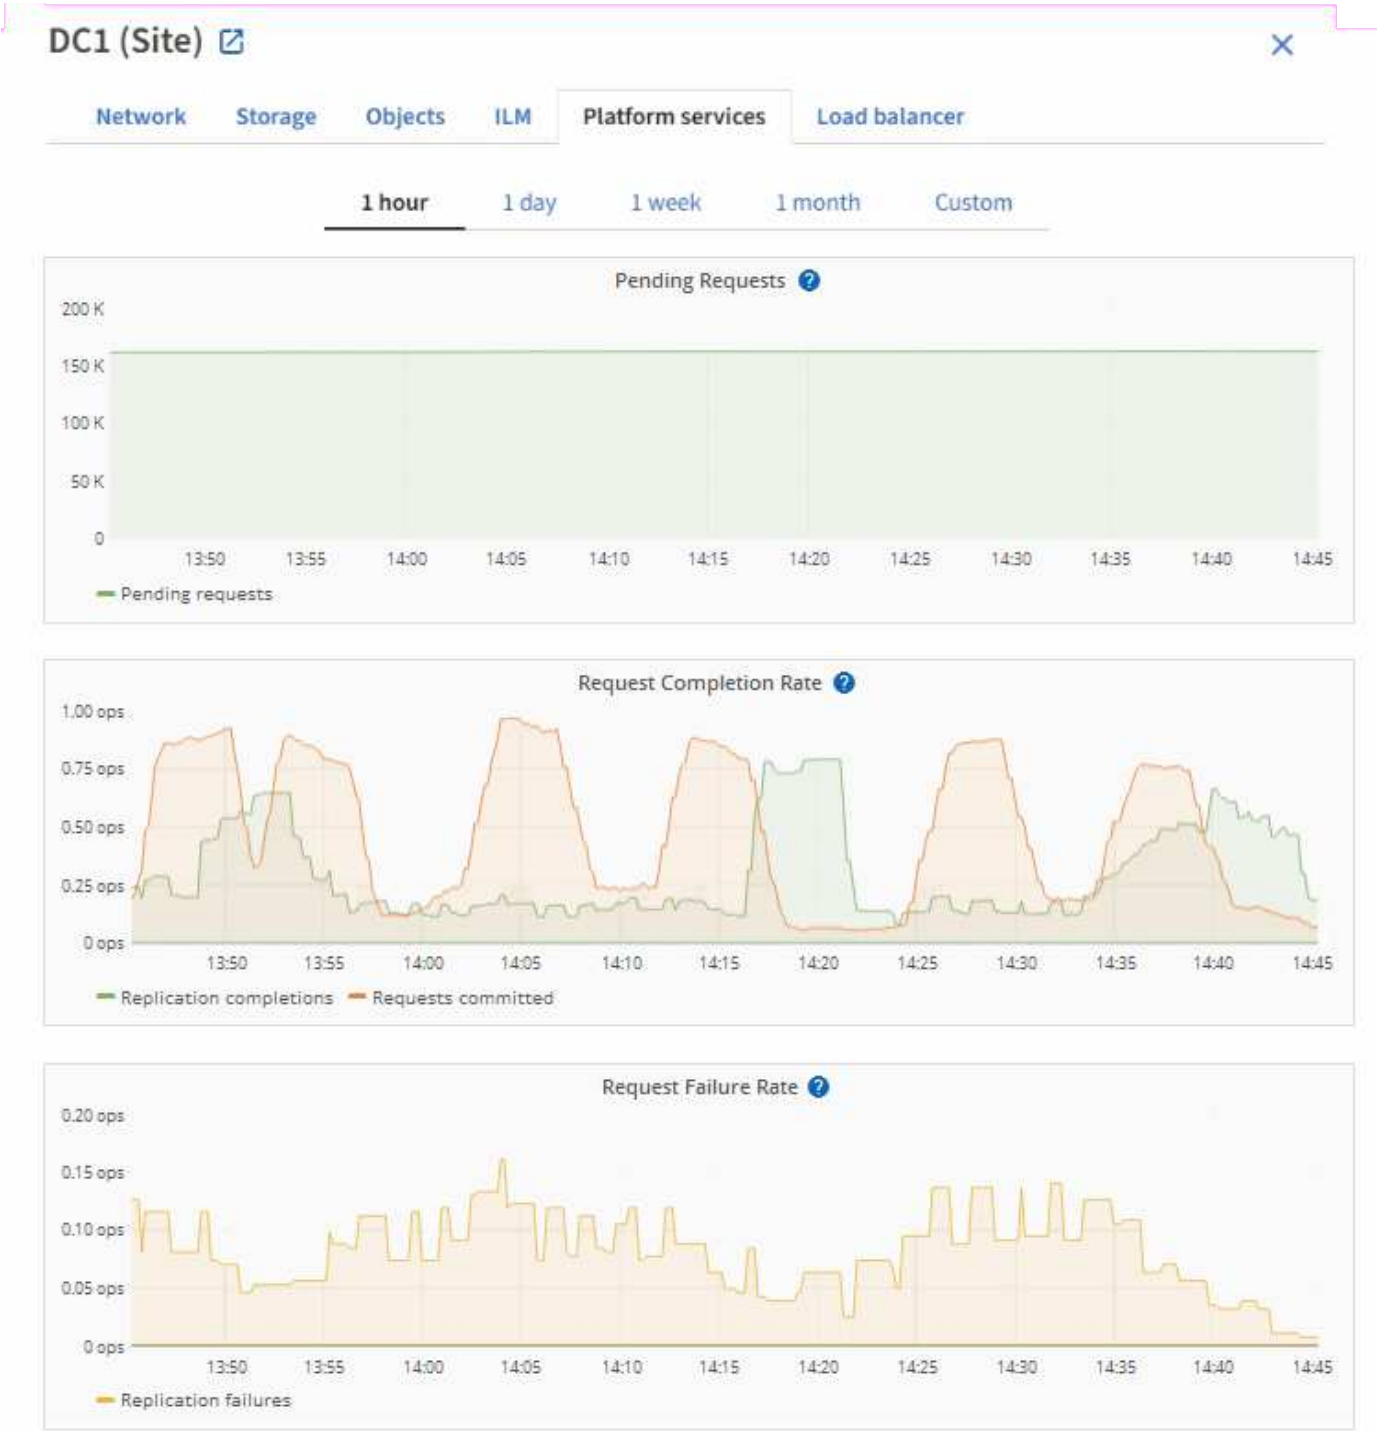

For more information about S3 platform services, including troubleshooting details, see the instructions for administering StorageGRID.

View the Manage drives tab (SGF6112 only)

The Manage drives tab enables you to access details and perform troubleshooting and maintenance tasks on the drives in the SGF6112 appliance.

> The Manage drives tab is shown only for SGF6112 storage appliance nodes.

Using the Manage drives tab, you can do the following:

- View a layout of the data storage drives in the appliance
- View a table that lists each drive location, type, status, firmware version, and serial number
- Perform troubleshooting and maintenance functions on each drive

To access the Manage drives tab, you must have the Storage appliance administrator or Root access permission.

For information about using the Manage drives tab, see Use the Manage drives tab.

View the SANtricity System Manager tab (E-Series only)

The SANtricity System Manager tab enables you to access SANtricity System Manager without having to configure or connect the management port of the storage appliance. You can use this tab to review hardware diagnostic and environmental information as well as issues related to the drives.

> The SANtricity System Manager tab is shown only for storage appliance nodes using E-Series hardware.

Using SANtricity System Manager, you can do the following:

- View performance data such as storage array level performance, I/O latency, storage controller CPU utilization, and throughput.
- Check hardware component status.
- Perform support functions including viewing diagnostic data, and configuring E-Series AutoSupport.

> To use SANtricity System Manager to configure a proxy for E-Series AutoSupport, see Send E-Series AutoSupport packages through StorageGRID.

To access SANtricity System Manager through Grid Manager, you must have the Storage appliance administrator or Root access permission.

> You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager.

> Accessing SANtricity System Manager from the Grid Manager is generally meant only to monitor appliance hardware and configure E-Series AutoSupport. Many features and operations within SANtricity System Manager such as upgrading firmware don't apply to monitoring your StorageGRID appliance. To avoid issues, always follow the hardware maintenance instructions for your appliance.

The tab displays the home page of SANtricity System Manager.

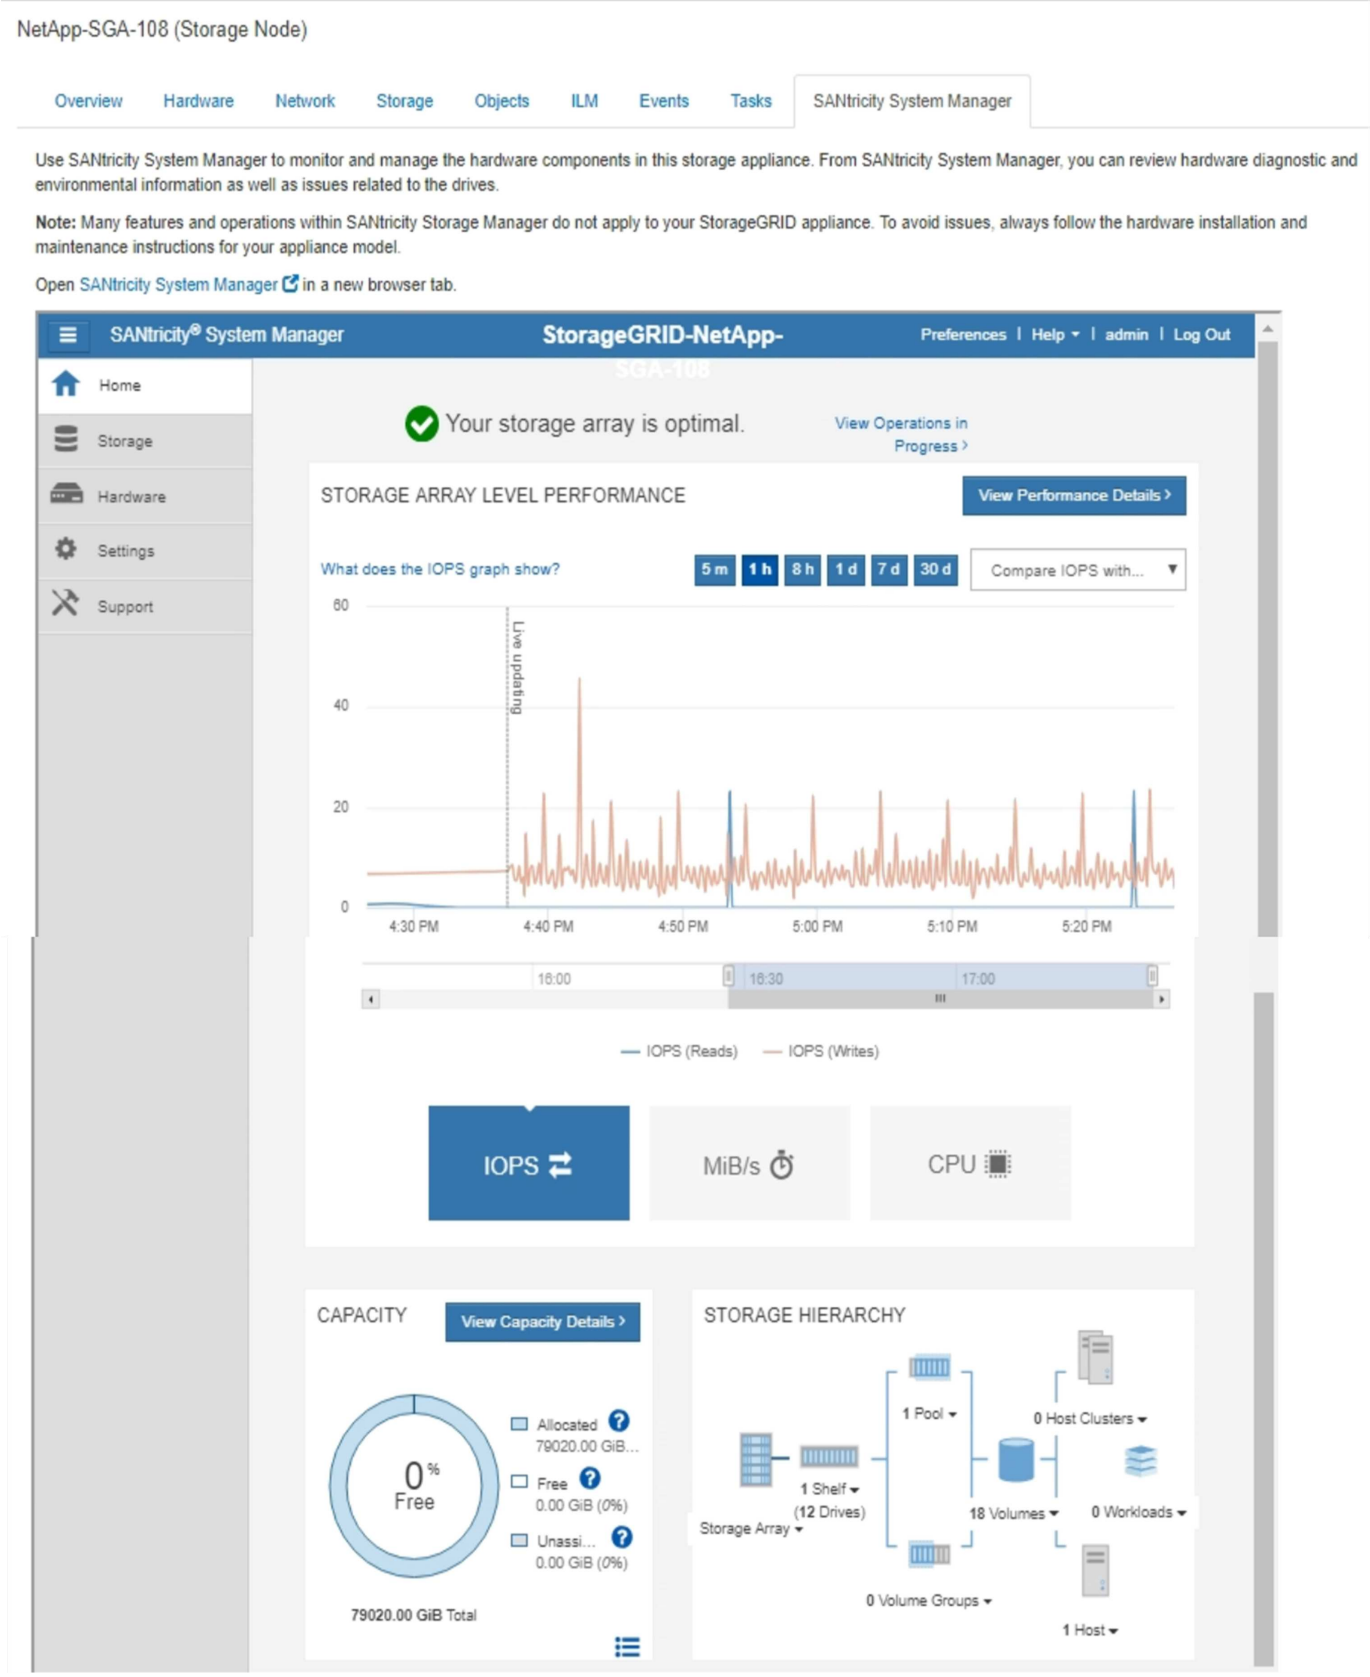

You can use the SANtricity System Manager link to open the SANtricity System Manager in a new browser window for easier viewing.

To see details for storage array level performance and capacity usage, position your cursor over each graph.

For more details on viewing the information accessible from the SANtricity System Manager tab, see NetApp E-Series and SANtricity documentation.

Information to monitor regularly

What and when to monitor

Even though the StorageGRID system can continue to operate when errors occur or parts of the grid are unavailable, you should monitor and address potential issues before they affect the grid's efficiency or availability.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About monitoring tasks

A busy system generates large amounts of information. The following list provides guidance about the most important information to monitor on an ongoing basis.

What to monitor	Frequency
System health status	Daily
Rate at which Storage Node object and metadata capacity is being consumed	Weekly
Information lifecycle management operations	Weekly
Networking and system resources	Weekly
Tenant activity	Weekly
S3 and Swift client operations	Weekly
Load balancing operations	After the initial configuration and after any configuration changes
Grid federation connections	Weekly
Capacity of the external archival storage system	Weekly

Monitor system health

Monitor the overall health of your StorageGRID system on a daily basis.

About this task

The StorageGRID system can continue to operate when parts of the grid are unavailable. Potential issues indicated by alerts or alarms (legacy system) aren't necessarily issues with system operations. Investigate

issues summarized on the Health status card of the Grid Manager Dashboard.

To be notified of alerts as soon as they are triggered, you can set up email notifications for alerts or configure SNMP traps.

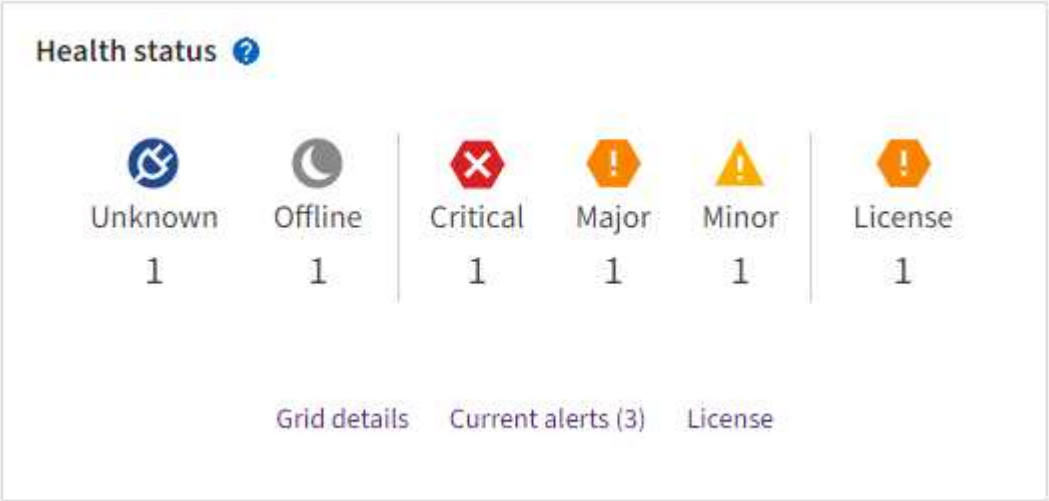

When issues exist, links appear that allow you to view additional details:

Link	Appears when…
Grid details	Any nodes are disconnected (connection state Unknown or Administratively Down).
Current alerts (Critical, Major, Minor)	Alerts are currently active.
Recently resolved alerts	Alerts triggered in the past week are now resolved.
License	There is an issue with the software license for this StorageGRID system. You can update license information as needed.

Monitor node connection states

If one or more nodes are disconnected from the grid, critical StorageGRID operations might be affected. Monitor node connection states and address any issues promptly.

Icon	Description	Action required
	Not connected - Unknown For an unknown reason, a node is disconnected or services on the node are unexpectedly down. For example, a service on the node might be stopped, or the node might have lost its network connection because of a power failure or unexpected outage. The **Unable to communicate with node** alert might also be triggered. Other alerts might also be active.	Requires immediate attention. Select each alert and follow the recommended actions. For example, you might need to restart a service that has stopped or restart the host for the node. **Note**: A node might appear as Unknown during managed shutdown operations. You can ignore the Unknown state in these cases.
	Not connected - Administratively down For an expected reason, node is not connected to grid. 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. Based on the underlying issue, these nodes often go back online with no intervention.	Determine if any alerts are affecting this node. If one or more alerts are active, select each alert and follow the recommended actions.
	Connected The node is connected to the grid.	No action required.

View current and resolved alerts

Current alerts: When an alert is triggered, an alert icon is displayed on the dashboard. An alert icon is also displayed for the node on the Nodes page. If alert email notifications are configured, an email notification will also be sent, unless the alert has been silenced.

Resolved alerts: You can search and view a history of alerts that have been resolved.

Optionally, you have watched the video: Video: Alerts overview for StorageGRID 11.8

The following table describes the information shown in the Grid Manager for current and resolved alerts.

Column header	Description
Name or title	The name of the alert and its description.
Severity	The severity of the alert. For current alerts, if multiple alerts are grouped the title row shows how many instances of that alert are occurring at each severity. **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 don't clear on their own to ensure they don't result in a more serious problem.
Time triggered	**Current alerts**: The date and time the alert was triggered in your local time and in UTC. If multiple alerts are grouped, the title row shows times for the most recent instance of the alert (*newest*) and the oldest instance of the alert (*oldest*). **Resolved alerts**: How long ago the alert was triggered.
Site/Node	The name of the site and node where the alert is occurring or has occurred.
Status	Whether the alert is active, silenced, or resolved. If multiple alerts are grouped and **All alerts** is selected in the drop-down, the title row shows how many instances of that alert are active and how many instances have been silenced.
Time resolved (resolved alerts only)	How long ago the alert was resolved.

Column header	Description
Current values or *data values*	The value of the metric that caused the alert to be triggered. For some alerts, additional values are shown to help you understand and investigate the alert. For example, the values shown for a **Low object data storage** alert include the percentage of disk space used, the total amount of disk space, and the amount of disk space used. **Note:** If multiple current alerts are grouped, current values aren't shown in the title row.
Triggered values (resolved alerts only)	The value of the metric that caused the alert to be triggered. For some alerts, additional values are shown to help you understand and investigate the alert. For example, the values shown for a **Low object data storage** alert include the percentage of disk space used, the total amount of disk space, and the amount of disk space used.

Steps

1. Select the **Current alerts** or **Resolved alerts** link to view a list of alerts in those categories. You can also view the details for an alert by selecting **Nodes** > ***node*** > **Overview** and then selecting the alert from the Alerts table.

   By default, current alerts are shown as follows:

   - The most recently triggered alerts are shown first.
   - Multiple alerts of the same type are shown as a group.
   - Alerts that have been silenced aren't shown.
   - For a specific alert on a specific node, if the thresholds are reached for more than one severity, only the most severe alert is shown. That is, if alert thresholds are reached for the minor, major, and critical severities, only the critical alert is shown.

     The Current alerts page is refreshed every two minutes.

2. To expand groups of alerts, select the down caret ⌄. To collapse individual alerts in a group, select the up caret ⌃, or select the group's name.
3. To display individual alerts instead of groups of alerts, clear the **Group alerts** checkbox.
4. To sort current alerts or alert groups, select the up/down arrows ⇅ in each column header.
   - When **Group alerts** is selected, both the alert groups and the individual alerts within each group are sorted. For example, you might want to sort the alerts in a group by **Time triggered** to find the most recent instance of a specific alert.
   - When **Group alerts** is cleared, the entire list of alerts is sorted. For example, you might want to sort all alerts by **Node/Site** to see all alerts affecting a specific node.
5. To filter current alerts by status (**All alerts**, **Active**, or **Silenced**, use the drop-down menu at the top of the table.

   See Silence alert notifications.

6. To sort resolved alerts:
   - Select a time period from the **When triggered** drop-down menu.

- Select one or more severities from the **Severity** drop-down menu.
- Select one or more default or custom alert rules from the **Alert rule** drop-down menu to filter on resolved alerts related to a specific alert rule.
- Select one or more nodes from the **Node** drop-down menu to filter on resolved alerts related to a specific node.

7. To view details for a specific alert, select the alert. A dialog box provides details and recommended actions for the alert you selected.
8. (Optional) For a specific alert, select silence this alert to silence the alert rule that caused this alert to be triggered.

   You must have the Manage alerts or Root access permission to silence an alert rule.

   > Be careful when deciding to silence an alert rule. If an alert rule is silenced, you might not detect an underlying problem until it prevents a critical operation from completing.

9. To view the current conditions for the alert rule:
   a. From the alert details, select **View conditions**.

      A pop-up appears, listing the Prometheus expression for each defined severity.

   b. To close the pop-up, click anywhere outside of the pop-up.
10. Optionally, select **Edit rule** to edit the alert rule that caused this alert to be triggered.

    You must have the Manage alerts or Root access permission to edit an alert rule.

    > Be careful when deciding to edit an alert rule. If you change trigger values, you might not detect an underlying problem until it prevents a critical operation from completing.

11. To close the alert details, select **Close**.

Monitor storage capacity

Monitor the total usable space available to ensure that the StorageGRID system does not run out of storage space for objects or for object metadata.

StorageGRID stores object data and object metadata separately, and reserves a specific amount of space for a distributed Cassandra database that contains object metadata. Monitor the total amount of space consumed for objects and for object metadata, as well as trends in the amount of space consumed for each. This will enable you to plan ahead for the addition of nodes and avoid any service outages.

You can view storage capacity information for the entire grid, for each site, and for each Storage Node in your StorageGRID system.

Monitor storage capacity for the entire grid

Monitor the overall storage capacity for your grid to ensure that adequate free space remains for object data and object metadata. Understanding how storage capacity changes over time can help you plan to add Storage Nodes or storage volumes before the grid's usable storage capacity is consumed.

The Grid Manager dashboard lets you quickly assess how much storage is available for the entire grid and for each data center. The Nodes page provides more detailed values for object data and object metadata.

Steps

1. Assess how much storage is available for the entire grid and for each data center.
   a. Select **Dashboard > Overview**.
   b. Note the values on the Data space usage breakdown and the Metadata allowed space usage breakdown cards. Each card lists a percentage of storage usage, the capacity of used space, and the total space available or allowed by site.

      The summary does not include archival media.

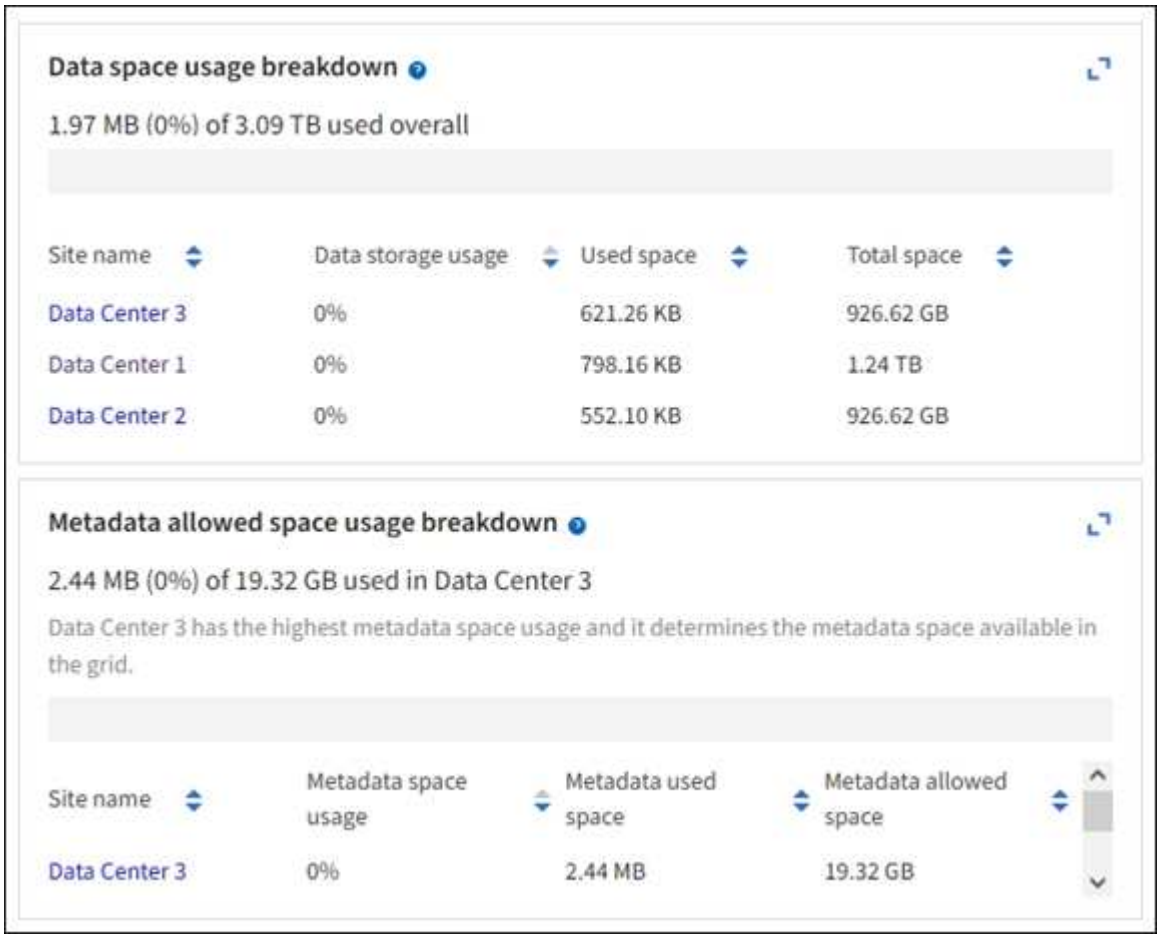

   c. Note the chart on the Storage over time card. Use the time period drop-down to help you determine how quickly storage is consumed.

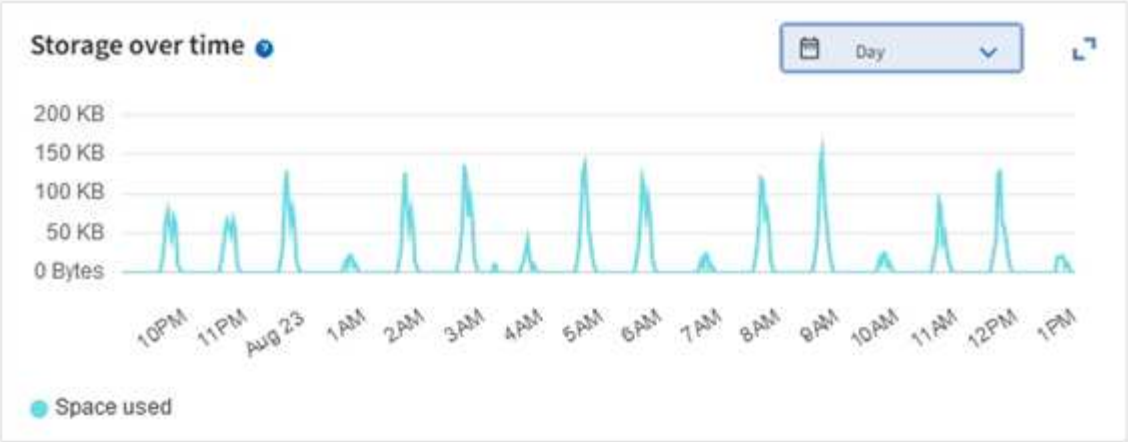

2. Use the Nodes page for additional details on how much storage has been used and how much storage remains available on the grid for object data and object metadata.
   a. Select **NODES**.
   b. Select ***grid*** > **Storage**.

   c. Position your cursor over the **Storage used - object data** and the **Storage used - object metadata** charts to see how much object storage and object metadata storage is available for the entire grid, and how much has been used over time.

> The total values for a site or the grid don't include nodes that have not reported metrics for at least five minutes, such as offline nodes.

3. Plan to perform an expansion to add Storage Nodes or storage volumes before the grid's usable storage capacity is consumed.

   When planning the timing of an expansion, consider how long it will take to procure and install additional storage.

> If your ILM policy uses erasure coding, you might prefer to expand when existing Storage Nodes are approximately 70% full to reduce the number of nodes that must be added.

For more information about planning a storage expansion, see the instructions for expanding StorageGRID.

Monitor storage capacity for each Storage Node

Monitor the total usable space for each Storage Node to ensure that the node has enough space for new object data.

About this task

Usable space is the amount of storage space available to store objects. The total usable space for a Storage Node is calculated by adding together the available space on all object stores within the node.

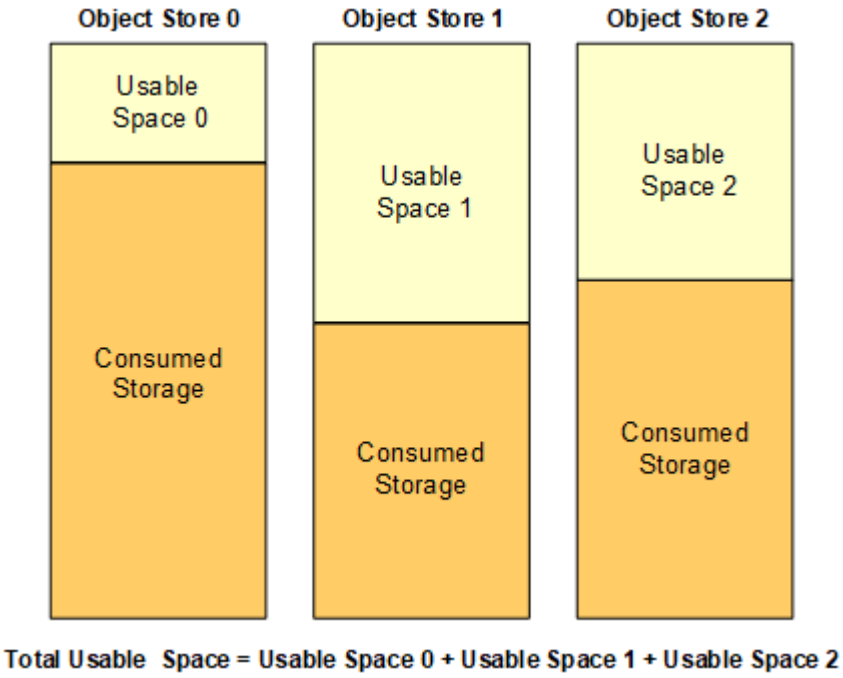

Steps

1. Select **NODES** > ***Storage Node*** > **Storage**.

   The graphs and tables for the node appear.

2. Position your cursor over the Storage used - object data graph.

   The following values are shown:

   - **Used (%)**: The percentage of the Total usable space that has been used for object data.
   - **Used**: The amount of the Total usable space that has been used for object data.
   - **Replicated data**: An estimate of the amount of replicated object data on this node, site, or grid.
   - **Erasure-coded data**: An estimate of the amount of erasure-coded object data on this node, site, or grid.
   - **Total**: The total amount of usable space on this node, site, or grid.

The Used value is the `storagegrid_storage_utilization_data_bytes` metric.

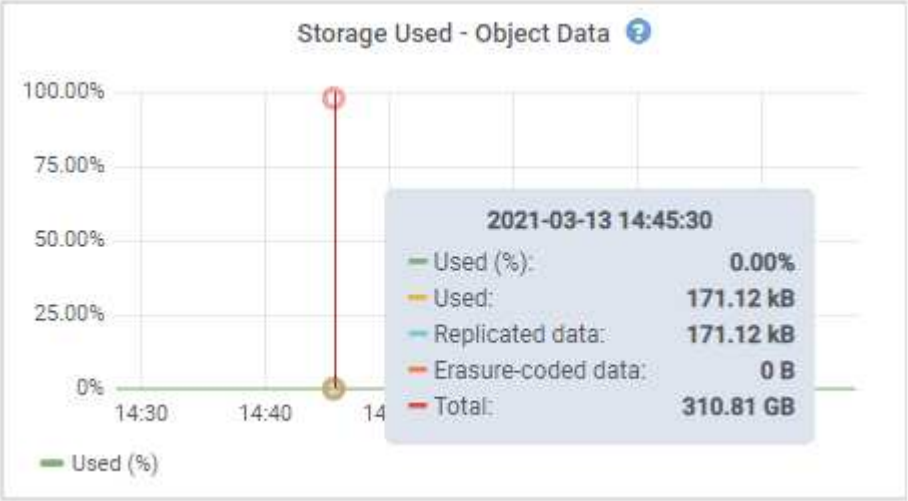

3. Review the Available values in the Volumes and Object stores tables, below the graphs.

To view graphs of these values, click the chart icons in the Available columns.

Disk devices

Name	World Wide Name	I/O load	Read rate	Write rate
croot(8:1,sda1)	N/A	0.04%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.67%	0 bytes/s	50 KB/s
sdc(8:16,sdb)	N/A	0.03%	0 bytes/s	4 KB/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s

Volumes

Mount point	Device	Status	Size	Available	Write cache status
/	croot	Online	21.00 GB	14.75 GB	Unknown
/var/local	cvloc	Online	85.86 GB	84.05 GB	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.17 GB	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/2	sde	Online	107.32 GB	107.18 GB	Enabled

Object stores

ID	Size	Available	Replicated data	EC data	Object data (%)	Health
0000	107.32 GB	96.44 GB	124.60 KB	0 bytes	0.00%	No Errors
0001	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0002	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors

4. Monitor the values over time to estimate the rate at which usable storage space is being consumed.
5. To maintain normal system operations, add Storage Nodes, add storage volumes, or archive object data before usable space is consumed.

   When planning the timing of an expansion, consider how long it will take to procure and install additional storage.

   > If your ILM policy uses erasure coding, you might prefer to expand when existing Storage Nodes are approximately 70% full to reduce the number of nodes that must be added.

   For more information about planning a storage expansion, see the instructions for expanding StorageGRID.

The Low object data storage alert is triggered when insufficient space remains for storing object data on a Storage Node.

Monitor object metadata capacity for each Storage Node

Monitor the metadata usage for each Storage Node to ensure that adequate space remains available for essential database operations. You must add new Storage Nodes at each site before object metadata exceeds 100% of the allowed metadata space.

About this task

StorageGRID maintains three copies of object metadata at each site to provide redundancy and to protect object metadata from loss. The three copies are evenly distributed across all Storage Nodes at each site using the space reserved for metadata on storage volume 0 of each Storage Node.

In some cases, the grid's object metadata capacity might be consumed faster than its object storage capacity. For example, if you typically ingest large numbers of small objects, you might need to add Storage Nodes to increase metadata capacity even though sufficient object storage capacity remains.

Some of the factors that can increase metadata usage include the size and quantity of user metadata and tags, the total number of parts in a multipart upload, and the frequency of changes to ILM storage locations.

Steps

1. Select **NODES** > ***Storage Node*** > **Storage**.
2. Position your cursor over the Storage used - object metadata graph to see the values for a specific time.

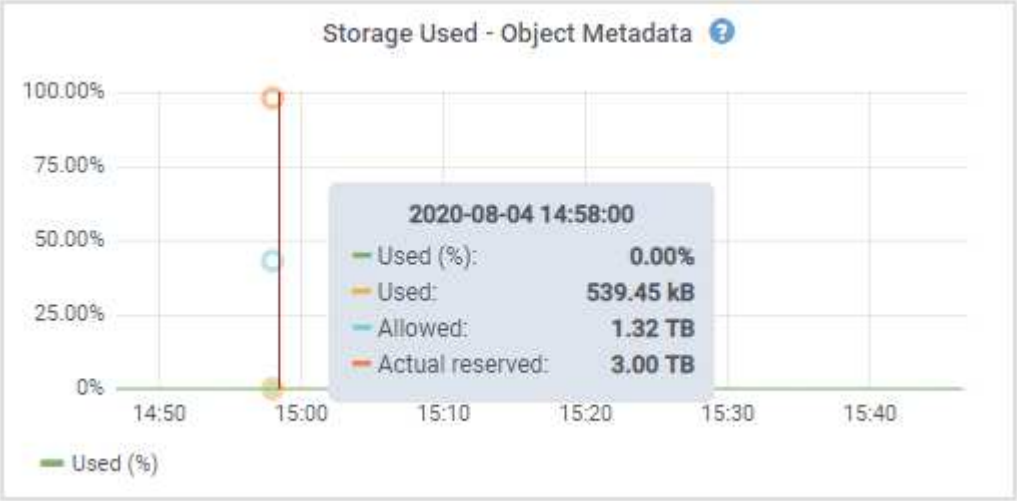

Used (%)

The percentage of the allowed metadata space that has been used on this Storage Node.

Prometheus metrics: `storagegrid_storage_utilization_metadata_bytes` and `storagegrid_storage_utilization_metadata_allowed_bytes`

Used

The bytes of the allowed metadata space that have been used on this Storage Node.

Prometheus metric: `storagegrid_storage_utilization_metadata_bytes`

Allowed

The space allowed for object metadata on this Storage Node. To learn how this value is determine for each Storage Node, see the full description of Allowed metadata space.

Prometheus metric: `storagegrid_storage_utilization_metadata_allowed_bytes`

Actual reserved

The actual space reserved for metadata on this Storage Node. Includes the allowed space and the required space for essential metadata operations. To learn how this value is calculated for each Storage Node, see the full description of Actual reserved space for metadata.

Prometheus metric will be added in a future release.

> The total values for a site or the grid don't include nodes that have not reported metrics for at least five minutes, such as offline nodes.

3. If the **Used (%)** value is 70% or higher, expand your StorageGRID system by adding Storage Nodes to each site.

> The **Low metadata storage** alert is triggered when the **Used (%)** value reaches certain thresholds. Undesirable results can occur if object metadata uses more than 100% of the allowed space.

When you add the new nodes, the system automatically rebalances object metadata across all Storage Nodes within the site. See the instructions for expanding a StorageGRID system.

Monitor space usage forecasts

Monitor space usage forecasts for user data and metadata to estimate when you will need to expand a grid.

If you notice that the rate of consumption changes over time, select a shorter range from the **Averaged over** pull-down to reflect only the most recent ingest patterns. If you notice seasonal patterns, select a longer range.

If you have a new StorageGRID installation, allow data and metadata to accumulate before evaluating the space usage forecasts.

Steps

1. On the dashboard, select **Storage**.
2. View the dashboard cards, Forecast of data usage by storage pool and Forecast of metadata usage by site.
3. Use these values to estimate when you will need to add new Storage Nodes for data and metadata storage.

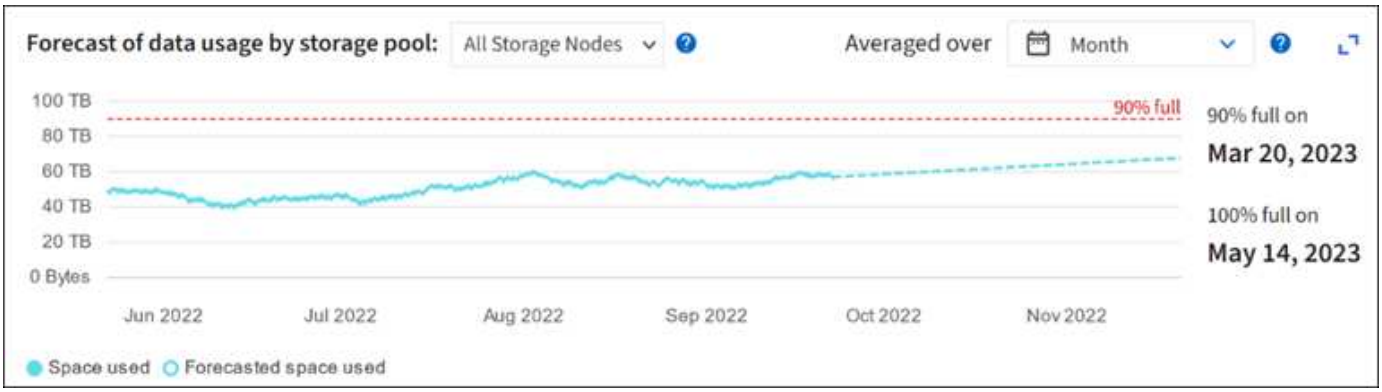

Monitor information lifecycle management

The information lifecycle management (ILM) system provides data management for all objects stored on the grid. You must monitor ILM operations to understand if the grid can handle the current load, or if more resources are needed.

About this task

The StorageGRID system manages objects by applying the active ILM policies. The ILM policies and associated ILM rules determine how many copies are made, the type of copies that are created, where copies are placed, and the length of time each copy is retained.

Object ingest and other object-related activities can exceed the rate at which StorageGRID can evaluate ILM, causing the system to queue objects whose ILM placement instructions can't be fulfilled in near real time. You should monitor whether StorageGRID is keeping up with client actions.

Use Grid Manager dashboard tab

Steps

Use the ILM tab on the Grid Manager dashboard to monitor ILM operations:

1. Sign in to the Grid Manager.
2. From the dashboard, select the ILM tab and note the values on the ILM queue (Objects) card and ILM evaluation rate card.

   Temporary spikes in the ILM queue (Objects) card on the dashboard are to be expected. But if the queue continues to increase and never declines, the grid needs more resources to operate efficiently: either more Storage Nodes, or, if the ILM policy places objects in remote locations, more network bandwidth.

Use the NODES page

Steps

Additionally, investigate ILM queues using the **NODES** page:

> The charts on the **NODES** page will be replaced with the corresponding dashboard cards in a future StorageGRID release.

1. Select **NODES**.
2. Select ***grid name*** > **ILM**.
3. Position your cursor over the ILM queue graph to see the value of following attributes at a given point in

time:

- **Objects queued (from client operations)**: The total number of objects awaiting ILM evaluation because of client operations (for example, ingest).
- **Objects queued (from all operations)**: The total number of objects awaiting ILM evaluation.
- **Scan rate (objects/sec)**: The rate at which objects in the grid are scanned and queued for ILM.
- **Evaluation rate (objects/sec)**: The current rate at which objects are being evaluated against the ILM policy in the grid.

4. In the ILM Queue section, look at the following attributes.

> The ILM queue section is included for the grid only. This information is not shown on the ILM tab for a site or Storage Node.

- **Scan period - estimated**: The estimated time to complete a full ILM scan of all objects.

> A full scan does not guarantee that ILM has been applied to all objects.

- **Repairs attempted**: The total number of object repair operations for replicated data that have been attempted. This count increments each time a Storage Node tries to repair a high-risk object. High-risk ILM repairs are prioritized if the grid becomes busy.

> The same object repair might increment again if replication failed after the repair.

These attributes can be useful when you are monitoring the progress of Storage Node volume recovery. If the number of Repairs attempted has stopped increasing and a full scan has been completed, the repair has probably completed.

Monitor networking and system resources

The integrity and bandwidth of the network between nodes and sites, and the resource usage by individual grid nodes, are critical to efficient operations.

Monitor network connections and performance

Network connectivity and bandwidth are especially important if your information lifecycle management (ILM) policy copies replicated objects between sites or stores erasure-coded objects using a scheme that provides site-loss protection. If the network between sites is not available, network latency is too high, or network bandwidth is insufficient, some ILM rules might not be able to place objects where expected. This can lead to ingest failures (when the Strict ingest option is selected for ILM rules), or to poor ingest performance and ILM backlogs.

Use the Grid Manager to monitor connectivity and network performance, so you can address any issues promptly.

Additionally, consider creating network traffic classification policies so that you can monitor traffic related to specific tenants, buckets, subnets, or load balancer endpoints. You can set traffic limiting policies as needed.

Steps

1. Select **NODES**.

   The Nodes page appears. Each node in the grid is listed in table format.

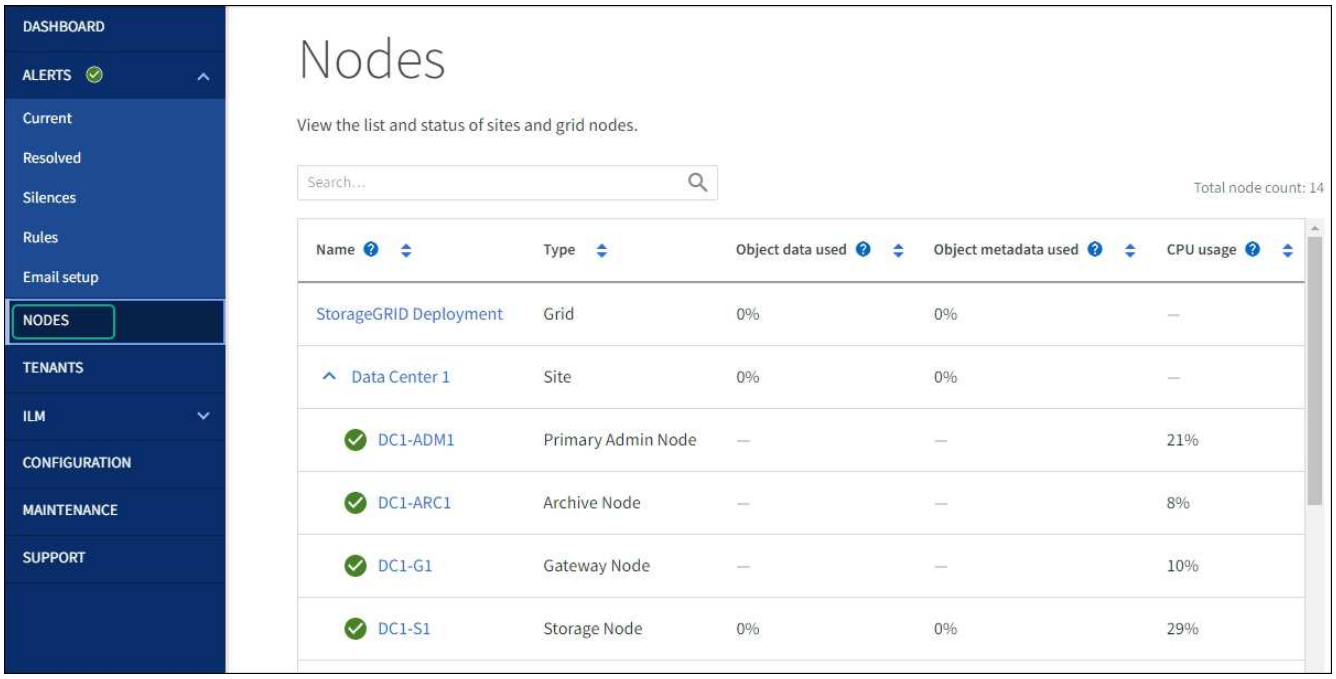

2. Select the grid name, a specific data center site, or a grid node, and then select the **Network** tab.

   The Network Traffic graph provides a summary of overall network traffic for the grid as a whole, the data center site, or for the node.

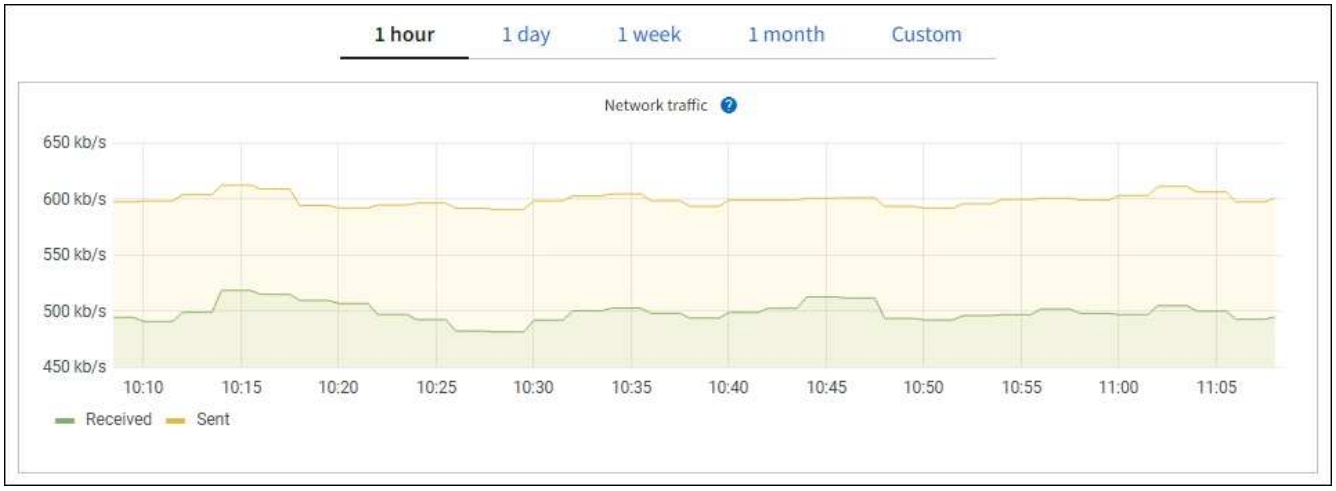

   a. If you selected a grid node, scroll down to review the **Network Interfaces** section of the page.

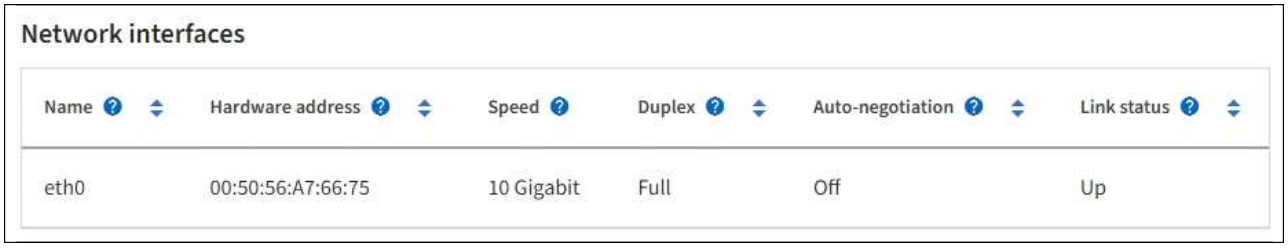

   b. For grid nodes, scroll down to review the **Network Communication** section of the page.

      The Receive and Transmit tables show how many bytes and packets have been received and sent across each network as well as other receive and transmission metrics.

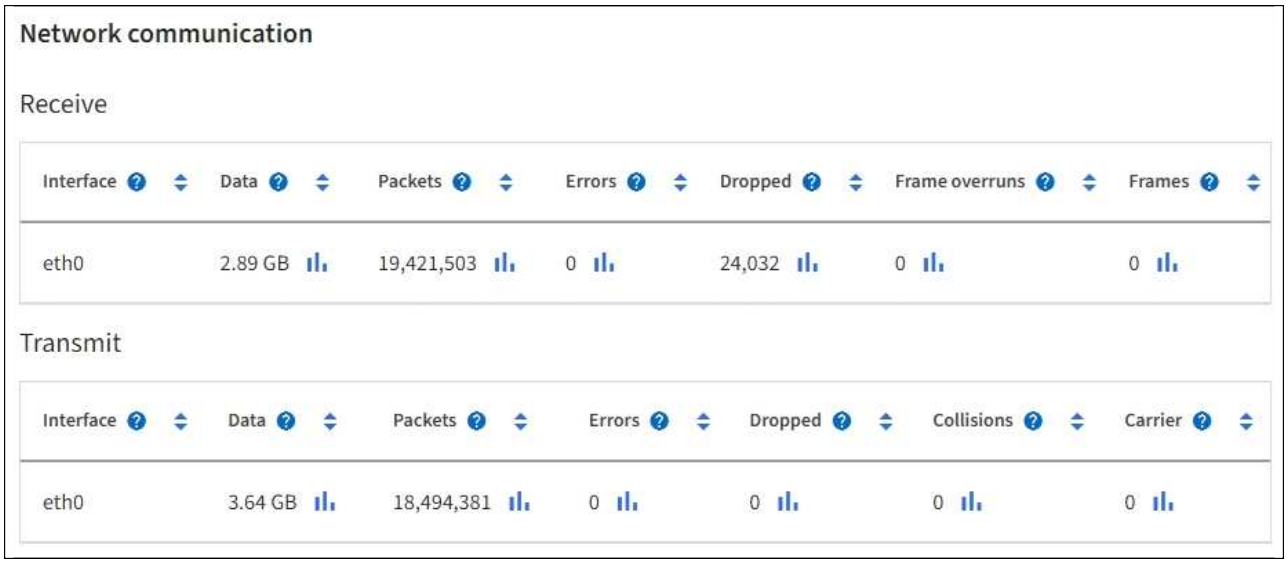

Network communication

Receive

Interface	Data	Packets	Errors	Dropped	Frame overruns	Frames
eth0	2.89 GB	19,421,503	0	24,032	0	0

Transmit

Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	3.64 GB	18,494,381	0	0	0	0

3. Use the metrics associated with your traffic classification policies to monitor network traffic.
   a. Select **CONFIGURATION** > **Network** > **Traffic classification**.

      The Traffic Classification Policies page appears, and the existing policies are listed in the table.

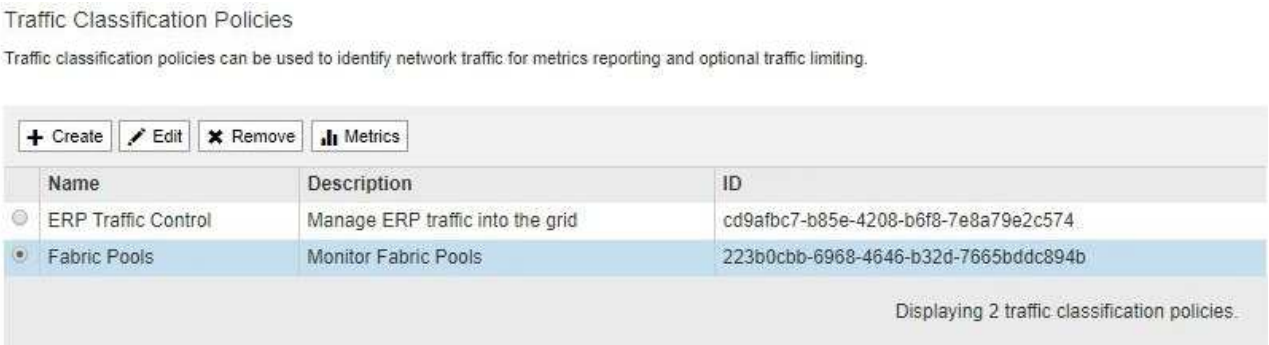

Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

+ Create | Edit | Remove | Metrics

Name	Description	ID
ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b

Displaying 2 traffic classification policies.

   b. To view graphs that show the networking metrics associated with a policy, select the radio button to the left of the policy, and then click **Metrics**.
   c. Review the graphs to understand the network traffic associated with the policy.

      If a traffic classification policy is designed to limit network traffic, analyze how often traffic is limited and decide if the policy continues to meet your needs. From time to time, adjust each traffic classification policy as needed.

Related information

View the Network tab

Monitor node connection states

Monitor node-level resources

Monitor individual grid nodes to check their resource usage levels. If nodes are consistently overloaded, more nodes might be required for efficient operations.

Steps

1. From the **NODES** page, select the node.

2. Select the **Hardware** tab to display graphs of CPU Utilization and Memory Usage.

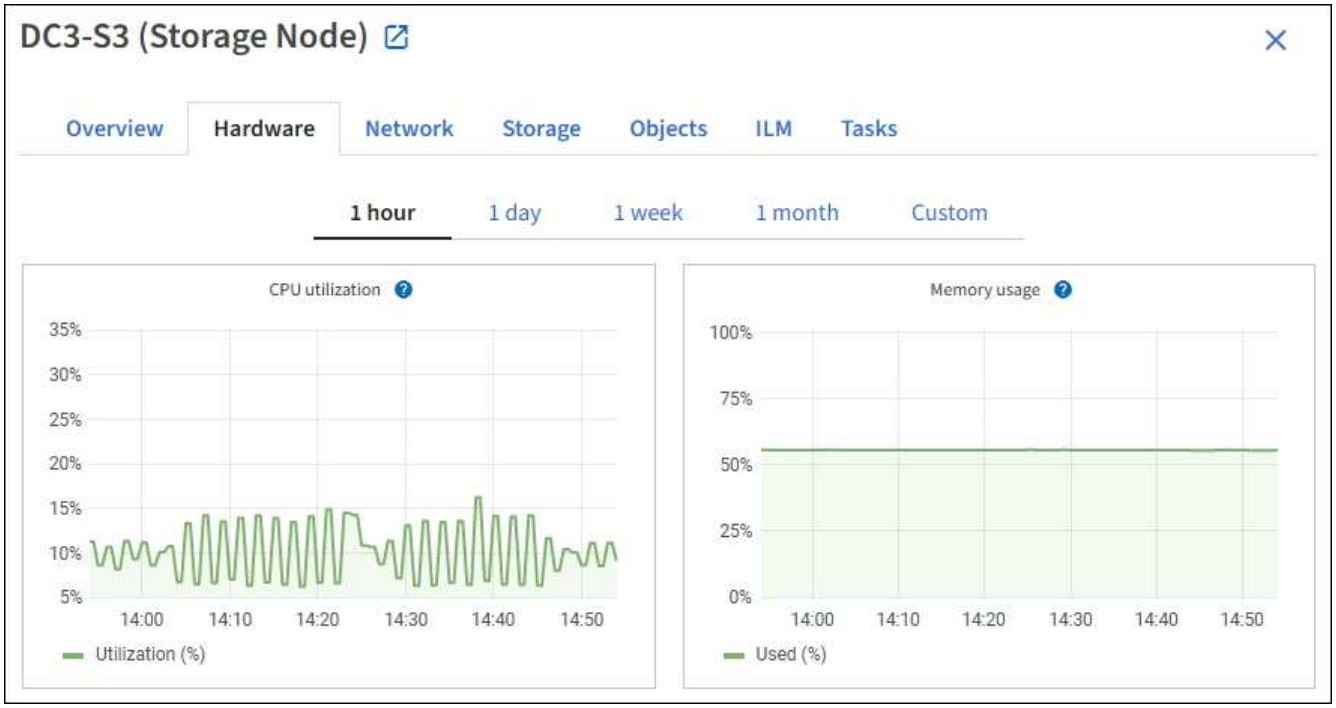

3. To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.
4. If the node is hosted on a storage appliance or a services appliance, scroll down to view the tables of components. The status of all components should be "Nominal." Investigate components that have any other status.

Related information

View information about appliance Storage Nodes

View information about appliance Admin Nodes and Gateway Nodes

Monitor tenant activity

All S3 and Swift client activity is associated with StorageGRID tenant accounts. You can use the Grid Manager to monitor the storage usage or network traffic for all tenants or a specific tenant. You can use the audit log or Grafana dashboards to gather more detailed information about how tenants are using StorageGRID.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access or Tenant accounts permission.

View all tenants

The Tenants page shows basic information for all current tenant accounts.

Steps

1. Select **TENANTS**.

2. Review the information shown on the Tenant pages.

   The Logical space used, Quota utilization, Quota, and Object count are listed for each tenant. If a quota is not set for a tenant, the Quota utilization and Quota fields contain a dash (—).

   > The space used values are estimates. These estimates are affected by the timing of ingests, network connectivity, and node status.

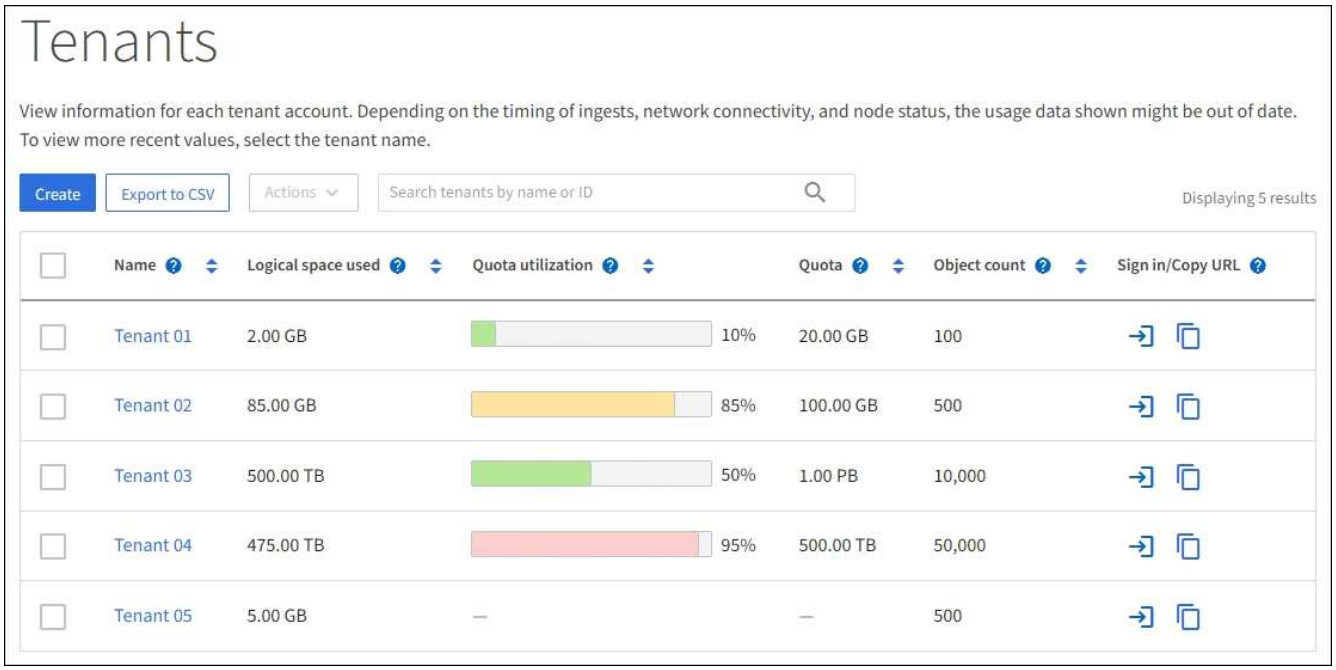

3. Optionally, sign in to a tenant account by selecting the sign-in link in the **Sign in/Copy URL** column.
4. Optionally, copy the URL for a tenant's sign-in page by selecting the copy URL link in the **Sign in/Copy URL** column.
5. Optionally, select **Export to CSV** to view and export a `.csv` file containing the usage values for all tenants.

   You are prompted to open or save the `.csv` file.

   The contents of the `.csv` file look like the following example:

Tenant ID	Display Name	Space Used (Bytes)	Quota utilization (%)	Quota (Bytes)	Object Count	Protocol
12659822378459233654	Tenant 01	2000000000	10	20000000000	100	S3
99658234112547853685	Tenant 02	85000000000	85	110000000	500	S3
03521145586975586321	Tenant 03	60500000000	50	150000	10000	S3
44251365987569885632	Tenant 04	4750000000	95	140000000	50000	S3
36521587546689565123	Tenant 05	5000000000	Infinity		500	S3

   You can open the `.csv` file in a spreadsheet application or use it in automation.

6. If no objects are listed, optionally, select **Actions** > **Delete** to remove one or more tenants. See Delete tenant account.

   You can't remove a tenant account if the account includes any buckets or containers.

View a specific tenant

You can view details for a specific tenant.

Steps

1. Select the tenant name from the Tenants page.

   The tenant details page appears.

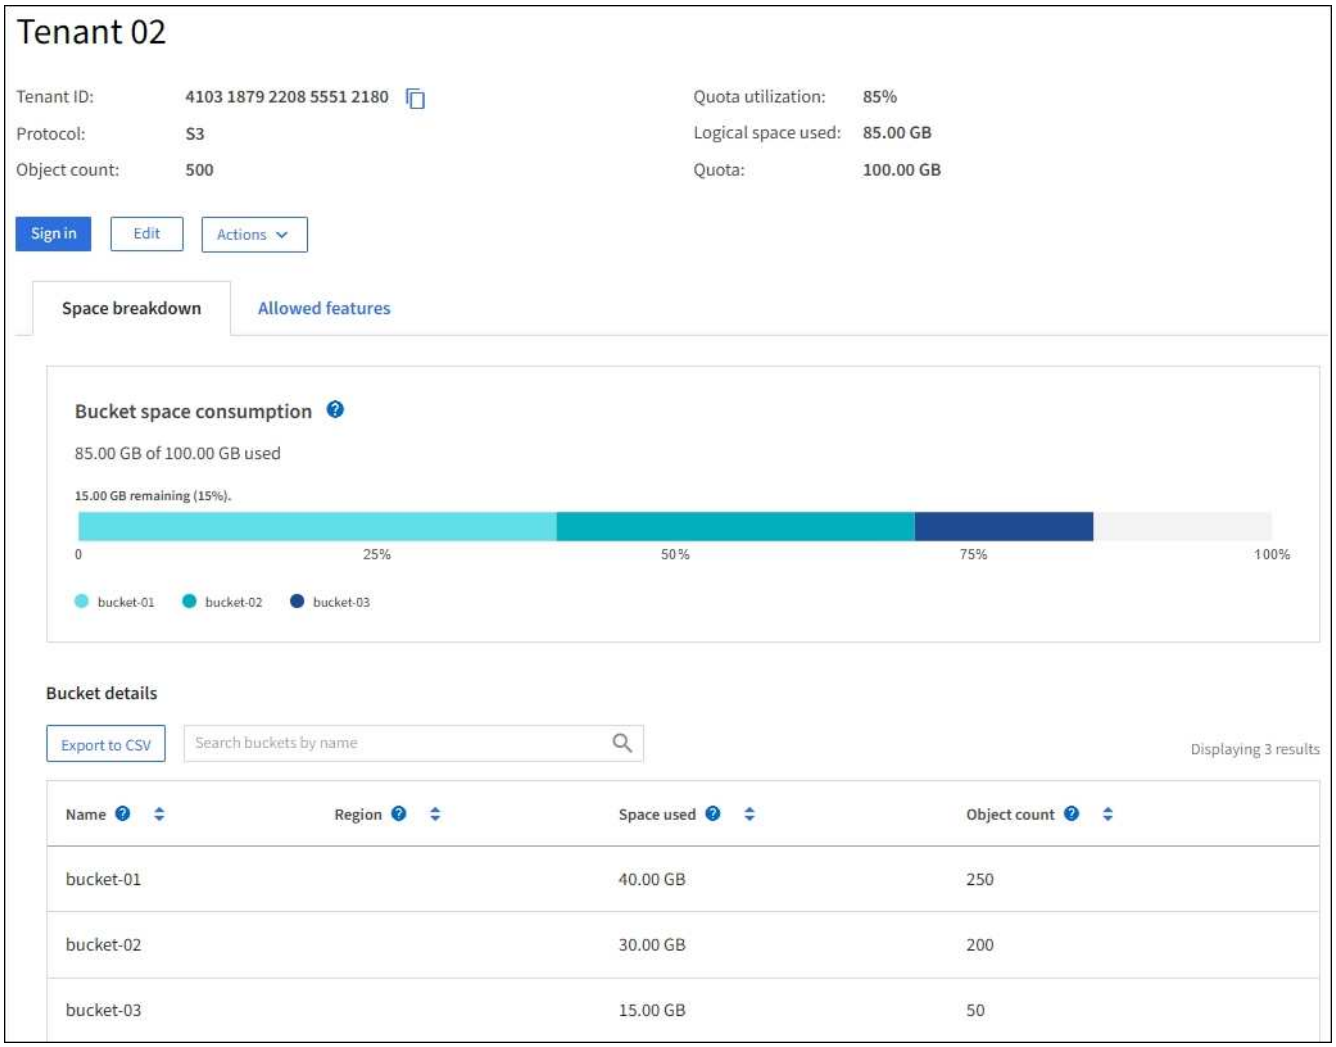

2. Review the tenant overview at the top of the page.

   This section of the details page provides summary information for the tenant, including the tenant's object count, quota utilization, logical space used, and quota setting.

3. From the **Space breakdown** tab, review the **Space consumption** chart.

   This chart shows the total space consumption for all of the tenant's S3 buckets (or Swift containers).

   If a quota was set for this tenant, the amount of quota used and remaining is displayed in text (for example, `85.00 GB of 100 GB used`). If no quota was set, the tenant has an unlimited quota, and the text includes only an amount of space used (for example, `85.00 GB used`). The bar chart shows the percentage of quota in each bucket or container. If the tenant has exceeded the storage quota by more than 1% and by at least 1 GB, the chart shows the total quota and the excess amount.

You can place your cursor over the bar chart to see the storage used by each bucket or container. You can place your cursor over the free space segment to see the amount of storage quota remaining.

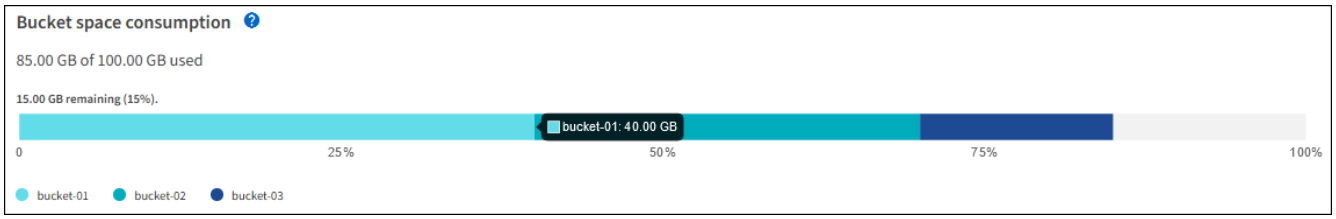

Quota utilization is based on internal estimates and might be exceeded in some cases. For example, StorageGRID checks the quota when a tenant starts uploading objects and rejects new ingests if the tenant has exceeded the quota. However, StorageGRID does not take into account the size of the current upload when determining if the quota has been exceeded. If objects are deleted, a tenant might be temporarily prevented from uploading new objects until the quota utilization is recalculated. Quota utilization calculations can take 10 minutes or longer.

A tenant's quota utilization indicates the total amount of object data the tenant has uploaded to StorageGRID (logical size). The quota utilization does not represent the space used to store copies of those objects and their metadata (physical size).

You can enable the **Tenant quota usage high** alert rule to determine if tenants are consuming their quotas. If enabled, this alert is triggered when a tenant has used 90% of its quota. For instructions, see Edit alert rules.

4. From the **Space breakdown** tab, review the **Bucket details**.

   This table lists the S3 buckets (or Swift containers) for the tenant. Space used is the total amount of object data in the bucket or container. This value does not represent the storage space required for ILM copies and object metadata.

5. Optionally, select **Export to CSV** to view and export a .csv file containing the usage values for each bucket or container.

   The contents of an individual S3 tenant's `.csv` file look like the following example:

Tenant ID	Bucket Name	Space Used (Bytes)	Number of Objects
64796966429038923647	bucket-01	88717711	14
64796966429038923647	bucket-02	21747507	11
64796966429038923647	bucket-03	15294070	3

   You can open the `.csv` file in a spreadsheet application or use it in automation.

6. Optionally, select the **Allowed features** tab to see a list of the permissions and features that are enabled for the tenant. See Edit tenant account if you need to change any of these settings.
7. If the tenant has the **Use grid federation connection** permission, optionally select the **Grid federation** tab to learn more about the connection.

   See What is grid federation? and Manage the permitted tenants for grid federation.

View network traffic

If traffic classification policies are in place for a tenant, review the network traffic for that tenant.

Steps

1. Select **CONFIGURATION** > **Network** > **Traffic classification**.

   The Traffic Classification Policies page appears, and the existing policies are listed in the table.

2. Review the list of policies to identify the ones that apply to a specific tenant.
3. To view metrics associated with a policy, select the radio button to the left of the policy, and select **Metrics**.
4. Analyze the graphs to determine how often the policy is limiting traffic and whether you need to adjust the policy.

See Manage traffic classification policies for more information.

Use the audit log

Optionally, you can use the audit log for more granular monitoring of a tenant's activities.

For instance, you can monitor the following types of information:

- Specific client operations, such as PUT, GET, or DELETE
- Object sizes
- The ILM rule applied to objects
- The source IP of client requests

Audit logs are written to text files that you can analyze using your choice of log analysis tool. This allows you to better understand client activities, or to implement sophisticated chargeback and billing models.

See Review audit logs for more information.

Use Prometheus metrics

Optionally, use Prometheus metrics to report on tenant activity.

- In the Grid Manager, select **SUPPORT** > **Tools** > **Metrics**. You can use existing dashboards, such as S3 Overview, to review client activities.

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

- From the top of the Grid Manager, select the help icon and select **API documentation**. You can use the metrics in the Metrics section of the Grid Management API to create custom alert rules and dashboards for tenant activity.

See Review support metrics for more information.

Monitor S3 and Swift client operations

You can monitor object ingest and retrieval rates as well as metrics for object counts, queries, and verification. You can view the number of successful and failed attempts by

client applications to read, write, and modify objects in the StorageGRID system.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.

Steps

1. From the dashboard, select the **Performance** tab.
2. Refer to the S3 and Swift charts, which summarize the number of client operations performed by Storage Nodes and the number of API requests received by Storage Nodes during the selected time frame.
3. Select **NODES** to access the Nodes page.
4. From the Nodes home page (grid level), select the **Objects** tab.

   The chart shows S3 and Swift ingest and retrieve rates for your entire StorageGRID system in bytes per second and the amount of data ingested or retrieved. You can select a time interval or apply a custom interval.

5. To see information for a particular Storage Node, select the node from the list on the left, and select the **Objects** tab.

   The chart shows the ingest and retrieve rates for the node. The tab also includes metrics for object counts, metadata queries, and verification operations.

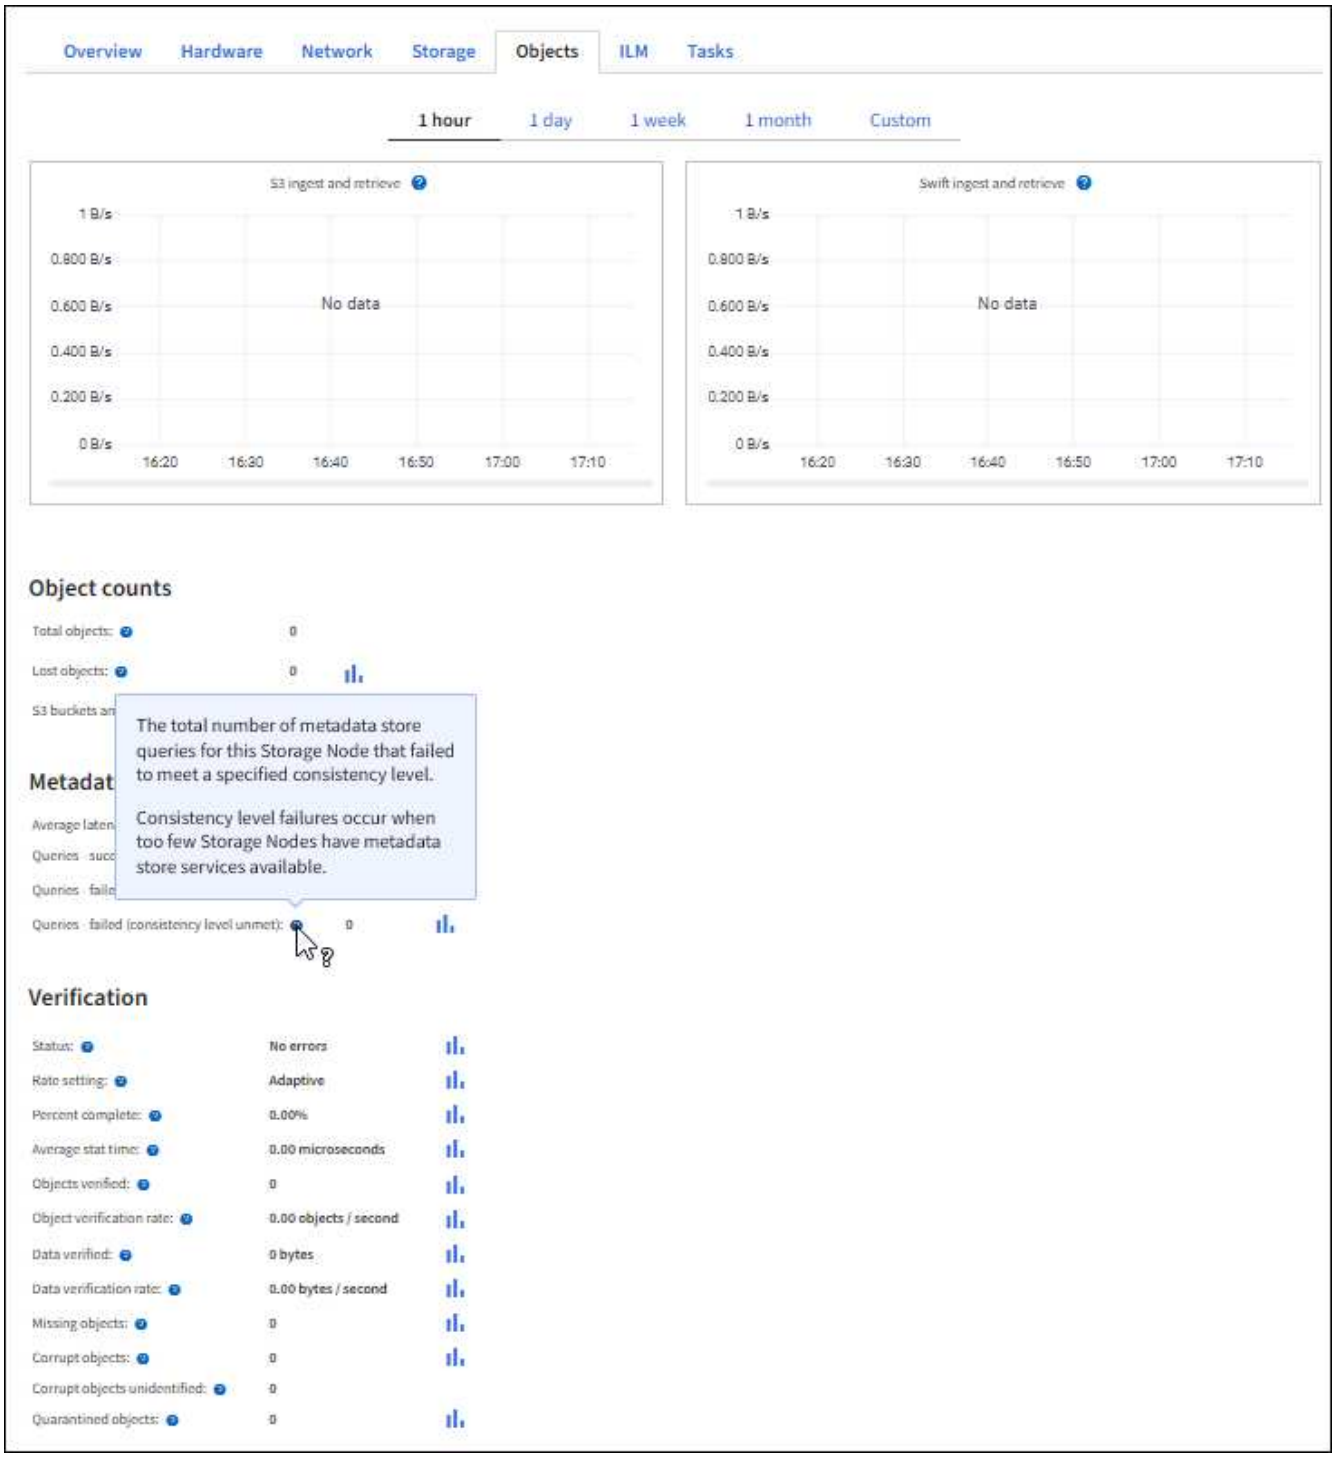

Monitor load balancing operations

If you are using a load balancer to manage client connections to StorageGRID, you should monitor load balancing operations after you configure the system initially and after you make any configuration changes or perform an expansion.

About this task

You can use the Load Balancer service on Admin Nodes or Gateway Nodes or an external third-party load balancer to distribute client requests across multiple Storage Nodes.

After configuring load balancing, you should confirm that object ingest and retrieval operations are being evenly distributed across Storage Nodes. Evenly distributed requests ensure that StorageGRID remains responsive to client requests under load and can help maintain client performance.

If you configured a high availability (HA) group of Gateway Nodes or Admin Nodes in active-backup mode, only one node in the group actively distributes client requests.

For more information, see Configure S3 and Swift client connections.

Steps

1. If S3 or Swift clients connect using the Load Balancer service, check that Admin Nodes or Gateway Nodes are actively distributing traffic as you expect:
   a. Select **NODES**.
   b. Select a Gateway Node or Admin Node.
   c. On the **Overview** tab, check if a node interface is in an HA group and if the node interface has the role of Primary.

      Nodes with the role of Primary and nodes that aren't in an HA group should be actively distributing requests to clients.

   d. For each node that should be actively distributing client requests, select the Load Balancer tab.
   e. Review the chart of Load Balancer Request Traffic for the last week to ensure that the node has been actively distributing requests.

      Nodes in an active-backup HA group might take the Backup role from time to time. During that time the nodes don't distribute client requests.

   f. Review the chart of Load Balancer Incoming Request Rate for the last week to review the object throughput of the node.
   g. Repeat these steps for each Admin Node or Gateway Node in the StorageGRID system.
   h. Optionally, use traffic classification policies to view a more detailed analysis of traffic being served by the Load Balancer service.
2. Verify that these requests are being evenly distributed to Storage Nodes.
   a. Select ***Storage Node*** > **LDR** > **HTTP**.
   b. Review the number of **Currently Established incoming Sessions**.
   c. Repeat for each Storage Node in the grid.

      The number of sessions should be roughly equal across all Storage Nodes.

Monitor grid federation connections

You can monitor basic information about all grid federation connections, detailed information about a specific connection, or Prometheus metrics about cross-grid replication operations. You can monitor a connection from either grid.

Before you begin

- You are signed in to the Grid Manager on either grid using a supported web browser.
- You have the Root access permission for the grid you are signed in to.

View all connections

The Grid federation page shows basic information about all grid federation connections and about all tenant accounts that are permitted to use grid federation connections.

Steps

1. Select **CONFIGURATION** > **System** > **Grid federation**.

   The Grid federation page appears.

2. To see basic information for all connections on this grid, select the **Connections** tab.

   From this tab, you can:

   - Create a new connection.
   - Select an existing connection to edit or test.

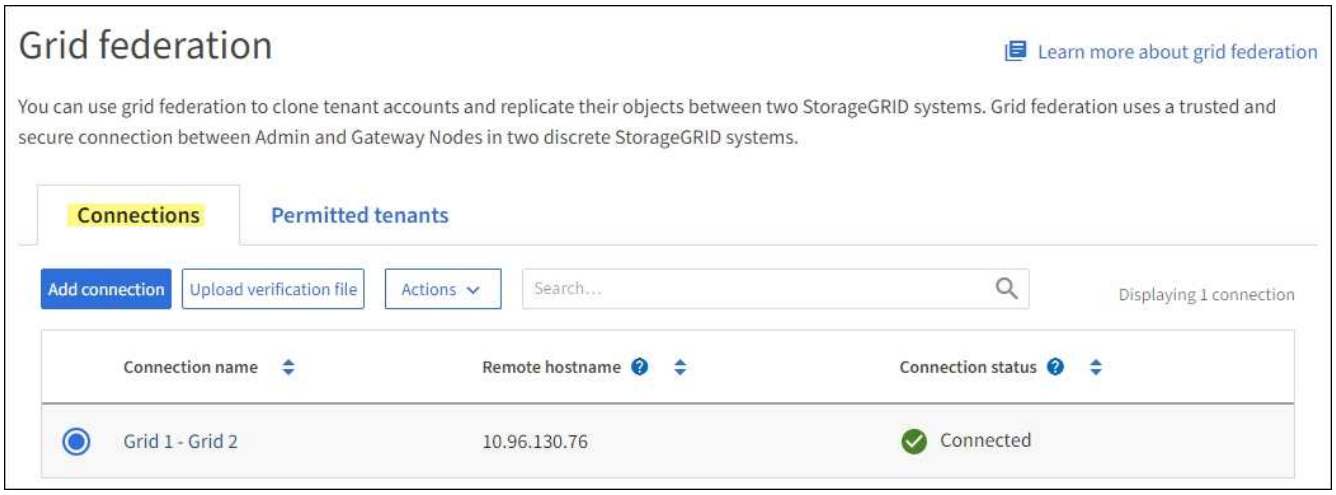

3. To see basic information for all tenant accounts on this grid that have the **Use grid federation connection** permission, select the **Permitted tenants** tab.

   From this tab, you can:

   - View the details page for each permitted tenant.
   - View the details page for each connection. See View a specific connection.
   - Select a permitted tenant and remove the permission.
   - Check for cross-grid replication errors and clear the last error, if any. See Troubleshoot grid federation errors.

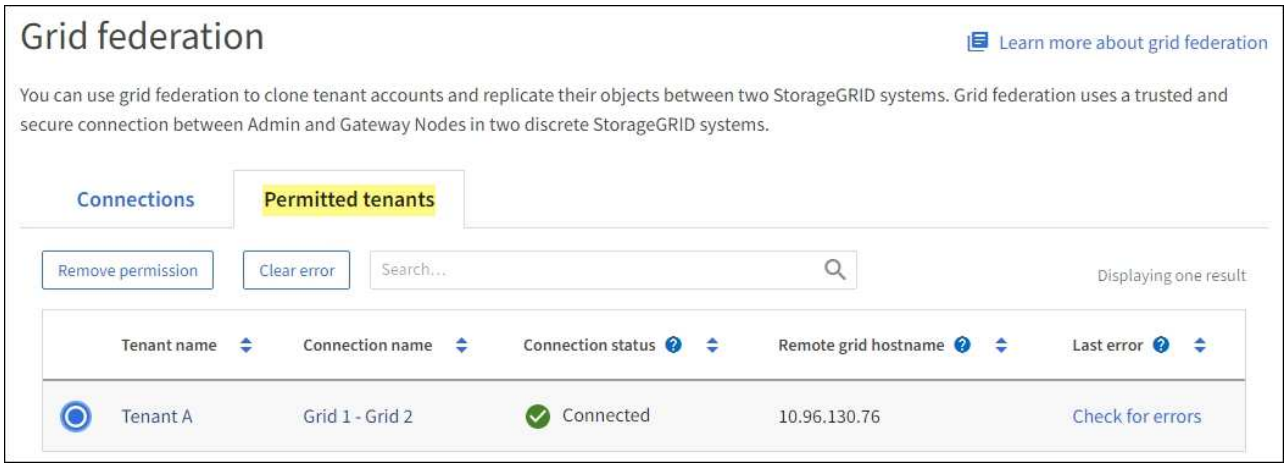

View a specific connection

You can view details for a specific grid federation connection.

Steps

1. Select either tab from the Grid federation page and then select the connection name from the table.

   From the details page for the connection, you can:

   - See basic status information about the connection, including the local and remote hostnames, port, and connection status.
   - Select a connection to edit, test, or remove.

2. When viewing a specific connection, select the **Permitted tenants** tab to view details about the permitted tenants for the connection.

   From this tab, you can:

   - View the details page for each permitted tenant.
   - Remove a tenant's permission to use the connection.
   - Check for cross-grid replication errors and clear the last error. See Troubleshoot grid federation errors.

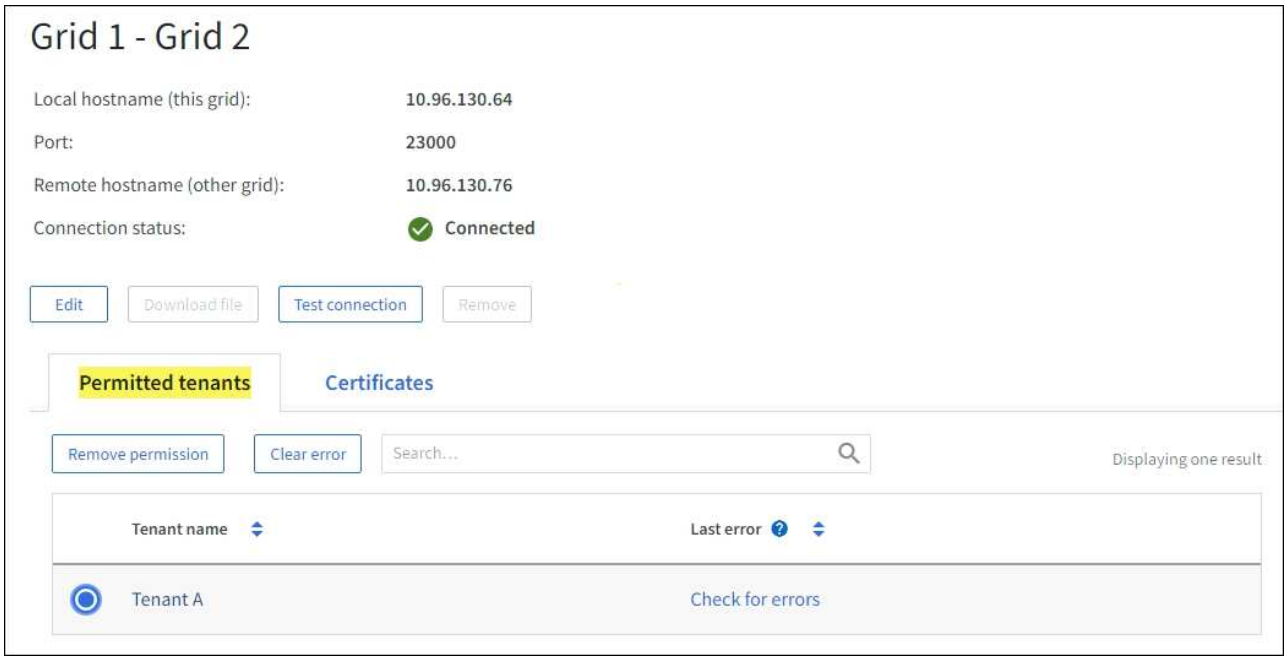

3. When viewing a specific connection, select the **Certificates** tab to view the system-generated server and client certificates for this connection.

   From this tab, you can:

   - Rotate connection certificates.
   - Select **Server** or **Client** to view or download the associated certificate or copy the certificate PEM.

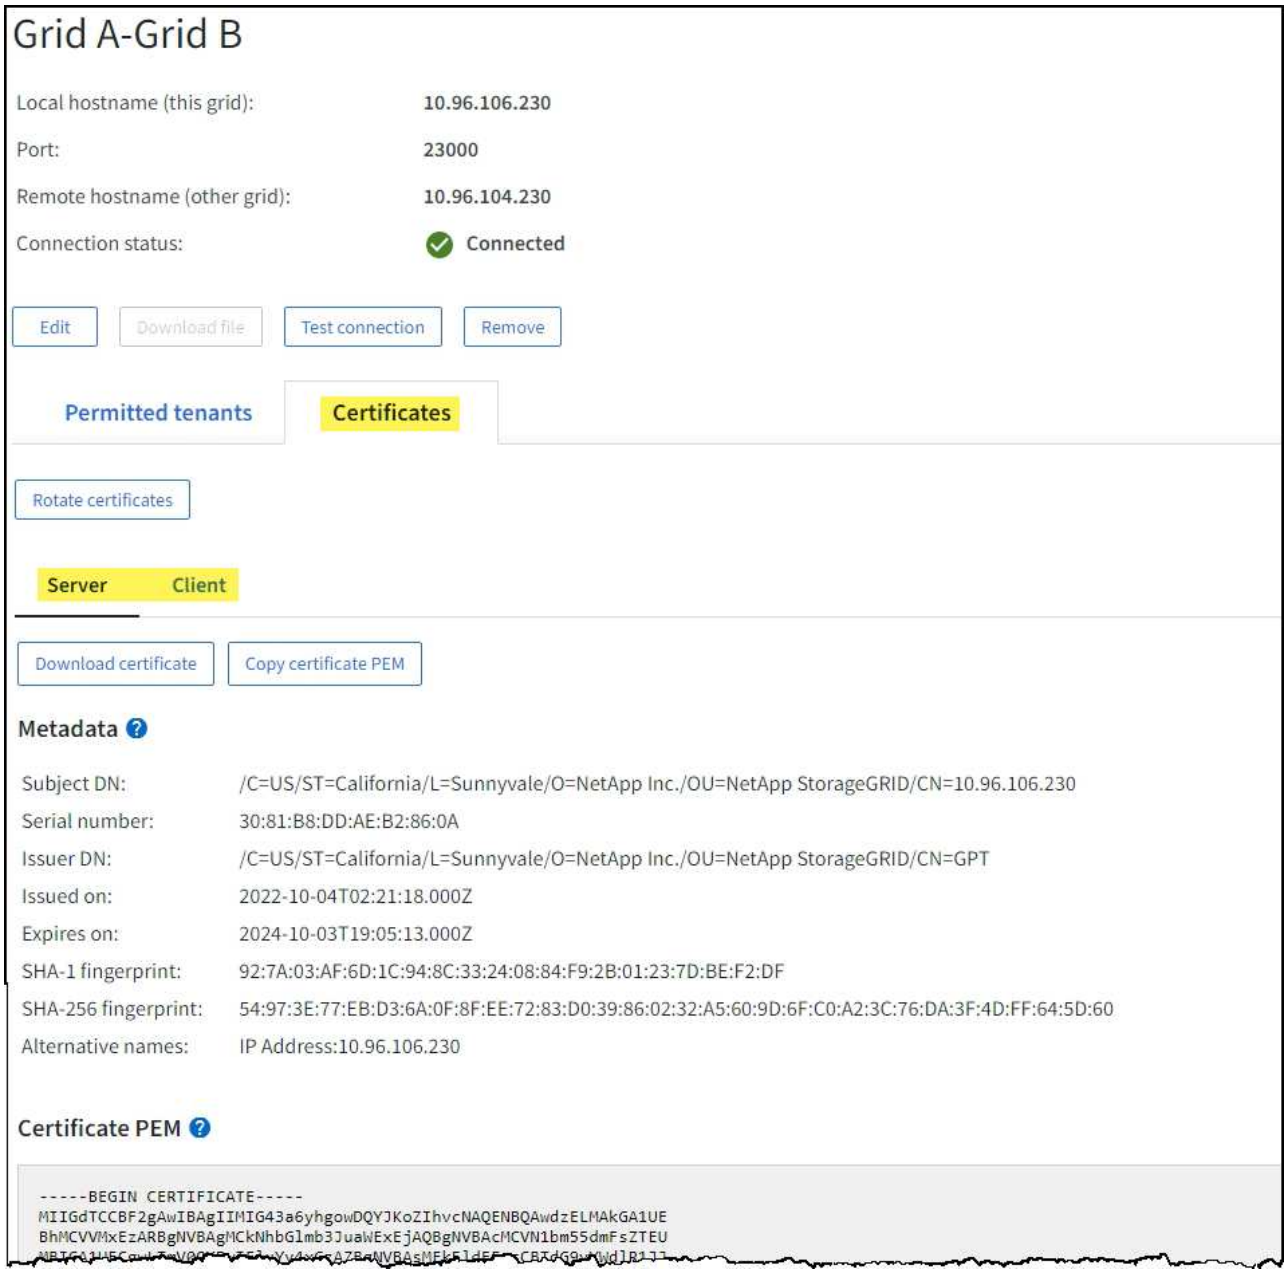

Review cross-grid replication metrics

You can use the Cross-Grid Replication dashboard in Grafana to view Prometheus metrics about cross-grid replication operations on your grid.

Steps

1. From the Grid Manager, select **SUPPORT** > **Tools** > **Metrics**.

> 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. See the list of commonly used Prometheus metrics.

2. In the Grafana section of the page, select **Cross Grid Replication**.

   For detailed instructions, see Review support metrics.

3. To retry replication of objects that failed to replicate, see Identify and retry failed replication operations.

Monitor archival capacity

You can't directly monitor an external archival storage system's capacity through the StorageGRID system. However, you can monitor whether the Archive Node can still send object data to the archival destination, which might indicate that an expansion of archival media is required.

About this task

You can monitor the Store component to check if the Archive Node can still send object data to the targeted archival storage system. The Store Failures (ARVF) alarm might also indicate that the targeted archival storage system has reached capacity and can no longer accept object data.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select ***Archive Node*** > **ARC> Overview> Main**.
3. Check the Store State and Store Status attributes to confirm that the Store component is Online with No Errors.

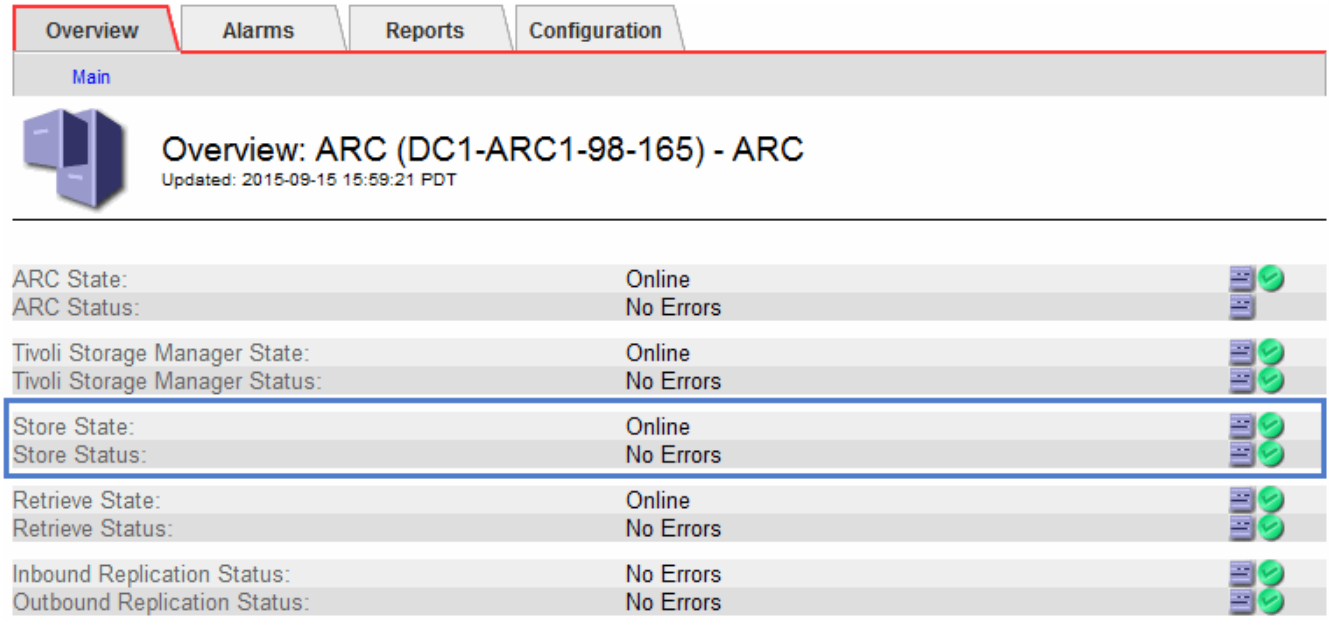

An offline Store component or one with errors might indicate that targeted archival storage system can no longer accept object data because it has reached capacity.

Alerts and alarms

Manage alerts and alarms: Overview

The StorageGRID alert system is designed to inform you about operational issues that require your attention. The legacy alarm system is deprecated.

Alert system

The alert system is designed to be your primary tool for monitoring any issues that might occur in your StorageGRID system. The alert system provides an easy-to-use interface for detecting, evaluating, and resolving issues.

Alerts are triggered at specific severity levels when alert rule conditions evaluate as true. When an alert is triggered, the following actions occur:

- An alert severity icon is shown on the dashboard in the Grid Manager, and the count of Current Alerts is incremented.
- The alert is shown on the **NODES** summary page and on the **NODES** > ***node*** > **Overview** tab.
- An email notification is sent, assuming you have configured an SMTP server and provided email addresses for the recipients.
- An Simple Network Management Protocol (SNMP) notification is sent, assuming you have configured the StorageGRID SNMP agent.

Legacy alarm system

Like alerts, alarms are triggered at specific severity levels when attributes reach defined threshold values. However, unlike alerts, many alarms are triggered for events that you can safely ignore, which might result in an excessive number of email or SNMP notifications.

> The alarm system is deprecated and will be removed in a future release. If you are still using legacy alarms, you should fully transition to the alert system as soon as possible.

When an alarm is triggered, the following actions occur:

- The alarm appears on the **SUPPORT** > **Alarms (legacy)** > **Current alarms** page.
- An email notification is sent, assuming you have configured an SMTP server and configured one or more mailing lists.
- An SNMP notification might be sent, assuming you have configured the StorageGRID SNMP agent. (SNMP notifications aren't sent for all alarms or alarm severities.)

Compare alerts and alarms

There are several similarities between the alert system and the legacy alarm system, but the alert system offers significant benefits and is easier to use.

Refer to the following table to learn how to perform similar operations.

	Alerts	Alarms (legacy system)
How do I see which alerts or alarms are active?	• Select the **Current alerts** link on the dashboard. • Select the alert on the **NODES** > **Overview** page. • Select **ALERTS** > **Current**. View current alerts	Select **SUPPORT** > **Alarms (legacy)** > **Current alarms**. Manage alarms (legacy system)

	Alerts	Alarms (legacy system)
What causes an alert or an alarm to be triggered?	Alerts are triggered when a Prometheus expression in an alert rule evaluates as true for the specific trigger condition and duration. View alert rules	Alarms are triggered when a StorageGRID attribute reaches a threshold value. Manage alarms (legacy system)
If an alert or alarm is triggered, how do I resolve the underlying problem?	The recommended actions for an alert are included in email notifications and are available from the Alerts pages in the Grid Manager. As required, additional information is provided in the StorageGRID documentation. Alerts reference	You can learn about an alarm by selecting the attribute name, or you can search for an alarm code in the StorageGRID documentation. Alarms reference (legacy system)
Where can I see a list of alerts or alarms that have been resolved?	Select **ALERTS** > **Resolved**. View current and resolved alerts	Select **SUPPORT** > **Alarms (legacy)** > **Historical alarms**. Manage alarms (legacy system)
Where do I manage the settings?	Select **ALERTS** > **Rules**. Manage alerts	Select **SUPPORT**. Then, use the options in the **Alarms (legacy)** section of the menu. Manage alarms (legacy system)
What user group permissions do I need?	• Anyone who can sign in to the Grid Manager can view current and resolved alerts. • You must have the Manage alerts permission to manage silences, alert notifications, and alert rules. Administer StorageGRID	• Anyone who can sign in to the Grid Manager can view legacy alarms. • You must have the Acknowledge alarms permission to acknowledge alarms. • You must have both the Grid topology page configuration and Other grid configuration permissions to manage global alarms and email notifications. Administer StorageGRID

	Alerts	Alarms (legacy system)
How do I manage email notifications?	Select **ALERTS** > **Email setup**. **Note:** Because alarms and alerts are independent systems, the email setup used for alarm and AutoSupport notifications is not used for alert notifications. However, you can use the same mail server for all notifications. Set up email notifications for alerts	Select **SUPPORT** > **Alarms (legacy)** > **Legacy email setup**. Manage alarms (legacy system)
How do I manage SNMP notifications?	Select **CONFIGURATION** > **Monitoring** > **SNMP agent**. Use SNMP monitoring	*Not supported*
How do I control who receives notifications?	1. Select **ALERTS** > **Email setup**. 2. In the **Recipients** section, enter an email address for each email list or person who should receive an email when an alert occurs. Set up email notifications for alerts	1. Select **SUPPORT** > **Alarms (legacy)** > **Legacy email setup**. 2. Creating a mailing list. 3. Select **Notifications**. 4. Select the mailing list. Manage alarms (legacy system)
Which Admin Nodes send notifications?	A single Admin Node (the preferred sender). What is an Admin Node?	A single Admin Node (the preferred sender). What is an Admin Node?
How do I suppress some notifications?	1. Select **ALERTS** > **Silences**. 2. Select the alert rule you want to silence. 3. Specify a duration for the silence. 4. Select the severity of alert you want to silence. 5. Select to apply the silence to the entire grid, a single site, or a single node. **Note**: If you have enabled the SNMP agent, silences also suppress SNMP traps and informs. Silence alert notifications	1. Select **SUPPORT** > **Alarms (legacy)** > **Legacy email setup**. 2. Select **Notifications**. 3. Select a mailing list, and select **Suppress**. Manage alarms (legacy system)

	Alerts	Alarms (legacy system)
How do I suppress all notifications?	Select **ALERTS** > **Silences**.Then, select **All rules**. **Note**: If you have enabled the SNMP agent, silences also suppress SNMP traps and informs. Silence alert notifications	*Not supported*
How do I customize the conditions and triggers?	1. Select **ALERTS** > **Rules**. 2. Select a default rule to edit, or select **Create custom rule**. Edit alert rules Create custom alert rules	1. Select **SUPPORT** > **Alarms (legacy)** > **Global alarms**. 2. Create a Global Custom alarm to override a Default alarm or to monitor an attribute that does not have a Default alarm. Manage alarms (legacy system)
How do I disable an individual alert or alarm?	1. Select **ALERTS** > **Rules**. 2. Select the rule, and select **Edit rule**. 3. Clear the **Enabled** checkbox. Disable alert rules	1. Select **SUPPORT** > **Alarms (legacy)** > **Global alarms**. 2. Select the rule, and select the Edit icon. 3. Clear the **Enabled** checkbox. Manage alarms (legacy system)

Manage alerts

Manage alerts: overview

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

You can create custom alerts, edit or disable alerts, and manage alert notifications.

To learn more:

- Review the video: Video: Alerts overview for StorageGRID 11.8

- Review the video: Video: Using metrics to create custom alerts in StorageGRID 11.8

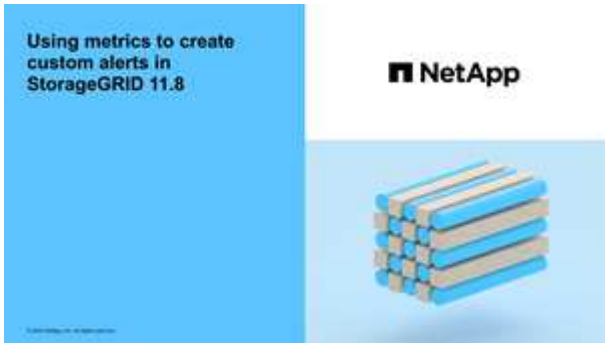

- See the Alerts reference.

View alert rules

Alert rules define the conditions that trigger specific alerts. StorageGRID includes a set of default alert rules, which you can use as is or modify, or you can create custom alert rules.

You can view the list of all default and custom alert rules to learn which conditions will trigger each alert and to see whether any alerts are disabled.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Manage alerts or Root access permission.
- Optionally, you have watched the video: Video: Alerts overview for StorageGRID 11.8

Steps

1. Select **ALERTS** > **Rules**.

   The Alert Rules page appears.

Alert Rules Learn more

Alert rules define which conditions trigger specific alerts.

You can edit the conditions for default alert rules to better suit your environment, or create custom alert rules that use your own conditions for triggering alerts.

Create custom rule | Edit rule | Remove custom rule

Name	Conditions	Type	Status
Appliance battery expired The battery in the appliance's storage controller has expired.	storagegrid_appliance_component_failure{type="REC_EXPIRED_BATTERY"} *Major* > 0	Default	Enabled
Appliance battery failed The battery in the appliance's storage controller has failed.	storagegrid_appliance_component_failure{type="REC_FAILED_BATTERY"} *Major* > 0	Default	Enabled
Appliance battery has insufficient learned capacity The battery in the appliance's storage controller has insufficient learned capacity.	storagegrid_appliance_component_failure{type="REC_BATTERY_WARN"} *Major* > 0	Default	Enabled
Appliance battery near expiration The battery in the appliance's storage controller is nearing expiration.	storagegrid_appliance_component_failure{type="REC_BATTERY_NEAR_EXPIRATION"} *Major* > 0	Default	Enabled
Appliance battery removed The battery in the appliance's storage controller is missing.	storagegrid_appliance_component_failure{type="REC_REMOVED_BATTERY"} *Major* > 0	Default	Enabled
Appliance battery too hot The battery in the appliance's storage controller is overheated.	storagegrid_appliance_component_failure{type="REC_BATTERY_OVERTEMP"} *Major* > 0	Default	Enabled
Appliance cache backup device failed A persistent cache backup device has failed.	storagegrid_appliance_component_failure{type="REC_CACHE_BACKUP_DEVICE_FAILED"} *Major* > 0	Default	Enabled
Appliance cache backup device insufficient capacity There is insufficient cache backup device capacity.	storagegrid_appliance_component_failure{type="REC_CACHE_BACKUP_DEVICE_INSUFFICIENT_CAPACITY"} *Major* > 0	Default	Enabled
Appliance cache backup device write-protected A cache backup device is write-protected.	storagegrid_appliance_component_failure{type="REC_CACHE_BACKUP_DEVICE_WRITE_PROTECTED"} *Major* > 0	Default	Enabled
Appliance cache memory size mismatch The two controllers in the appliance have different cache sizes.	storagegrid_appliance_component_failure{type="REC_CACHE_MEM_SIZE_MISMATCH"} *Major* > 0	Default	Enabled

Displaying 62 alert rules.

2. Review the information in the alert rules table:

Column header	Description
Name	The unique name and description of the alert rule. Custom alert rules are listed first, followed by default alert rules. The alert rule name is the subject for email notifications.
Conditions	The Prometheus expressions that determine when this alert is triggered. An alert can be triggered at one or more of the following severity levels, but a condition for each severity is not required. • **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 don't clear on their own to ensure they don't result in a more serious problem.

Column header	Description
Type	The type of alert rule: • **Default**: An alert rule provided with the system. You can disable a default alert rule or edit the conditions and duration for a default alert rule. You can't remove a default alert rule. • **Default***: A default alert rule that includes an edited condition or duration. As required, you can easily revert a modified condition back to the original default. • **Custom**: An alert rule that you created. You can disable, edit, and remove custom alert rules.
Status	Whether this alert rule is currently enabled or disabled. The conditions for disabled alert rules aren't evaluated, so no alerts are triggered.

Create custom alert rules

You can create custom alert rules to define your own conditions for triggering alerts.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Manage alerts or Root access permission.
- You are familiar with the commonly used Prometheus metrics.
- You understand the syntax of Prometheus queries.
- Optionally, you have watched the video: Video: Using metrics to create custom alerts in StorageGRID 11.8.

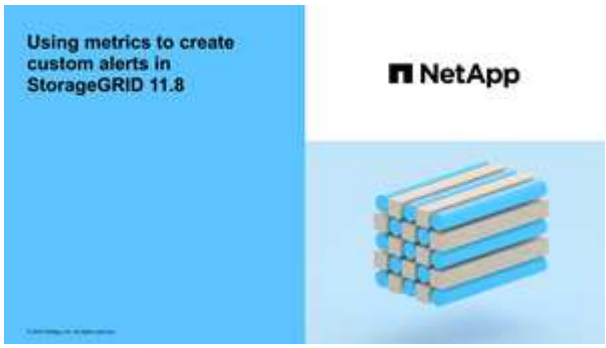

About this task

StorageGRID does not validate custom alerts. If you decide to create custom alert rules, follow these general guidelines:

- Look at the conditions for the default alert rules, and use them as examples for your custom alert rules.
- If you define more than one condition for an alert rule, use the same expression for all conditions. Then, change the threshold value for each condition.
- Carefully check each condition for typos and logic errors.
- Use only the metrics listed in the Grid Management API.

- When testing an expression using the Grid Management API, be aware that a "successful" response might be an empty response body (no alert triggered). To see if the alert is actually triggered, you can temporarily set a threshold to a value you expect to be true currently.

  For example, to test the expression `node_memory_MemTotal_bytes < 24000000000`, first execute `node_memory_MemTotal_bytes >= 0` and ensure you get the expected results (all nodes return a value). Then, change the operator and the threshold back to the intended values and execute again. No results indicate there are no current alerts for this expression.

- Don’t assume a custom alert is working unless you have validated that the alert is triggered when expected.

Steps

1. Select **ALERTS** > **Rules**.

   The Alert Rules page appears.

2. Select **Create custom rule**.

   The Create Custom Rule dialog box appears.

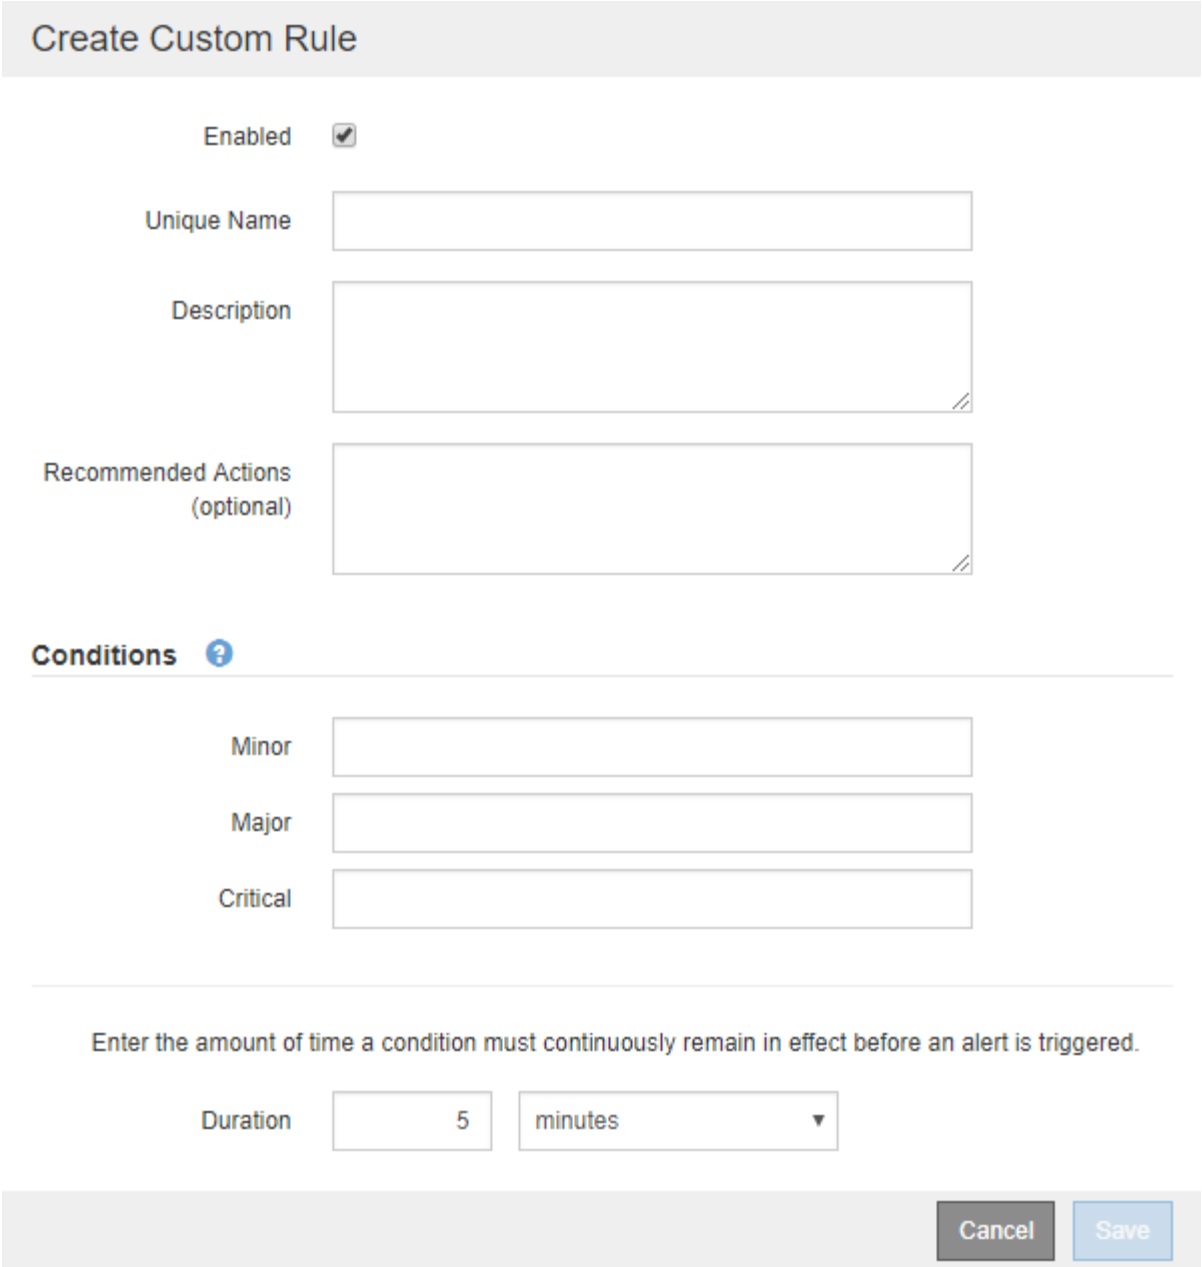

3. Select or clear the **Enabled** checkbox to determine if this alert rule is currently enabled.

   If an alert rule is disabled, its expressions aren't evaluated and no alerts are triggered.

4. Enter the following information:

Field	Description
Unique Name	A unique name for this rule. The alert rule name is shown on the Alerts page and is also the subject for email notifications. Names for alert rules can be between 1 and 64 characters.
Description	A description of the problem that is occurring. The description is the alert message shown on the Alerts page and in email notifications. Descriptions for alert rules can be between 1 and 128 characters.

Field	Description
Recommended Actions	Optionally, the recommended actions to take when this alert is triggered. Enter recommended actions as plain text (no formatting codes). Recommended actions for alert rules can be between 0 and 1,024 characters.

5. In the Conditions section, enter a Prometheus expression for one or more of the alert severity levels.

   A basic expression is usually of the form:

   ```
   [metric] [operator] [value]
   ```

   Expressions can be any length, but appear on a single line in the user interface. At least one expression is required.

   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_bytes < 24000000000
   ```

   To see available metrics and to test Prometheus expressions, select the help icon and follow the link to the Metrics section of the Grid Management API.

6. In the **Duration** field, enter the amount of time a condition must continuously remain in effect before the alert is triggered, and select a unit of time.

   To trigger an alert immediately when a condition becomes true, enter **0**. Increase this value to prevent temporary conditions from triggering alerts.

   The default is 5 minutes.

7. Select **Save**.

   The dialog box closes, and the new custom alert rule appears in the Alert Rules table.

Edit alert rules

You can edit an alert rule to change the trigger conditions, For a custom alert rule, you can also update the rule name, description, and recommended actions.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Manage alerts or Root access permission.

About this task

When you edit a default alert rule, you can change the conditions for minor, major, and critical alerts; and the duration. When you edit a custom alert rule, you can also edit the rule's name, description, and recommended actions.

> Be careful when deciding to edit an alert rule. If you change trigger values, you might not detect an underlying problem until it prevents a critical operation from completing.

Steps

1. Select **ALERTS** > **Rules**.

   The Alert Rules page appears.

2. Select the radio button for the alert rule you want to edit.
3. Select **Edit rule**.

   The Edit Rule dialog box appears. This example shows a default alert rule—the Unique Name, Description, and Recommended Actions fields are disabled and can't be edited.

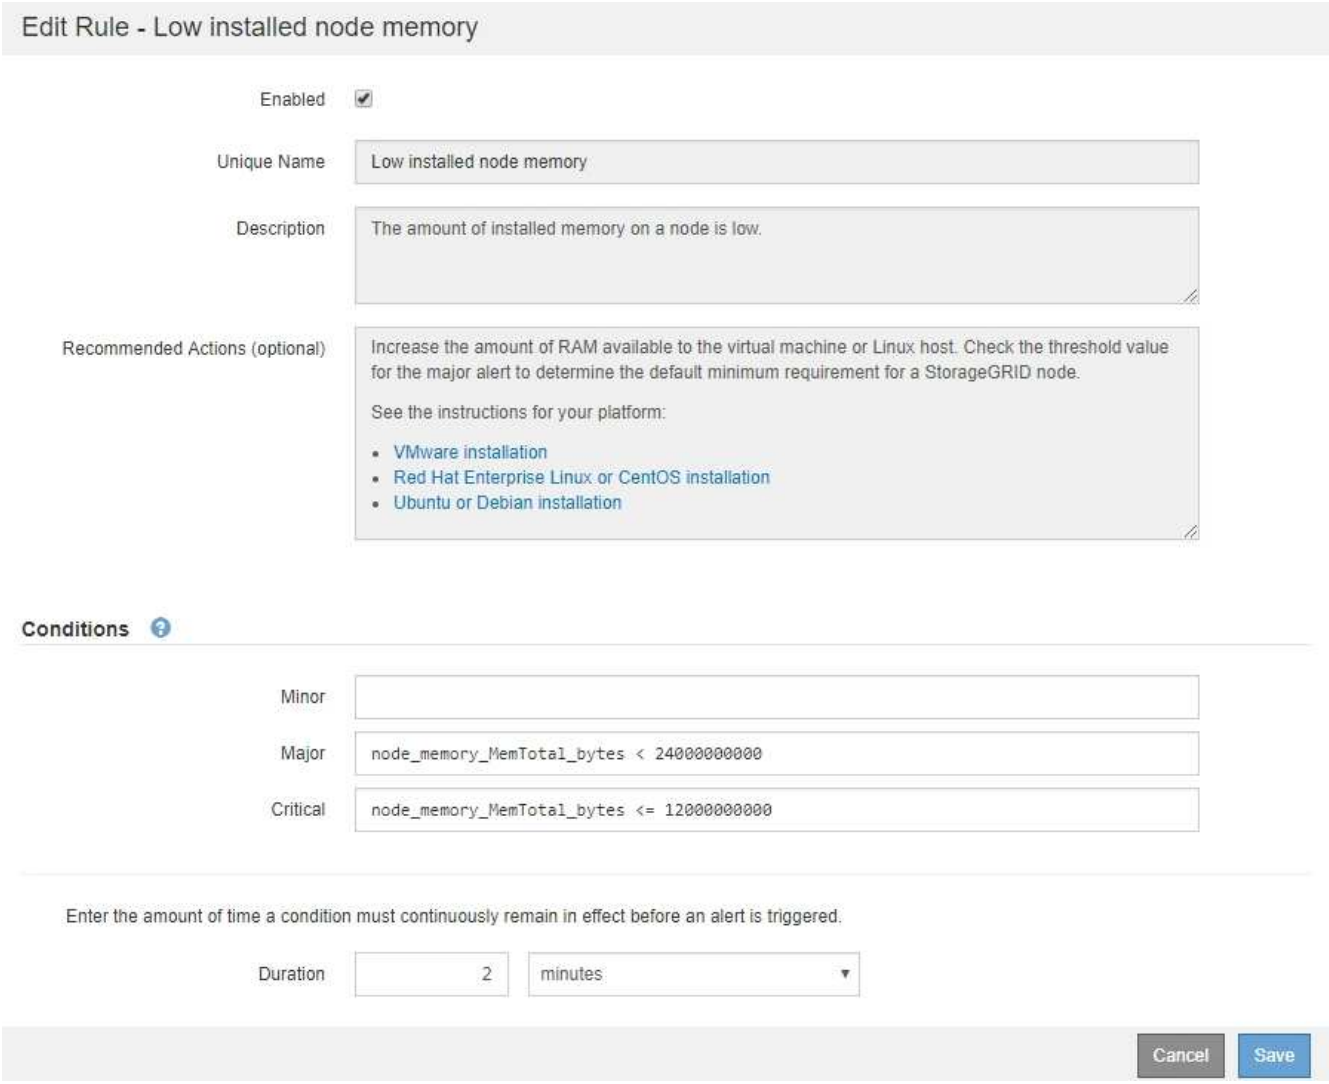

4. Select or clear the **Enabled** checkbox to determine if this alert rule is currently enabled.

   If an alert rule is disabled, its expressions aren't evaluated and no alerts are triggered.

   If you disable the alert rule for a current alert, you must wait a few minutes for the alert to no longer appear as an active alert.

   In general, disabling a default alert rule is not recommended. If an alert rule is disabled, you might not detect an underlying problem until it prevents a critical operation from completing.

5. For custom alert rules, update the following information as required.

> You can't edit this information for default alert rules.

Field	Description
Unique Name	A unique name for this rule. The alert rule name is shown on the Alerts page and is also the subject for email notifications. Names for alert rules can be between 1 and 64 characters.
Description	A description of the problem that is occurring. The description is the alert message shown on the Alerts page and in email notifications. Descriptions for alert rules can be between 1 and 128 characters.
Recommended Actions	Optionally, the recommended actions to take when this alert is triggered. Enter recommended actions as plain text (no formatting codes). Recommended actions for alert rules can be between 0 and 1,024 characters.

6. In the Conditions section, enter or update the Prometheus expression for one or more of the alert severity levels.

> If you want to restore a condition for an edited default alert rule back to its original value, select the three dots to the right of the modified condition.

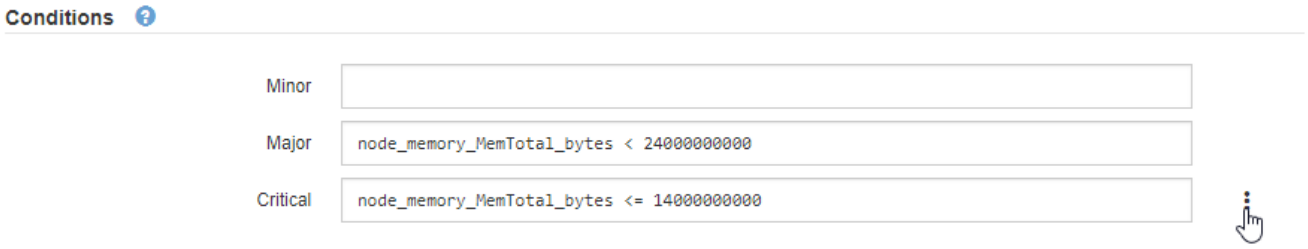

> If you update the conditions for a current alert, your changes might not be implemented until the previous condition is resolved. The next time one of the conditions for the rule is met, the alert will reflect the updated values.

A basic expression is usually of the form:

```
[metric] [operator] [value]
```

Expressions can be any length, but appear on a single line in the user interface. At least one expression is required.

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_bytes < 24000000000
```

7. In the **Duration** field, enter the amount of time a condition must continuously remain in effect before the

alert is triggered, and select the unit of time.

To trigger an alert immediately when a condition becomes true, enter **0**. Increase this value to prevent temporary conditions from triggering alerts.

The default is 5 minutes.

8. Select **Save**.

   If you edited a default alert rule, **Default*** appears in the Type column. If you disabled a default or custom alert rule, **Disabled** appears in the **Status** column.

Disable alert rules

You can change the enabled/disabled state for a default or custom alert rule.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Manage alerts or Root access permission.

About this task

When an alert rule is disabled, its expressions aren't evaluated and no alerts are triggered.

> In general, disabling a default alert rule is not recommended. If an alert rule is disabled, you might not detect an underlying problem until it prevents a critical operation from completing.

Steps

1. Select **ALERTS** > **Rules**.

   The Alert Rules page appears.

2. Select the radio button for the alert rule you want to disable or enable.
3. Select **Edit rule**.

   The Edit Rule dialog box appears.

4. Select or clear the **Enabled** checkbox to determine if this alert rule is currently enabled.

   If an alert rule is disabled, its expressions aren't evaluated and no alerts are triggered.

   > If you disable the alert rule for a current alert, you must wait a few minutes for the alert to no longer display as an active alert.

5. Select **Save**.

   Disabled appears in the **Status** column.

Remove custom alert rules

You can remove a custom alert rule if you no longer want to use it.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Manage alerts or Root access permission.

Steps

1. Select **ALERTS** > **Rules**.

   The Alert Rules page appears.

2. Select the radio button for the custom alert rule you want to remove.

   You can't remove a default alert rule.

3. Select **Remove custom rule**.

   A confirmation dialog box appears.

4. Select **OK** to remove the alert rule.

   Any active instances of the alert will be resolved within 10 minutes.

Manage alert notifications

Set up SNMP notifications for alerts

If you want StorageGRID to send SNMP notifications when alerts occur, you must enable the StorageGRID SNMP agent and configure one or more trap destinations.

You can use the **CONFIGURATION** > **Monitoring** > **SNMP agent** option in the Grid Manager or the SNMP endpoints for the Grid Management API to enable and configure the StorageGRID SNMP agent. The SNMP agent supports all three versions of the SNMP protocol.

To learn how to configure the SNMP agent, see Use SNMP monitoring.

After you configure the StorageGRID SNMP agent, two types of event-driven notifications can be sent:

- Traps are notifications sent by the SNMP agent that don't 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 when 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. See Silence alert notifications.

If your StorageGRID deployment includes multiple Admin Nodes, the primary Admin Node is the preferred sender for alert notifications, AutoSupport packages, SNMP traps and informs, and legacy alarm notifications. If the primary Admin Node becomes unavailable, notifications are temporarily sent by other Admin Nodes. See What is an Admin Node?.

Set up email notifications for alerts

If you want email notifications to be sent when alerts occur, you must provide information about your SMTP server. You must also enter email addresses for the recipients of alert notifications.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Manage alerts or Root access permission.

About this task

Because alarms and alerts are independent systems, the email setup used for alert notifications is not used for alarm notifications and AutoSupport packages. However, you can use the same email server for all notifications.

If your StorageGRID deployment includes multiple Admin Nodes, the primary Admin Node is the preferred sender for alert notifications, AutoSupport packages, SNMP traps and informs, and legacy alarm notifications. If the primary Admin Node becomes unavailable, notifications are temporarily sent by other Admin Nodes. See What is an Admin Node?.

Steps

1. Select **ALERTS** > **Email setup**.

   The Email Setup page appears.

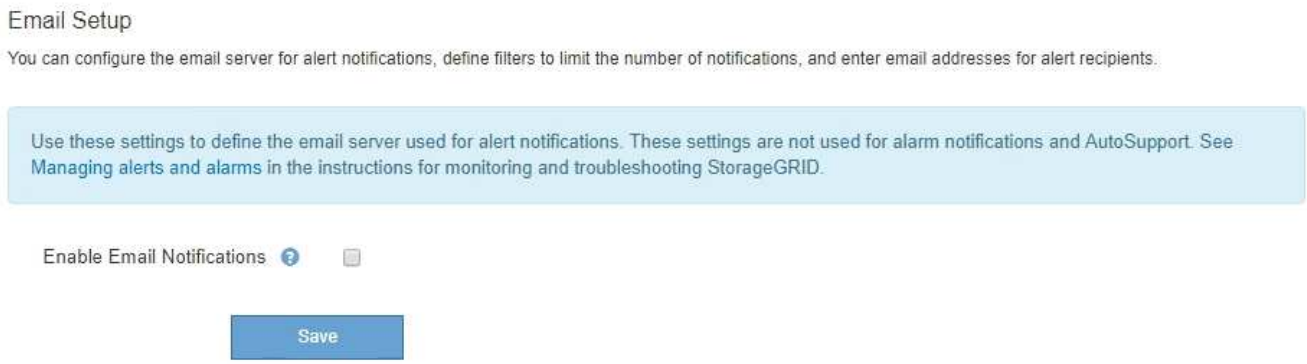

2. Select the **Enable Email Notifications** checkbox to indicate that you want notification emails to be sent when alerts reach configured thresholds.

   The Email (SMTP) Server, Transport Layer Security (TLS), Email Addresses, and Filters sections appear.

3. In the Email (SMTP) Server section, enter the information StorageGRID needs to access your SMTP server.

   If your SMTP server requires authentication, you must provide both a username and a password.

Field	Enter
Mail Server	The fully qualified domain name (FQDN) or IP address of the SMTP server.

Field	Enter
Port	The port used to access the SMTP server. Must be between 1 and 65535.
Username (optional)	If your SMTP server requires authentication, enter the username to authenticate with.
Password (optional)	If your SMTP server requires authentication, enter the password to authenticate with.

4. In the Email Addresses section, enter email addresses for the sender and for each recipient.

   a. For the **Sender Email Address**, specify a valid email address to use as the From address for alert notifications.

      For example: `storagegrid-alerts@example.com`

   b. In the Recipients section, enter an email address for each email list or person who should receive an email when an alert occurs.

      Select the plus icon ✚ to add recipients.

5. If Transport Layer Security (TLS) is required for communications with the SMTP server, select **Require TLS** in the Transport Layer Security (TLS) section.

   a. In the **CA Certificate** field, provide the CA certificate that will be used to verify the identify of the SMTP server.

      You can copy and paste the contents into this field, or select **Browse** and select the file.

      You must provide a single file that contains the certificates from each intermediate issuing certificate

authority (CA). The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.

b. Select the **Send Client Certificate** checkbox if your SMTP email server requires email senders to provide client certificates for authentication.

c. In the **Client Certificate** field, provide the PEM-encoded client certificate to send to the SMTP server.

   You can copy and paste the contents into this field, or select **Browse** and select the file.

d. In the **Private Key** field, enter the private key for the client certificate in unencrypted PEM encoding.

   You can copy and paste the contents into this field, or select **Browse** and select the file.

> If you need to edit the email setup, select the pencil icon to update this field.

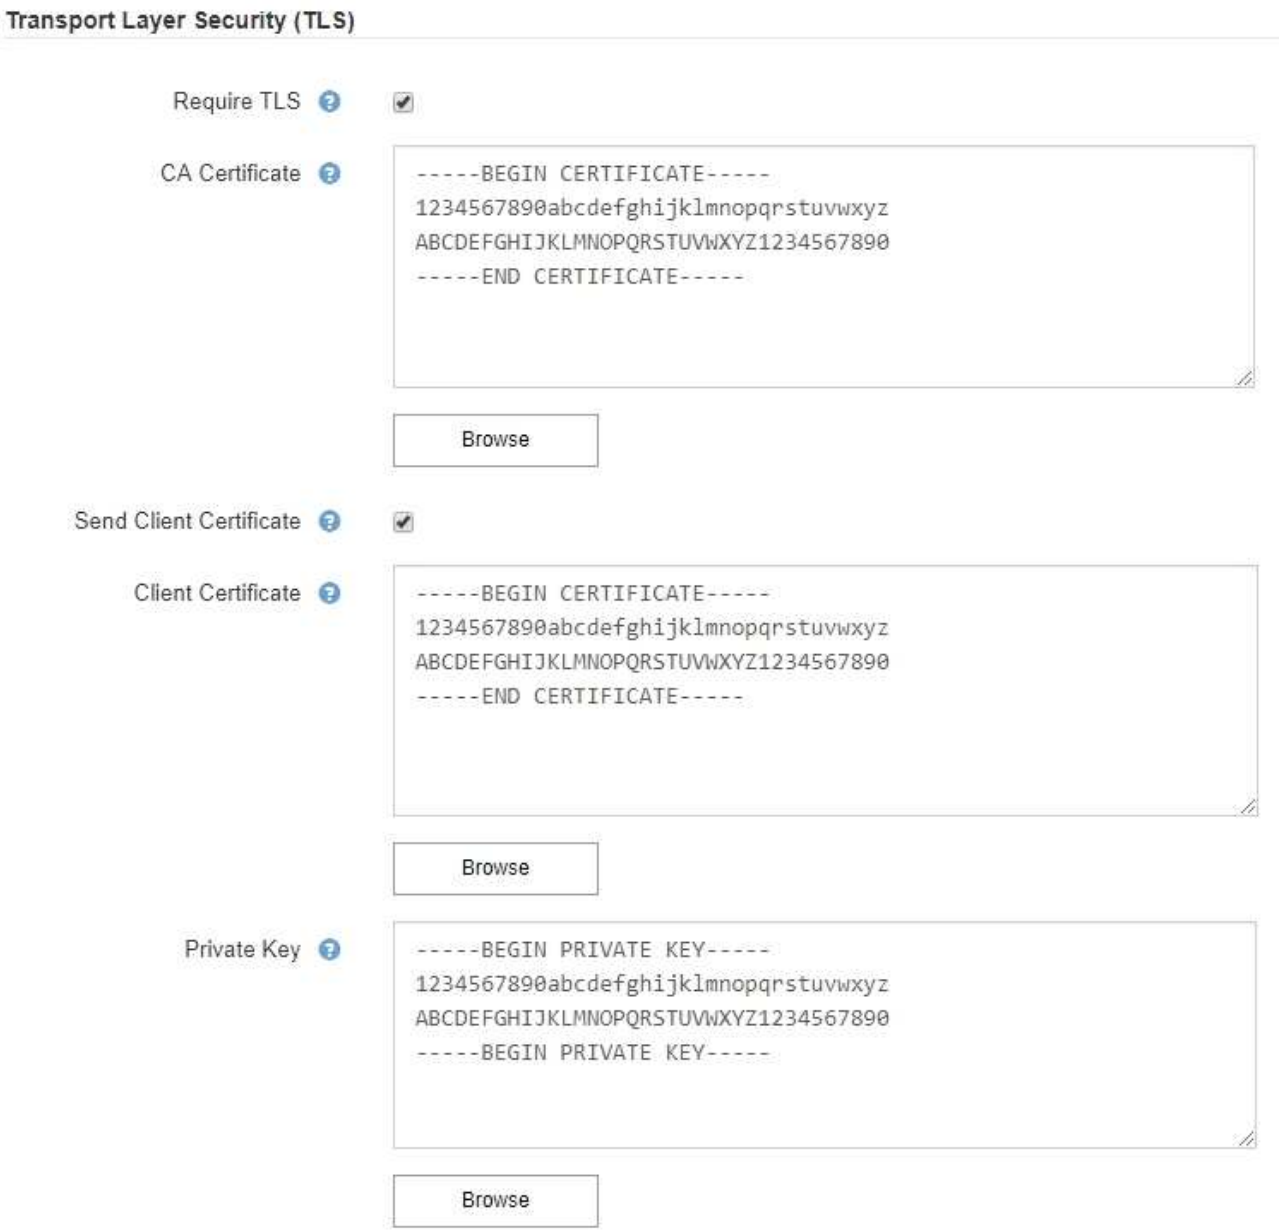

6. In the Filters section, select which alert severity levels should result in email notifications, unless the rule for a specific alert has been silenced.

Severity	Description
Minor, major, critical	An email notification is sent when the minor, major, or critical condition for an alert rule is met.
Major, critical	An email notification is sent when the major or critical condition for an alert rule is met. Notifications aren't sent for minor alerts.
Critical only	An email notification is sent only when the critical condition for an alert rule is met. Notifications aren't sent for minor or major alerts.

7. When you are ready to test your email settings, perform these steps:
   a. Select **Send Test Email**.

      A confirmation message appears, indicating that a test email was sent.

   b. Check the inboxes of all email recipients and confirm that a test email was received.

      If the email is not received within a few minutes or if the **Email notification failure** alert is triggered, check your settings and try again.

   c. Sign in to any other Admin Nodes and send a test email to verify connectivity from all sites.

      When you test alert notifications, you must sign in to every Admin Node to verify connectivity. This is in contrast to testing AutoSupport packages and legacy alarm notifications, where all Admin Nodes send the test email.

8. Select **Save**.

   Sending a test email does not save your settings. You must select **Save**.

   The email settings are saved.

Information included in alert email notifications

After you configure the SMTP email server, email notifications are sent to the designated recipients when an alert is triggered, unless the alert rule is suppressed by a silence. See Silence alert notifications.

Email notifications include the following information:

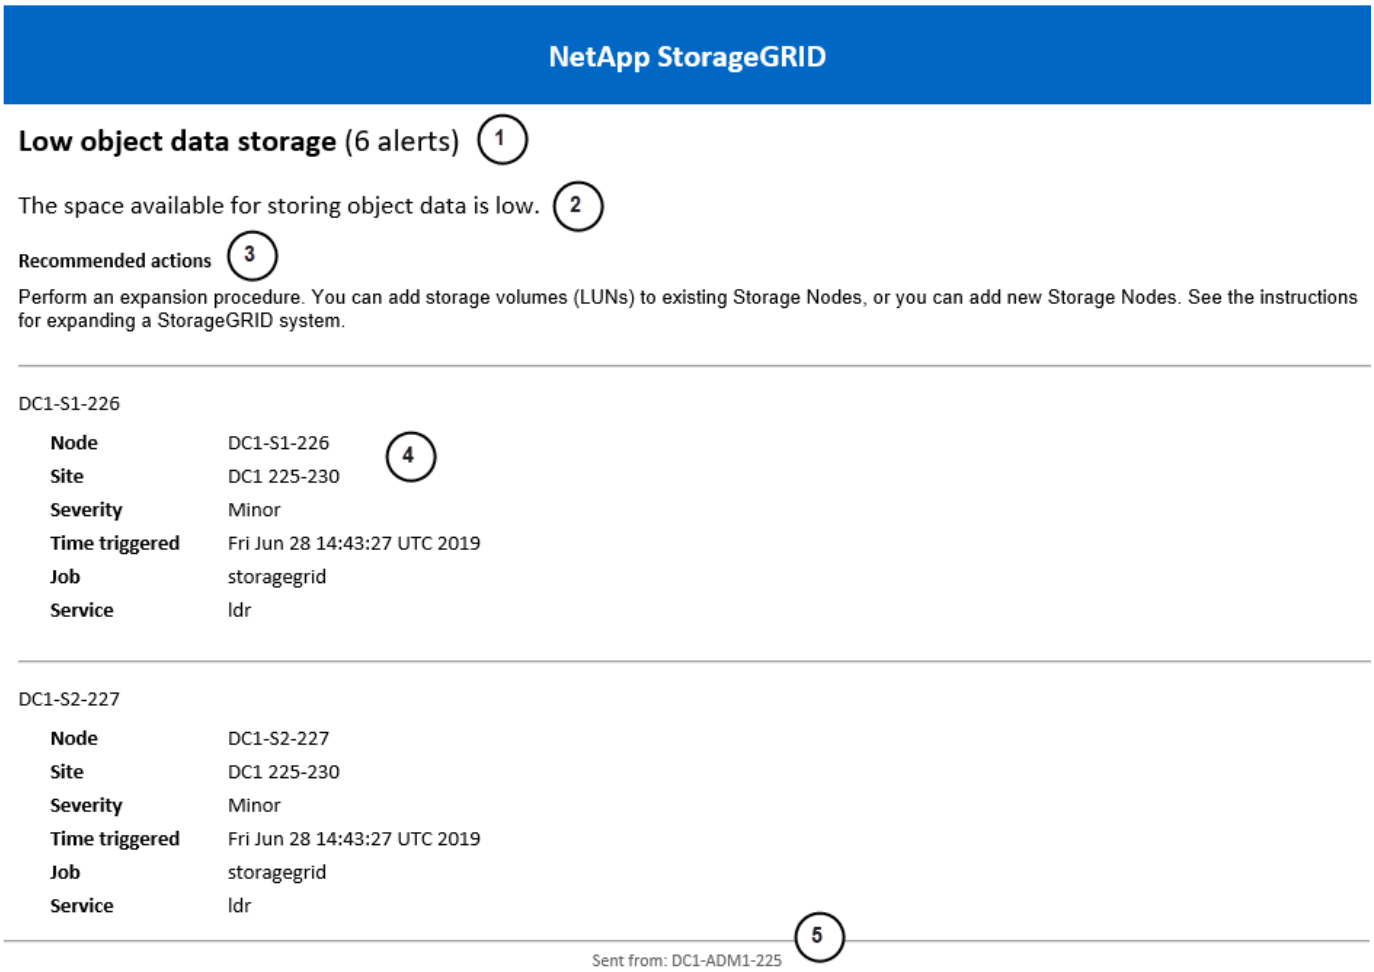

Callout	Description
1	The name of the alert, followed by the number of active instances of this alert.
2	The description of the alert.
3	Any recommended actions for the alert.
4	Details about each active instance of the alert, including the node and site affected, the alert severity, the UTC time when the alert rule was triggered, and the name of the affected job and service.
5	The hostname of the Admin Node that sent the notification.

How alerts are grouped

To prevent an excessive number of email notifications from being sent when alerts are triggered, StorageGRID attempts to group multiple alerts in the same notification.

Refer to the following table for examples of how StorageGRID groups multiple alerts in email notifications.

Behavior	Example
Each alert notification applies only to alerts that have the same name. If two alerts with different names are triggered at the same time, two email notifications are sent.	• Alert A is triggered on two nodes at the same time. Only one notification is sent. • Alert A is triggered on node 1, and Alert B is triggered on node 2 at the same time. Two notifications are sent—one for each alert.
For a specific alert on a specific node, if the thresholds are reached for more than one severity, a notification is sent only for the most severe alert.	• Alert A is triggered and the minor, major, and critical alert thresholds are reached. One notification is sent for the critical alert.
The first time an alert is triggered, StorageGRID waits 2 minutes before sending a notification. If other alerts with the same name are triggered during that time, StorageGRID groups all of the alerts in the initial notification.	1. Alert A is triggered on node 1 at 08:00. No notification is sent. 2. Alert A is triggered on node 2 at 08:01. No notification is sent. 3. At 08:02, a notification is sent to report both instances of the alert.
If an another alert with the same name is triggered, StorageGRID waits 10 minutes before sending a new notification. The new notification reports all active alerts (current alerts that have not been silenced), even if they were reported previously.	1. Alert A is triggered on node 1 at 08:00. A notification is sent at 08:02. 2. Alert A is triggered on node 2 at 08:05. A second notification is sent at 08:15 (10 minutes later). Both nodes are reported.
If there are multiple current alerts with the same name and one of those alerts is resolved, a new notification is not sent if the alert reoccurs on the node for which the alert was resolved.	1. Alert A is triggered for node 1. A notification is sent. 2. Alert A is triggered for node 2. A second notification is sent. 3. Alert A is resolved for node 2, but it remains active for node 1. 4. Alert A is triggered again for node 2. No new notification is sent because the alert is still active for node 1.
StorageGRID continues to send email notifications once every 7 days until all instances of the alert are resolved or the alert rule is silenced.	1. Alert A is triggered for node 1 on March 8. A notification is sent. 2. Alert A is not resolved or silenced. Additional notifications are sent on March 15, March 22, March 29, and so on.

Troubleshoot alert email notifications

If the **Email notification failure** alert is triggered or you are unable to receive the test alert email notification, follow these steps to resolve the issue.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Manage alerts or Root access permission.

Steps

1. Verify your settings.
   a. Select **ALERTS** > **Email setup**.
   b. Verify that the Email (SMTP) Server settings are correct.
   c. Verify that you have specified valid email addresses for the recipients.
2. Check your spam filter, and make sure that the email was not sent to a junk folder.
3. Ask your email administrator to confirm that emails from the sender address aren't being blocked.
4. Collect a log file for the Admin Node, and then contact technical support.

   Technical support can use the information in the logs to help determine what went wrong. For example, the prometheus.log file might show an error when connecting to the server you specified.

   See Collect log files and system data.

Silence alert notifications

Optionally, you can configure silences to temporarily suppress alert notifications.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Manage alerts or Root access permission.

About this task

You can silence alert rules on the entire grid, a single site, or a single node and for one or more severities. Each silence suppresses all notifications for a single alert rule or for all alert rules.

If you have enabled the SNMP agent, silences also suppress SNMP traps and informs.

> Be careful when deciding to silence an alert rule. If you silence an alert, you might not detect an underlying problem until it prevents a critical operation from completing.

> Because alarms and alerts are independent systems, you can't use this functionality to suppress alarm notifications.

Steps

1. Select **ALERTS** > **Silences**.

   The Silences page appears.

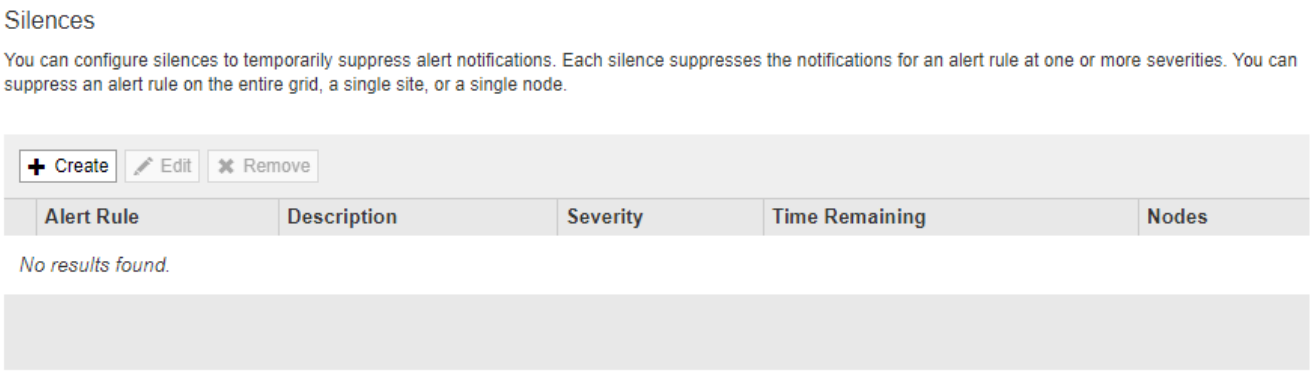

2. Select **Create**.

   The Create Silence dialog box appears.

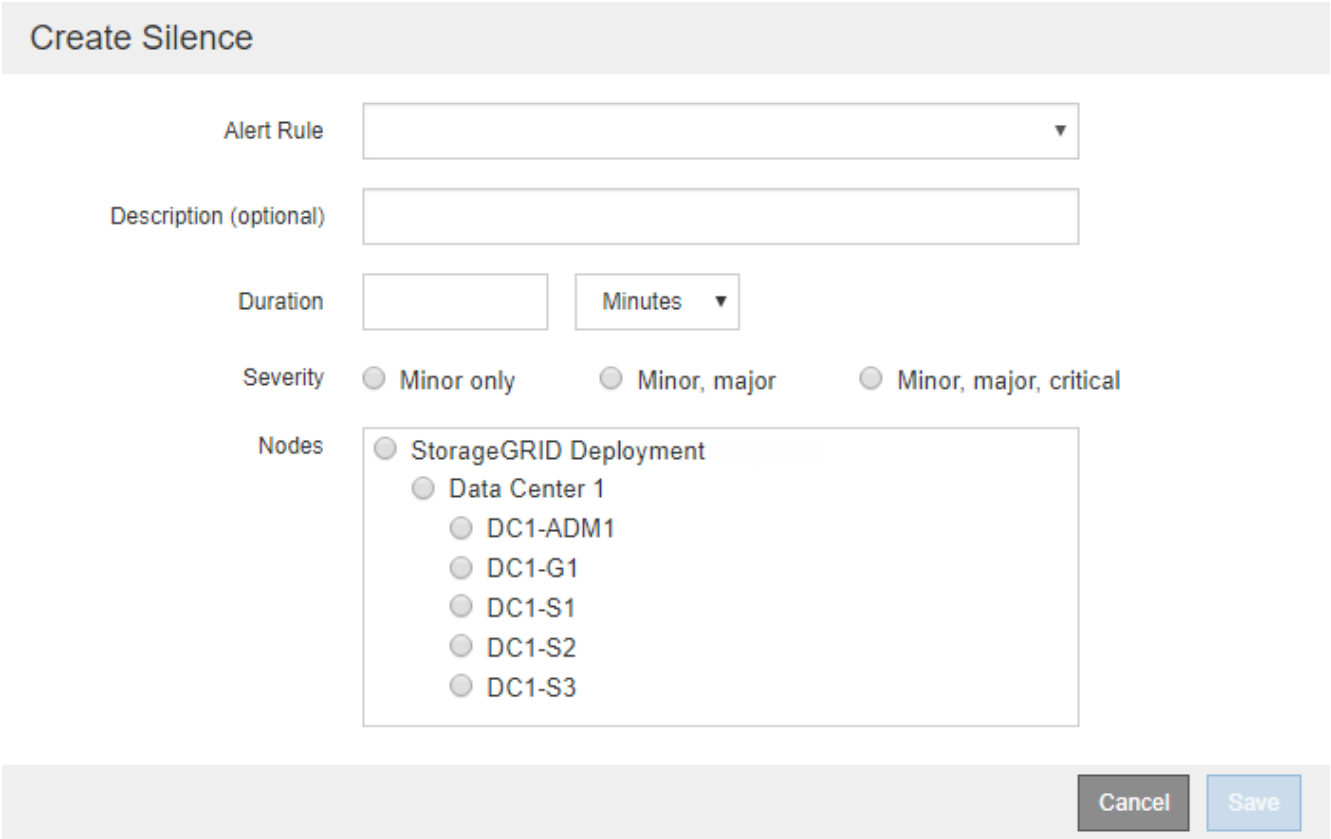

3. Select or enter the following information:

Field	Description
Alert Rule	The name of the alert rule you want to silence. You can select any default or custom alert rule, even if the alert rule is disabled. **Note:** Select **All rules** if you want to silence all alert rules using the criteria specified in this dialog box.
Description	Optionally, a description of the silence. For example, describe the purpose of this silence.

Field	Description
Duration	How long you want this silence to remain in effect, in minutes, hours, or days. A silence can be in effect from 5 minutes to 1,825 days (5 years). **Note:** You should not silence an alert rule for an extended amount of time. If an alert rule is silenced, you might not detect an underlying problem until it prevents a critical operation from completing. However, you might need to use an extended silence if an alert is triggered by a specific, intentional configuration, such as might be the case for the **Services appliance link down** alerts and the **Storage appliance link down** alerts.
Severity	Which alert severity or severities should be silenced. If the alert is triggered at one of the selected severities, no notifications are sent.
Nodes	Which node or nodes you want this silence to apply to. You can suppress an alert rule or all rules on the entire grid, a single site, or a single node. If you select the entire grid, the silence applies to all sites and all nodes. If you select a site, the silence applies only to the nodes at that site. **Note:** You can’t select more than one node or more than one site for each silence. You must create additional silences if you want to suppress the same alert rule on more than one node or more than one site at one time.

4. Select **Save**.
5. If you want to modify or end a silence before it expires, you can edit or remove it.

Option	Description
Edit a silence	a. Select **ALERTS** > **Silences**. b. From the table, select the radio button for the silence you want to edit. c. Select **Edit**. d. Change the description, the amount of time remaining, the selected severities, or the affected node. e. Select **Save**.
Remove a silence	a. Select **ALERTS** > **Silences**. b. From the table, select the radio button for the silence you want to remove. c. Select **Remove**. d. Select **OK** to confirm you want to remove this silence. **Note**: Notifications will now be sent when this alert is triggered (unless suppressed by another silence). If this alert is currently triggered, it might take few minutes for email or SNMP notifications to be sent and for the Alerts page to update.

Related information

- Configure the SNMP agent

Alerts reference

This reference lists the default alerts that appear in the Grid Manager. Recommended actions are in the alert message you receive.

As required, you can create custom alert rules to fit your system management approach.

Some of the default alerts use Prometheus metrics.

Appliance alerts

Alert name	Description
Appliance battery expired	The battery in the appliance's storage controller has expired.
Appliance battery failed	The battery in the appliance's storage controller has failed.
Appliance battery has insufficient learned capacity	The battery in the appliance's storage controller has insufficient learned capacity.
Appliance battery near expiration	The battery in the appliance's storage controller is nearing expiration.
Appliance battery removed	The battery in the appliance's storage controller is missing.
Appliance battery too hot	The battery in the appliance's storage controller is overheated.
Appliance BMC communication error	Communication with the baseboard management controller (BMC) has been lost.
Appliance cache backup device failed	A persistent cache backup device has failed.
Appliance cache backup device insufficient capacity	There is insufficient cache backup device capacity.
Appliance cache backup device write-protected	A cache backup device is write-protected.
Appliance cache memory size mismatch	The two controllers in the appliance have different cache sizes.
Appliance compute controller chassis temperature too high	The temperature of the compute controller in a StorageGRID appliance has exceeded a nominal threshold.
Appliance compute controller CPU temperature too high	The temperature of the CPU in the compute controller in a StorageGRID appliance has exceeded a nominal threshold.

Alert name	Description
Appliance compute controller needs attention	A hardware fault has been detected in the compute controller of a StorageGRID appliance.
Appliance compute controller power supply A has a problem	Power supply A in the compute controller has a problem.
Appliance compute controller power supply B has a problem	Power supply B in the compute controller has a problem.
Appliance compute hardware monitor service stalled	The service that monitors storage hardware status has stalled.
Appliance DAS drive exceeding limit for data written per day	An excessive amount of data is being written to a drive each day, which might void its warranty.
Appliance DAS drive fault detected	A problem was detected with a direct-attached storage (DAS) drive in the appliance.
Appliance DAS drive locator light on	The drive locator light for one or more direct-attached storage (DAS) drives in an appliance Storage Node is on.
Appliance DAS drive rebuilding	A direct-attached storage (DAS) drive is rebuilding. This is expected if it was recently replaced or removed/reinserted.
Appliance fan fault detected	A problem with a fan unit in the appliance was detected.
Appliance Fibre Channel fault detected	A Fibre Channel link problem has been detected between the appliance storage controller and compute controller
Appliance Fibre Channel HBA port failure	A Fibre Channel HBA port is failing or has failed.
Appliance flash cache drives non-optimal	The drives used for the SSD cache are non-optimal.
Appliance interconnect/battery canister removed	The interconnect/battery canister is missing.
Appliance LACP port missing	A port on a StorageGRID appliance is not participating in the LACP bond.
Appliance NIC fault detected	A problem with a network interface card (NIC) in the appliance was detected.

Alert name	Description
Appliance overall power supply degraded	The power of a StorageGRID appliance has deviated from the recommended operating voltage.
Appliance SSD critical warning	An appliance SSD is reporting a critical warning.
Appliance storage controller A failure	Storage controller A in a StorageGRID appliance has failed.
Appliance storage controller B failure	Storage controller B in a StorageGRID appliance has failed.
Appliance storage controller drive failure	One or more drives in a StorageGRID appliance has failed or is not optimal.
Appliance storage controller hardware issue	SANtricity software is reporting "Needs attention" for a component in a StorageGRID appliance.
Appliance storage controller power supply A failure	Power supply A in a StorageGRID appliance has deviated from the recommended operating voltage.
Appliance storage controller power supply B failure	Power supply B in a StorageGRID appliance has deviated from the recommended operating voltage.
Appliance storage hardware monitor service stalled	The service that monitors storage hardware status has stalled.
Appliance storage shelves degraded	The status of one of the components in the storage shelf for a storage appliance is degraded.
Appliance temperature exceeded	The nominal or maximum temperature for the appliance’s storage controller has been exceeded.
Appliance temperature sensor removed	A temperature sensor has been removed.
Appliance UEFI secure boot error	An appliance has not been booted securely.
Disk I/O is very slow	Very slow disk I/O may be impacting grid performance.
Storage appliance fan fault detected	A problem with a fan unit in the storage controller for an appliance was detected.
Storage appliance storage connectivity degraded	There is a problem with one or more connections between the compute controller and storage controller.

Alert name	Description
Storage device inaccessible	A storage device cannot be accessed.

Audit and syslog alerts

Alert name	Description
Audit logs are being added to the in-memory queue	Node cannot send logs to the local syslog server and the in-memory queue is filling up.
External syslog server forwarding error	Node cannot forward logs to the external syslog server.
Large audit queue	The disk queue for audit messages is full. If this condition is not addressed, S3 or Swift operations might fail.
Logs are being added to the on-disk queue	Node cannot forward logs to the external syslog server and the on-disk queue is filling up.

Bucket alerts

Alert name	Description
FabricPool bucket has unsupported bucket consistency setting	A FabricPool bucket uses the Available or Strong-site consistency level, which is not supported.

Cassandra alerts

Alert name	Description
Cassandra auto-compactor error	The Cassandra auto-compactor has experienced an error.
Cassandra auto-compactor metrics out of date	The metrics that describe the Cassandra auto-compactor are out of date.
Cassandra communication error	The nodes that run the Cassandra service are having trouble communicating with each other.
Cassandra compactions overloaded	The Cassandra compaction process is overloaded.
Cassandra oversize write error	An internal StorageGRID process sent a write request to Cassandra that was too large.
Cassandra repair metrics out of date	The metrics that describe Cassandra repair jobs are out of date.

Alert name	Description
Cassandra repair progress slow	The progress of Cassandra database repairs is slow.
Cassandra repair service not available	The Cassandra repair service is not available.
Cassandra table corruption	Cassandra has detected table corruption. Cassandra automatically restarts if it detects table corruption.

Cloud Storage Pool alerts

Alert name	Description
Cloud Storage Pool connectivity error	The health check for Cloud Storage Pools detected one or more new errors.

Cross-grid replication alerts

Alert name	Description
Cross-grid replication permanent failure	A cross-grid replication error occurred that requires user intervention to resolve.
Cross-grid replication resources unavailable	Cross-grid replication requests are pending because a resource is unavailable.

DHCP alerts

Alert name	Description
DHCP lease expired	The DHCP lease on a network interface has expired.
DHCP lease expiring soon	The DHCP lease on a network interface is expiring soon.
DHCP server unavailable	The DHCP server is unavailable.

Debug and trace alerts

Alert name	Description
Debug performance impact	When debug mode is enabled, system performance might be negatively impacted.
Trace configuration enabled	When trace configuration is enabled, system performance might be negatively impacted.

Email and AutoSupport alerts

Alert name	Description
AutoSupport message failed to send	The most recent AutoSupport message failed to send.
Email notification failure	The email notification for an alert could not be sent.

Erasure coding (EC) alerts

Alert name	Description
EC rebalance failure	The EC rebalance procedure has failed or has been stopped.
EC repair failure	A repair job for EC data has failed or has been stopped.
EC repair stalled	A repair job for EC data has stalled.

Expiration of certificates alerts

Alert name	Description
Admin Proxy CA certificate expiration	One or more certificates in the admin proxy server CA bundle is about to expire.
Expiration of client certificate	One or more client certificates are about to expire.
Expiration of global server certificate for S3 and Swift	The global server certificate for S3 and Swift is about to expire.
Expiration of load balancer endpoint certificate	One or more load balancer endpoint certificates are about to expire.
Expiration of server certificate for Management interface	The server certificate used for the management interface is about to expire.
External syslog CA certificate expiration	The certificate authority (CA) certificate used to sign the external syslog server certificate is about to expire.
External syslog client certificate expiration	The client certificate for an external syslog server is about to expire.
External syslog server certificate expiration	The server certificate presented by the external syslog server is about to expire.

Grid Network alerts

Alert name	Description
Grid Network MTU mismatch	The MTU setting for the Grid Network interface (eth0) differs significantly across nodes in the grid.

Grid federation alerts

Alert name	Description
Expiration of grid federation certificate	One or more grid federation certificates are about to expire.
Grid federation connection failure	The grid federation connection between the local and remote grid is not working.

High usage or high latency alerts

Alert name	Description
High Java heap use	A high percentage of Java heap space is being used.
High latency for metadata queries	The average time for Cassandra metadata queries is too long.

Identity federation alerts

Alert name	Description
Identity federation synchronization failure	Unable to synchronize federated groups and users from the identity source.
Identity federation synchronization failure for a tenant	Unable to synchronize federated groups and users from the identity source configured by a tenant.

Information lifecycle management (ILM) alerts

Alert name	Description
ILM placement unachievable	A placement instruction in an ILM rule cannot be achieved for certain objects.
ILM scan period too long	The time required to scan, evaluate, and apply ILM to objects is too long.
ILM scan rate low	The ILM scan rate is set to less than 100 objects/second.

Key management server (KMS) alerts

Alert name	Description
KMS CA certificate expiration	The certificate authority (CA) certificate used to sign the key management server (KMS) certificate is about to expire.
KMS client certificate expiration	The client certificate for a key management server is about to expire
KMS configuration failed to load	The configuration for the key management server exists but failed to load.
KMS connectivity error	An appliance node could not connect to the key management server for its site.
KMS encryption key name not found	The configured key management server does not have an encryption key that matches the name provided.
KMS encryption key rotation failed	All appliance volumes were successfully decrypted, but one or more volumes could not rotate to the latest key.
KMS is not configured	No key management server exists for this site.
KMS key failed to decrypt an appliance volume	One or more volumes on an appliance with node encryption enabled could not be decrypted with the current KMS key.
KMS server certificate expiration	The server certificate used by the key management server (KMS) is about to expire.

Local clock offset alerts

Alert name	Description
Local clock large time offset	The offset between local clock and Network Time Protocol (NTP) time is too large.

Low memory or low space alerts

Alert name	Description
Low audit log disk capacity	The space available for audit logs is low. If this condition is not addressed, S3 or Swift operations might fail.
Low available node memory	The amount of RAM available on a node is low.
Low free space for storage pool	The space available for storing object data in the Storage Node is low.
Low installed node memory	The amount of installed memory on a node is low.

Alert name	Description
Low metadata storage	The space available for storing object metadata is low.
Low metrics disk capacity	The space available for the metrics database is low.
Low object data storage	The space available for storing object data is low.
Low read-only watermark override	The Storage Volume Soft Read-Only Watermark Override is less than the minimum optimized watermark for a Storage Node.
Low root disk capacity	The space available on the root disk is low.
Low system data capacity	The space available for /var/local is low. If this condition is not addressed, S3 or Swift operations might fail.
Low tmp directory free space	The space available in the /tmp directory is low.

Node or node network alerts

Alert name	Description
Admin Network receive usage	The receive usage on the Admin Network is high.
Admin Network transmit usage	The transmit usage on the Admin Network is high.
Firewall configuration failure	Failed to apply firewall configuration.
Management interface endpoints in fallback mode	All management interface endpoints have been falling back to the default ports for too long.
Node network connectivity error	Errors have occurred while transferring data between nodes.
Node network reception frame error	A high percentage of the network frames received by a node had errors.
Node not in sync with NTP server	The node is not in sync with the network time protocol (NTP) server.
Node not locked with NTP server	The node is not locked to a network time protocol (NTP) server.
Non-appliance node network down	One or more network devices are down or disconnected.
Services appliance link down on Admin Network	The appliance interface to the Admin Network (eth1) is down or disconnected.
Services appliance link down on Admin Network port 1	The Admin Network port 1 on the appliance is down or disconnected.

Alert name	Description
Services appliance link down on Client Network	The appliance interface to the Client Network (eth2) is down or disconnected.
Services appliance link down on network port 1	Network port 1 on the appliance is down or disconnected.
Services appliance link down on network port 2	Network port 2 on the appliance is down or disconnected.
Services appliance link down on network port 3	Network port 3 on the appliance is down or disconnected.
Services appliance link down on network port 4	Network port 4 on the appliance is down or disconnected.
Storage appliance link down on Admin Network	The appliance interface to the Admin Network (eth1) is down or disconnected.
Storage appliance link down on Admin Network port 1	The Admin Network port 1 on the appliance is down or disconnected.
Storage appliance link down on Client Network	The appliance interface to the Client Network (eth2) is down or disconnected.
Storage appliance link down on network port 1	Network port 1 on the appliance is down or disconnected.
Storage appliance link down on network port 2	Network port 2 on the appliance is down or disconnected.
Storage appliance link down on network port 3	Network port 3 on the appliance is down or disconnected.
Storage appliance link down on network port 4	Network port 4 on the appliance is down or disconnected.
Storage Node not in desired storage state	The LDR service on a Storage Node cannot transition to the desired state because of an internal error or volume related issue
TCP connection usage	The number of TCP connections on this node is approaching the maximum number that can be tracked.
Unable to communicate with node	One or more services are unresponsive, or the node cannot be reached.
Unexpected node reboot	A node rebooted unexpectedly within the last 24 hours.

Object alerts

Alert name	Description
Object existence check failed	The object existence check job has failed.
Object existence check stalled	The object existence check job has stalled.
Objects lost	One or more objects have been lost from the grid.
S3 PUT object size too large	A client is attempting a PUT Object operation that exceeds S3 size limits.
Unidentified corrupt object detected	A file was found in replicated object storage that could not be identified as a replicated object.

Platform services alerts

Alert name	Description
Platform Services pending request capacity low	The number of Platform Services pending requests is approaching capacity.
Platform services unavailable	Too few Storage Nodes with the RSM service are running or available at a site.

Storage volume alerts

Alert name	Description
Storage volume needs attention	A storage volume is offline and needs attention.
Storage volume needs to be restored	A storage volume has been recovered and needs to be restored.
Storage volume offline	A storage volume has been offline for more than 5 minutes, possibly because the node rebooted during the volume formatting step.
Volume Restoration failed to start replicated data repair	Replicated data repair for a repaired volume couldn’t be started automatically.

StorageGRID services alerts

Alert name	Description
nginx service using backup configuration	The configuration of the nginx service is invalid. The previous configuration is now being used.

Alert name	Description
nginx-gw service using backup configuration	The configuration of the nginx-gw service is invalid. The previous configuration is now being used.
Reboot required to disable FIPS	The security policy does not require FIPS mode, but the NetApp Cryptographic Security Module is enabled.
Reboot required to enable FIPS	The security policy requires FIPS mode, but the NetApp Cryptographic Security Module is disabled.
SSH service using backup configuration	The configuration of the SSH service is invalid. The previous configuration is now being used.

Tenant alerts

Alert name	Description
Tenant quota usage high	A high percentage of quota space is being used. This rule is disabled by default because it might cause too many notifications.

Commonly used Prometheus metrics

Refer to this list of commonly used Prometheus metrics to better understand conditions in the default alert rules or to construct the conditions for custom alert rules.

You can also obtain a complete list of all metrics.

For details on the syntax of Prometheus queries, see Querying Prometheus.

What are Prometheus metrics?

Prometheus metrics are time series measurements. The Prometheus service on Admin Nodes collects these metrics from the services on all nodes. Metrics are stored on each Admin Node until the space reserved for Prometheus data is full. When the `/var/local/mysql_ibdata/` volume reaches capacity, the oldest metrics are deleted first.

Where are Prometheus metrics used?

The metrics collected by Prometheus are used in several 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.

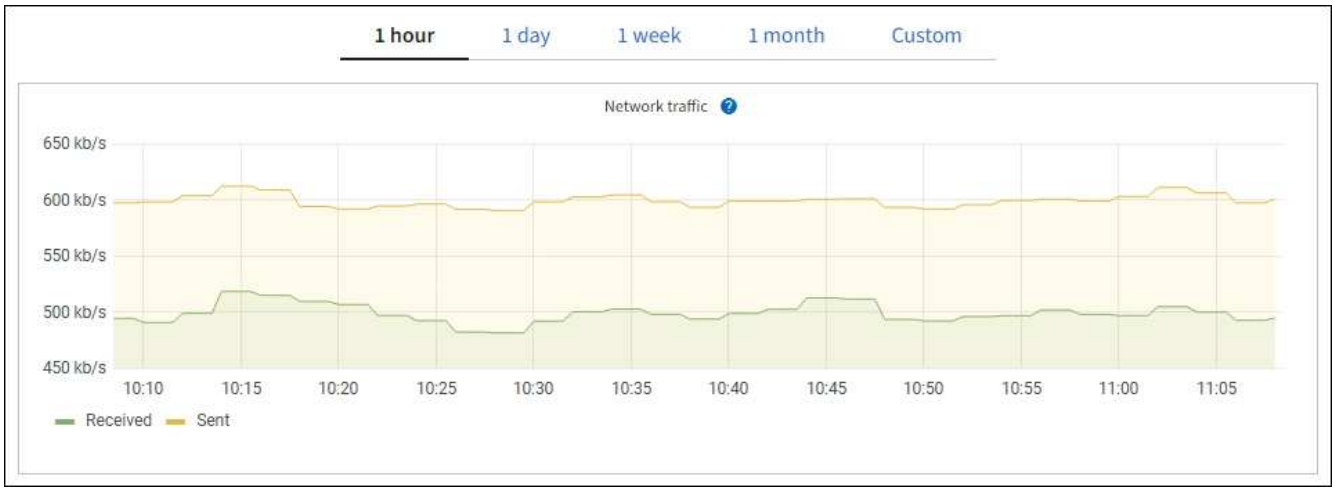

- **Alerts**: Alerts are triggered at specific severity levels when alert rule conditions that use Prometheus metrics evaluate as true.
- **Grid Management API**: You can use Prometheus metrics in custom alert rules or with external automation tools to monitor your StorageGRID system. A complete list 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.

  > 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 several 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.

List of most common metrics

The following list contains the most commonly used Prometheus metrics.

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

alertmanager_notifications_failed_total

The total number of failed alert notifications.

node_filesystem_avail_bytes

The amount of file system space available to non-root users in bytes.

node_memory_MemAvailable_bytes

Memory information field MemAvailable_bytes.

node_network_carrier

Carrier value of `/sys/class/net/`*`iface`*.

node_network_receive_errs_total

Network device statistic `receive_errs`.

node_network_transmit_errs_total

Network device statistic `transmit_errs`.

storagegrid_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.

storagegrid_appliance_compute_controller_hardware_status

The status of the compute controller hardware in an appliance.

storagegrid_appliance_failed_disks

For the storage controller in an appliance, the number of drives that aren't optimal.

storagegrid_appliance_storage_controller_hardware_status

The overall status of the storage controller hardware in an appliance.

storagegrid_content_buckets_and_containers

The total number of S3 buckets and Swift containers known by this Storage Node.

storagegrid_content_objects

The total number of S3 and Swift data objects known by this Storage Node. Count is valid only for data objects created by client applications that interface with the system through S3 or Swift.

storagegrid_content_objects_lost

The total number of objects this service detects as missing from the StorageGRID system. Action should be taken to determine the cause of the loss and if recovery is possible.

Troubleshoot lost and missing object data

storagegrid_http_sessions_incoming_attempted

The total number of HTTP sessions that have been attempted to a Storage Node.

storagegrid_http_sessions_incoming_currently_established

The number of HTTP sessions that are currently active (open) on the Storage Node.

storagegrid_http_sessions_incoming_failed

The total number of HTTP sessions that failed to complete successfully, either due to a malformed HTTP request or a failure while processing an operation.

storagegrid_http_sessions_incoming_successful

The total number of HTTP sessions that have completed successfully.

storagegrid_ilm_awaiting_background_objects

The total number of objects on this node awaiting ILM evaluation from the scan.

storagegrid_ilm_awaiting_client_evaluation_objects_per_second

The current rate at which objects are evaluated against the ILM policy on this node.

storagegrid_ilm_awaiting_client_objects

The total number of objects on this node awaiting ILM evaluation from client operations (for example, ingest).

storagegrid_ilm_awaiting_total_objects

The total number of objects awaiting ILM evaluation.

storagegrid_ilm_scan_objects_per_second

The rate at which objects owned by this node are scanned and queued for ILM.

storagegrid_ilm_scan_period_estimated_minutes

The estimated time to complete a full ILM scan on this node.

Note: A full scan does not guarantee that ILM has been applied to all objects owned by this node.

storagegrid_load_balancer_endpoint_cert_expiry_time

The expiration time of the load balancer endpoint certificate in seconds since the epoch.

storagegrid_metadata_queries_average_latency_milliseconds

The average time required to run a query against the metadata store through this service.

storagegrid_network_received_bytes

The total amount of data received since installation.

storagegrid_network_transmitted_bytes

The total amount of data sent since installation.

storagegrid_node_cpu_utilization_percentage

The percentage of available CPU time currently being used by this service. Indicates how busy the service is. The amount of available CPU time depends on the number of CPUs for the server.

storagegrid_ntp_chosen_time_source_offset_milliseconds

Systematic offset of time provided by a chosen time source. Offset is introduced when the delay to reach a time source is not equal to the time required for the time source to reach the NTP client.

storagegrid_ntp_locked

The node is not locked to a Network Time Protocol (NTP) server.

storagegrid_s3_data_transfers_bytes_ingested

The total amount of data ingested from S3 clients to this Storage Node since the attribute was last reset.

storagegrid_s3_data_transfers_bytes_retrieved

The total amount of data retrieved by S3 clients from this Storage Node since the attribute was last reset.

storagegrid_s3_operations_failed

The total number of failed S3 operations (HTTP status codes 4xx and 5xx), excluding those caused by S3 authorization failure.

storagegrid_s3_operations_successful

The total number of successful S3 operations (HTTP status code 2xx).

storagegrid_s3_operations_unauthorized

The total number of failed S3 operations that are the result of an authorization failure.

storagegrid_servercertificate_management_interface_cert_expiry_days

The number of days before the Management Interface certificate expires.

storagegrid_servercertificate_storage_api_endpoints_cert_expiry_days

The number of days before the Object Storage API certificate expires.

storagegrid_service_cpu_seconds

The cumulative amount of time that the CPU has been used by this service since installation.

storagegrid_service_memory_usage_bytes

The amount of memory (RAM) currently in use by this service. This value is identical to that displayed by the Linux top utility as RES.

storagegrid_service_network_received_bytes

The total amount of data received by this service since installation.

storagegrid_service_network_transmitted_bytes

The total amount of data sent by this service.

storagegrid_service_restarts

The total number of times the service has been restarted.

storagegrid_service_runtime_seconds

The total amount of time that the service has been running since installation.

storagegrid_service_uptime_seconds

The total amount of time the service has been running since it was last restarted.

storagegrid_storage_state_current

The current state of the storage services. Attribute values are:

- 10 = Offline
- 15 = Maintenance
- 20 = Read-only
- 30 = Online

storagegrid_storage_status

The current status of the storage services. Attribute values are:

- 0 = No Errors
- 10 = In Transition
- 20 = Insufficient Free Space
- 30 = Volume(s) Unavailable
- 40 = Error

storagegrid_storage_utilization_data_bytes

An estimate of the total size of replicated and erasure-coded object data on the Storage Node.

storagegrid_storage_utilization_metadata_allowed_bytes

The total space on volume 0 of each Storage Node that is allowed for object metadata. This value is always less than the actual space reserved for metadata on a node, because a portion of the reserved space is required for essential database operations (such as compaction and repair) and future hardware and software upgrades.The allowed space for object metadata controls overall object capacity.

storagegrid_storage_utilization_metadata_bytes

The amount of object metadata on storage volume 0, in bytes.

storagegrid_storage_utilization_total_space_bytes

The total amount of storage space allocated to all object stores.

storagegrid_storage_utilization_usable_space_bytes

The total amount of object storage space remaining. Calculated by adding together the amount of available space for all object stores on the Storage Node.

storagegrid_swift_data_transfers_bytes_ingested

The total amount of data ingested from Swift clients to this Storage Node since the attribute was last reset.

storagegrid_swift_data_transfers_bytes_retrieved

The total amount of data retrieved by Swift clients from this Storage Node since the attribute was last reset.

storagegrid_swift_operations_failed

The total number of failed Swift operations (HTTP status codes 4xx and 5xx), excluding those caused by Swift authorization failure.

storagegrid_swift_operations_successful

The total number of successful Swift operations (HTTP status code 2xx).

storagegrid_swift_operations_unauthorized

The total number of failed Swift operations that are the result of an authorization failure (HTTP status codes 401, 403, 405).

storagegrid_tenant_usage_data_bytes

The logical size of all objects for the tenant.

storagegrid_tenant_usage_object_count

The number of objects for the tenant.

storagegrid_tenant_usage_quota_bytes

The maximum amount of logical space available for the tenant’s objects. If a quota metric is not provided, an unlimited amount of space is available.

Get a list of all metrics

To obtain the complete list of metrics, use the Grid Management API.

1. From the top of the Grid Manager, select the help icon and select **API documentation**.
2. Locate the **metrics** operations.

3. Execute the `GET /grid/metric-names` operation.
4. Download the results.

Manage alarms (legacy system)

Manage alarms (legacy system)

The StorageGRID alarm system is the legacy system used to identify trouble spots that sometimes occur during normal operation.

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

Alarm classes (legacy system)

A legacy alarm can belong to one of two mutually exclusive alarm classes.

- Default alarms are provided with each StorageGRID system and can't be modified. However, you can disable Default alarms or override them by defining Global Custom alarms.
- Global Custom alarms monitor the status of all services of a given type in the StorageGRID system. You can create a Global Custom alarm to override a Default alarm. You can also create a new Global Custom alarm. This can be useful for monitoring any customized conditions of your StorageGRID system.

Alarm triggering logic (legacy system)

A legacy alarm is triggered when a StorageGRID attribute reaches a threshold value that evaluates to true against a combination of alarm class (Default or Global Custom) and alarm severity level.

Icon	Color	Alarm severity	Meaning
	Yellow	Notice	The node is connected to the grid, but an unusual condition exists that does not affect normal operations.
	Light Orange	Minor	The node is connected to the grid, but an abnormal condition exists that could affect operation in the future. You should investigate to prevent escalation.
	Dark Orange	Major	The node is connected to the grid, but an abnormal condition exists that currently affects operation. This requires prompt attention to prevent escalation.
	Red	Critical	The node is connected to the grid, but an abnormal condition exists that has stopped normal operations. You should address the issue immediately.

The alarm severity and corresponding threshold value can be set for every numerical attribute. The NMS service on each Admin Node continuously monitors current attribute values against configured thresholds. When an alarm is triggered, a notification is sent to all designated personnel.

Note that a severity level of Normal does not trigger an alarm.

Attribute values are evaluated against the list of enabled alarms defined for that attribute. The list of alarms is checked in the following order to find the first alarm class with a defined and enabled alarm for the attribute:

1. Global Custom alarms with alarm severities from Critical down to Notice.
2. Default alarms with alarm severities from Critical down to Notice.

After an enabled alarm for an attribute is found in the higher alarm class, the NMS service only evaluates within that class. The NMS service will not evaluate against the other lower priority classes. That is, if there is an enabled Global Custom alarm for an attribute, the NMS service only evaluates the attribute value against Global Custom alarms. Default alarms aren't evaluated. Thus, an enabled Default alarm for an attribute can meet the criteria needed to trigger an alarm, but it will not be triggered because a Global Custom alarm (that does not meet the specified criteria) for the same attribute is enabled. No alarm is triggered and no notification is sent.

Alarm triggering example

You can use this example to understand how Global Custom alarms and Default alarms are triggered.

For the following example, an attribute has a Global Custom alarm and a Default alarm defined and enabled as shown in the following table.

	Global Custom alarm threshold (enabled)	**Default alarm threshold (enabled)**
Notice	>= 1500	>= 1000
Minor	>= 15,000	>= 1000
Major	>=150,000	>= 250,000

If the attribute is evaluated when its value is 1000, no alarm is triggered and no notification is sent.

The Global Custom alarm takes precedence over the Default alarm. A value of 1000 does not reach the threshold value of any severity level for the Global Custom alarm. As a result, the alarm level is evaluated to be Normal.

After the above scenario, if the Global Custom alarm is disabled, nothing changes. The attribute value must be reevaluated before a new alarm level is triggered.

With the Global Custom alarm disabled, when the attribute value is reevaluated, the attribute value is evaluated against the threshold values for the Default alarm. The alarm level triggers a Notice level alarm and an email notification is sent to the designated personnel.

Alarms of same severity

If two Global Custom alarms for the same attribute have the same severity, the alarms are evaluated with a "top down" priority.

For instance, if UMEM drops to 50MB, the first alarm is triggered (= 50000000), but not the one below it (<=100000000).

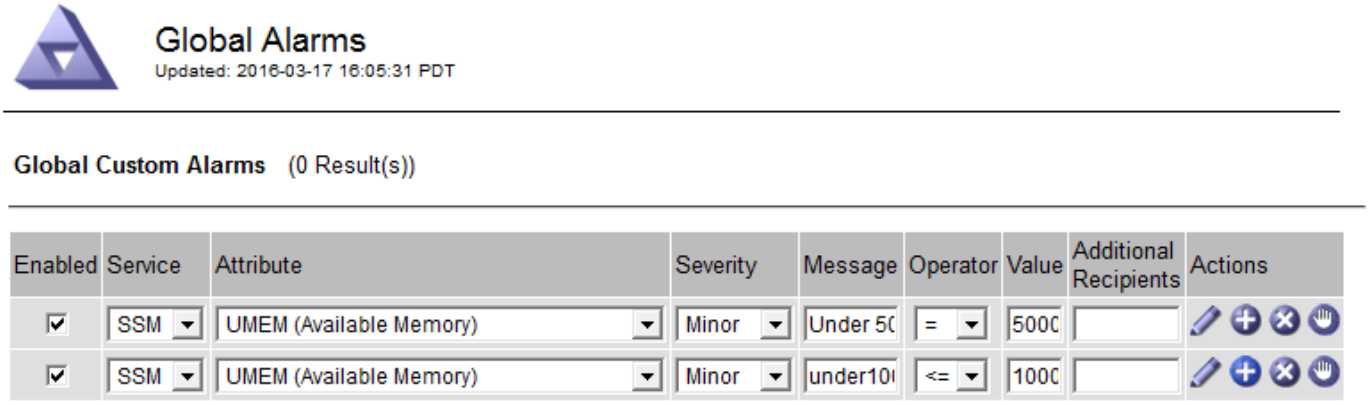

If the order is reversed, when UMEM drops to 100MB, the first alarm (<=100000000) is triggered, but not the one below it (= 50000000).

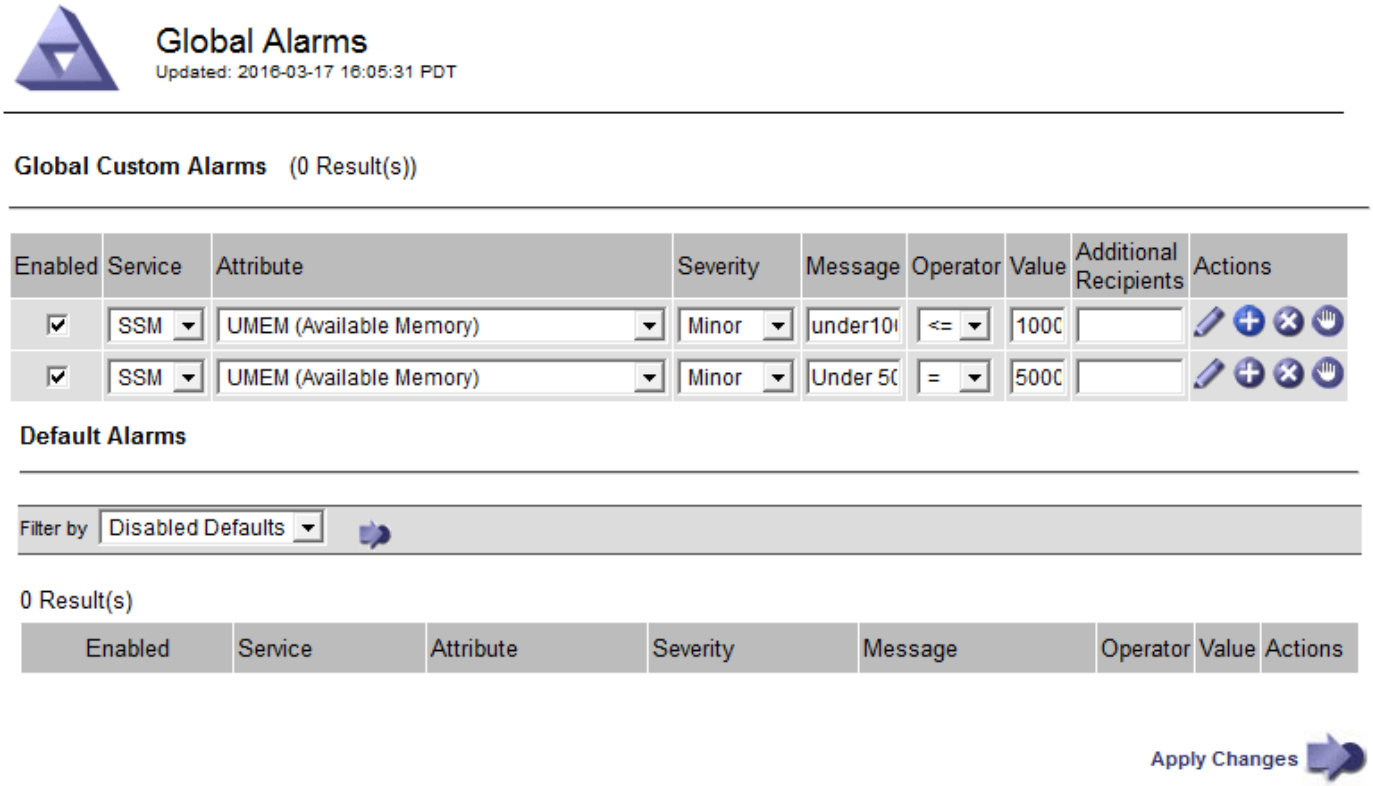

Notifications

A notification reports the occurrence of an alarm or the change of state for a service. Alarm notifications can be sent in email or using SNMP.

To avoid multiple alarms and notifications being sent when an alarm threshold value is reached, the alarm severity is checked against the current alarm severity for the attribute. If there is no change, then no further action is taken. This means that as the NMS service continues to monitor the system, it will only raise an alarm and send notifications the first time it notices an alarm condition for an attribute. If a new value threshold for the attribute is reached and detected, the alarm severity changes and a new notification is sent. Alarms are cleared when conditions return to the Normal level.

The trigger value shown in the notification of an alarm state is rounded to three decimal places. Therefore, an attribute value of 1.9999 triggers an alarm whose threshold is less than (<) 2.0, although the alarm notification shows the trigger value as 2.0.

New services

As new services are added through the addition of new grid nodes or sites, they inherit Default alarms and Global Custom alarms.

Alarms and tables

Alarm attributes displayed in tables can be disabled at the system level. Alarms can't be disabled for individual rows in a table.

For example, the following table shows two critical Entries Available (VMFI) alarms. (Select **SUPPORT** > **Tools** > **Grid topology**. Then, select ***Storage Node*** > **SSM** > **Resources**.)

You can disable the VMFI alarm so that the Critical level VMFI alarm is not triggered (both currently Critical alarms would appear in the table as green); however, you can't disable a single alarm in a table row so that one VMFI alarm displays as a Critical level alarm while the other remains green.

Volumes

Mount Point	Device	Status	Size	Space Available	Total Entries	Entries Available	Write Cache
/	sda1	Online	10.6 GB	7.46 GB	655,360	559,263	Enabled
/var/local	sda3	Online	63.4 GB	59.4 GB	3,932,160	3,931,842	Unknown
/var/local/rangedb/0	sdb	Online	53.4 GB	53.4 GB	52,428,800	52,427,856	Enabled
/var/local/rangedb/1	sdc	Online	53.4 GB	53.4 GB	52,428,800	52,427,848	Enabled
/var/local/rangedb/2	sdd	Online	53.4 GB	53.4 GB	52,428,800	52,427,856	Enabled

Acknowledge current alarms (legacy system)

Legacy alarms are triggered when system attributes reach alarm threshold values. Optionally, if you want to reduce or clear the list of legacy alarms, you can acknowledge the alarms.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You must have the Acknowledge alarms permission.

About this task

Because the legacy alarm system continues to be supported, the list of legacy alarms on the Current Alarms page is increased whenever a new alarm occurs. You can typically ignore the alarms (because alerts provide a better view of the system), or you can acknowledge the alarms.

> Optionally, when you have completely transitioned to the alert system, you can disable each legacy alarm to prevent it from being triggered and added to the count of legacy alarms.

When you acknowledge an alarm, it is no longer listed on the Current Alarms page in the Grid Manager, unless the alarm is triggered at the next severity level or it is resolved and occurs again.

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

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Current alarms**.

2. Select the service name in the table.

   The Alarms tab for the selected service appears (**SUPPORT** > **Tools** > **Grid topology** > ***Grid Node*** > ***Service*** > **Alarms**).

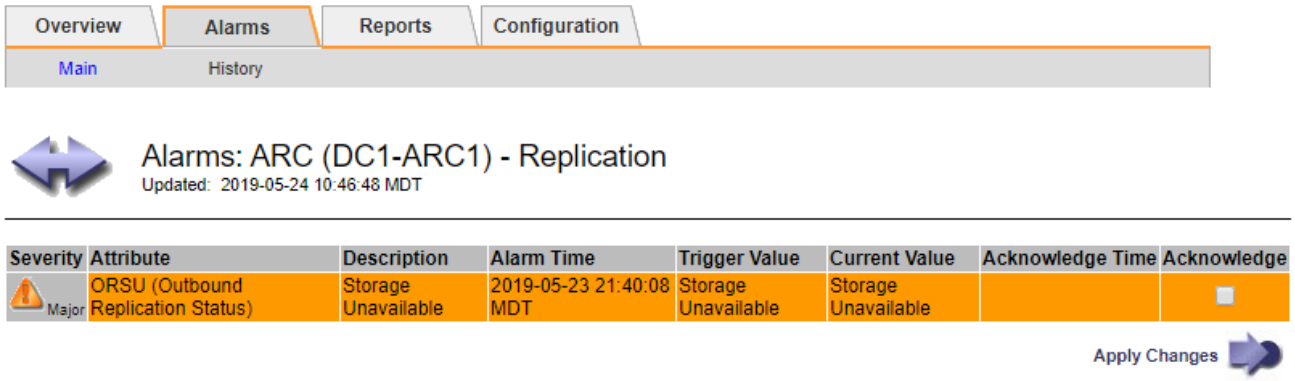

3. Select the **Acknowledge** checkbox for the alarm, and click **Apply Changes**.

   The alarm no longer appears on the dashboard or the Current Alarms page.

   > When you acknowledge an alarm, the acknowledgment is not copied to other Admin Nodes. For this reason, if you view the dashboard from another Admin Node, you might continue to see the active alarm.

4. As required, view acknowledged alarms.
   a. Select **SUPPORT** > **Alarms (legacy)** > **Current alarms**.
   b. Select **Show Acknowledged Alarms**.

      Any acknowledged alarms are shown.

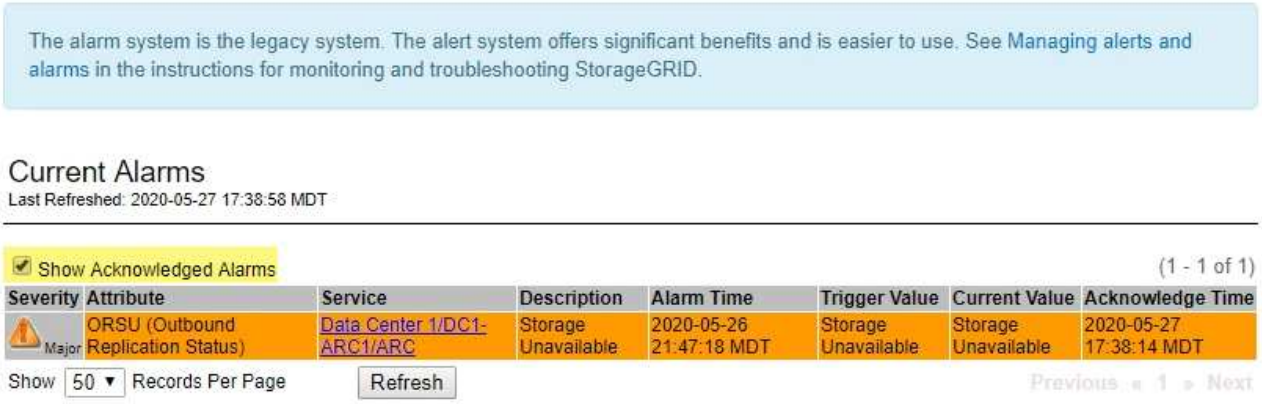

View Default alarms (legacy system)

You can view the list of all Default legacy alarms.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

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

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Global alarms**.
2. For Filter by, select **Attribute Code** or **Attribute Name**.
3. For equals, enter an asterisk: `*`
4. Click the arrow or press **Enter**.

   All Default alarms are listed.

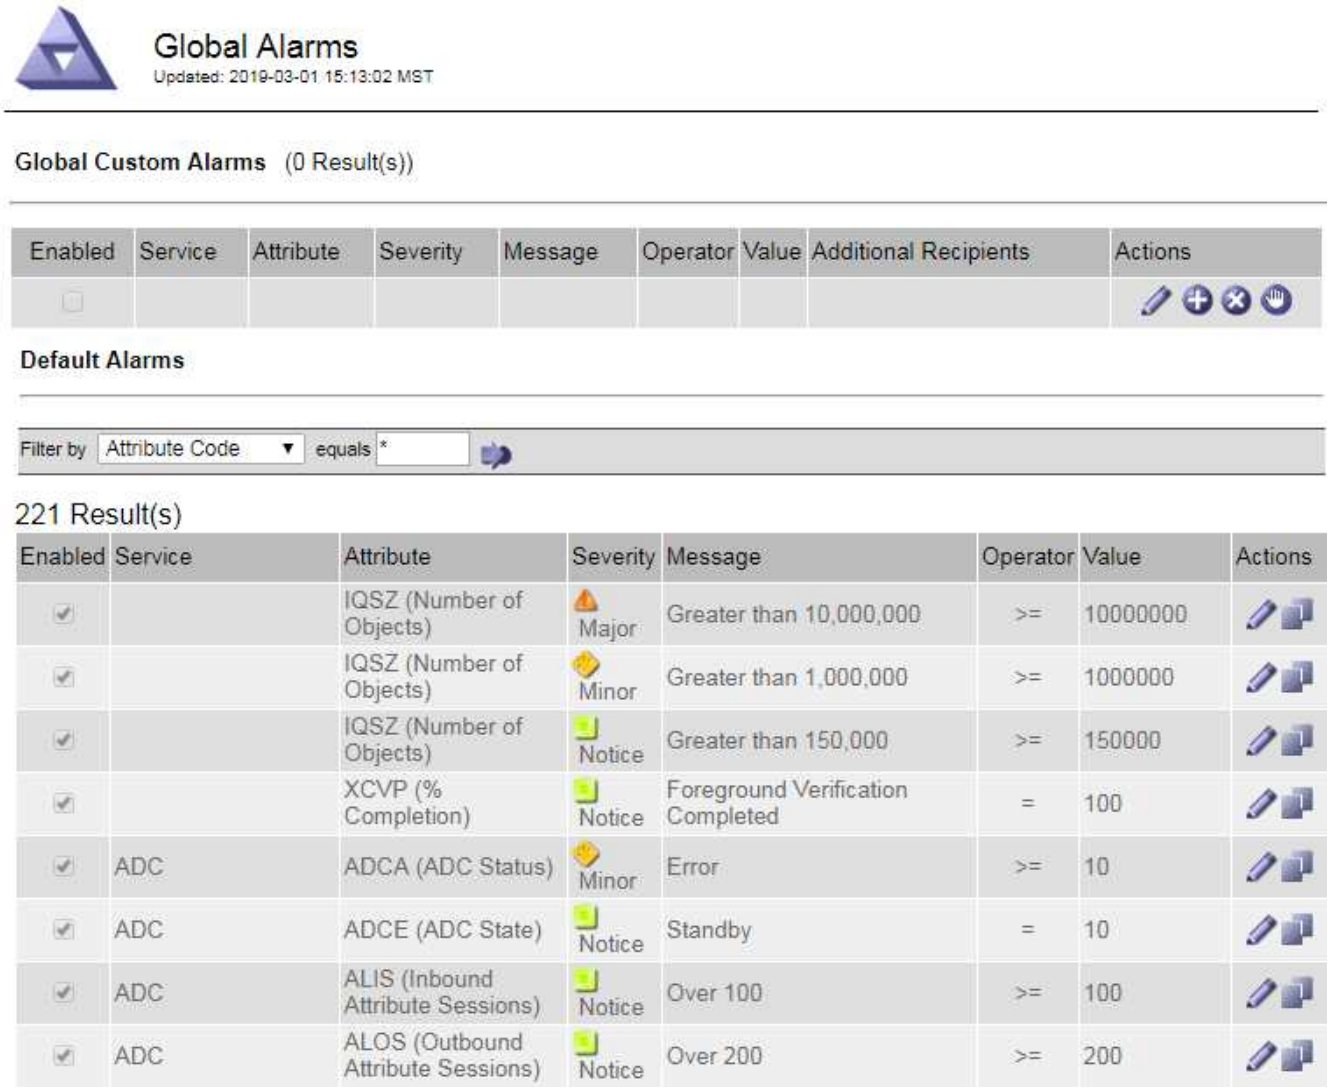

Review historical alarms and alarm frequency (legacy system)

When troubleshooting an issue, you can review how often a legacy alarm was triggered in the past.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

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

Steps

1. Follow these steps to get a list of all alarms triggered over a period of time.
   a. Select **SUPPORT** > **Alarms (legacy)** > **Historical alarms**.
   b. Do one of the following:
      - Click one of the time periods.
      - Enter a custom range, and click **Custom Query**.

2. Follow these steps to find out how often alarms have been triggered for a particular attribute.
   a. Select **SUPPORT** > **Tools** > **Grid topology**.
   b. Select ***grid node*** > ***service or component*** > **Alarms** > **History**.
   c. Select the attribute from the list.
   d. Do one of the following:
      - Click one of the time periods.
      - Enter a custom range, and click **Custom Query**.

        The alarms are listed in reverse chronological order.
   e. To return to the alarms history request form, click **History**.

Create Global Custom alarms (legacy system)

You might have used Global Custom alarms for the legacy system to address specific monitoring requirements. Global Custom alarms might have alarm levels that override Default alarms, or they might monitor attributes that don't have a Default alarm.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

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

Global Custom alarms override Default alarms. You should not change Default alarm values unless absolutely necessary. By changing Default alarms, you run the risk of concealing problems that might otherwise trigger an alarm.

> Be careful if you change alarm settings. For example, if you increase the threshold value for an alarm, you might not detect an underlying problem. Discuss your proposed changes with technical support before changing an alarm setting.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Global alarms**.
2. Add a new row to the Global Custom alarms table:
   - To add a new alarm, click **Edit** (if this is the first entry) or **Insert** .

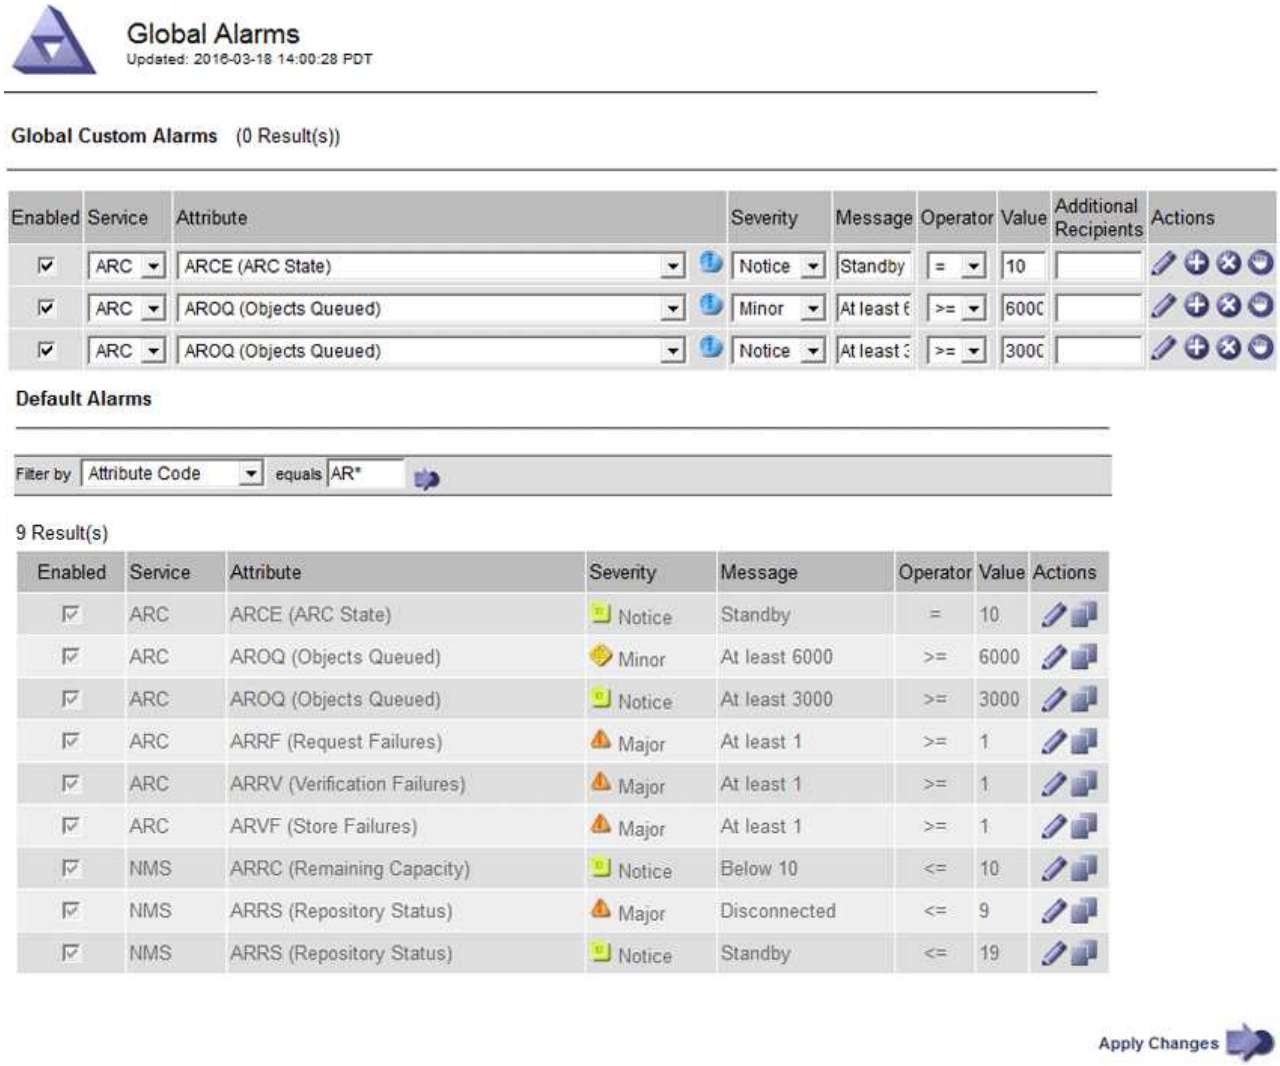

- To modify a Default alarm, search for the Default alarm.
  i. Under Filter by, select either **Attribute Code** or **Attribute Name**.
  ii. Type a search string.

     Specify four characters or use wildcards (for example, A??? or AB*). Asterisks (*) represent multiple characters, and question marks (?) represent a single character.

  iii. Click the arrow , or press **Enter**.
  iv. In the list of results, click **Copy** next to the alarm you want to modify.

     The Default alarm is copied to the Global Custom alarms table.

3. Make any necessary changes to the Global Custom alarms settings:

Heading	Description
Enabled	Select or clear the checkbox to enable or disable the alarm.

Heading	Description
Attribute	Select the name and code of the attribute being monitored from the list of all attributes applicable to the selected service or component. To display information about the attribute, click **Info** next to the attribute's name.
Severity	The icon and text indicating the level of the alarm.
Message	The reason for the alarm (connection lost, storage space below 10%, and so on).
Operator	Operators for testing the current attribute value against the Value threshold: • = equals • > greater than • < less than • >= greater than or equal to • <= less than or equal to • ≠ not equal to
Value	The alarm's threshold value used to test against the attribute's actual value using the operator. The entry can be a single number, a range of numbers specified with a colon (1:3), or a comma-delineated list of numbers and ranges.
Additional Recipients	A supplementary list of email addresses to be notified when the alarm is triggered. This is in addition to the mailing list configured on the **Alarms** > **Email Setup** page. Lists are comma delineated. **Note:** Mailing lists require SMTP server setup to operate. Before adding mailing lists, confirm that SMTP is configured. Notifications for Custom alarms can override notifications from Global Custom or Default alarms.
Actions	Control buttons to: Edit a row + Insert a row + Delete a row + Drag a row up or down + Copy a row

4. Click **Apply Changes**.

Disable alarms (legacy system)

The alarms in the legacy alarm system are enabled by default, but you can disable alarms that aren't required. You can also disable the legacy alarms after you have completely transitioned to the new alert system.

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

Disable a Default alarm (legacy system)

You can disable one of the legacy Default alarms for the entire system.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

Disabling an alarm for an attribute that currently has an alarm triggered does not clear the current alarm. The alarm will be disabled the next time the attribute crosses the alarm threshold, or you can clear the triggered alarm.

> Don't disable any of the legacy alarms until you have completely transitioned to the new alert system. Otherwise, you might not detect an underlying problem until it has prevented a critical operation from completing.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Global alarms**.
2. Search for the Default alarm to disable.
   a. In the Default Alarms section, select **Filter by** > **Attribute Code** or **Attribute Name**.
   b. Type a search string.

      Specify four characters or use wildcards (for example, A??? or AB*). Asterisks (*) represent multiple characters, and question marks (?) represent a single character.

   c. Click the arrow , or press **Enter**.

      > Selecting **Disabled Defaults** displays a list of all currently disabled Default alarms.

3. From the search results table, click the Edit icon for the alarm you want to disable.

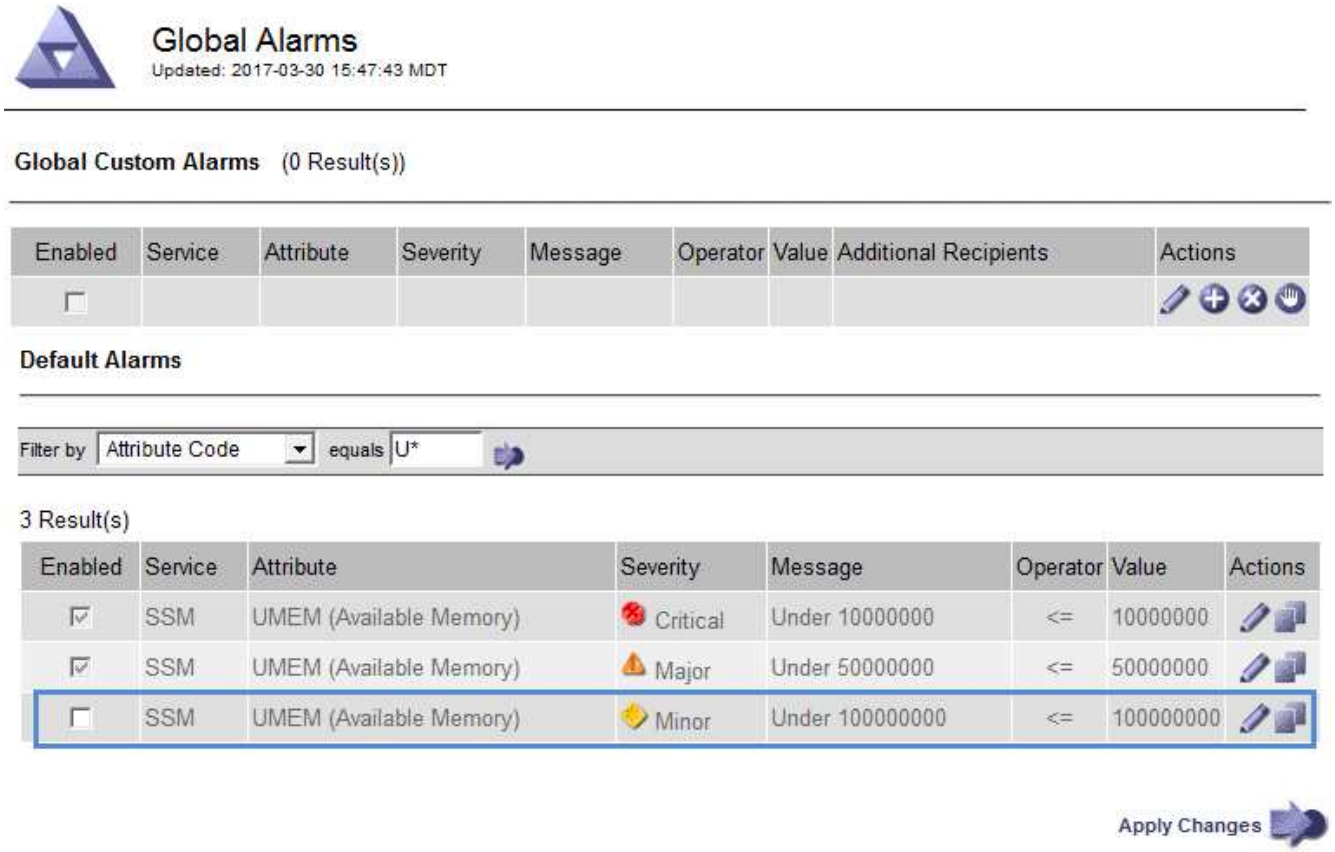

The **Enabled** checkbox for the selected alarm becomes active.

4. Clear the **Enabled** checkbox.
5. Click **Apply Changes**.

   The Default alarm is disabled.

Disable Global Custom alarms (legacy system)

You can disable a legacy Global Custom alarm for the entire system.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

Disabling an alarm for an attribute that currently has an alarm triggered does not clear the current alarm. The alarm will be disabled the next time the attribute crosses the alarm threshold, or you can clear the triggered alarm.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Global alarms**.
2. In the Global Custom Alarms table, click **Edit** next to the alarm you want to disable.
3. Clear the **Enabled** checkbox.

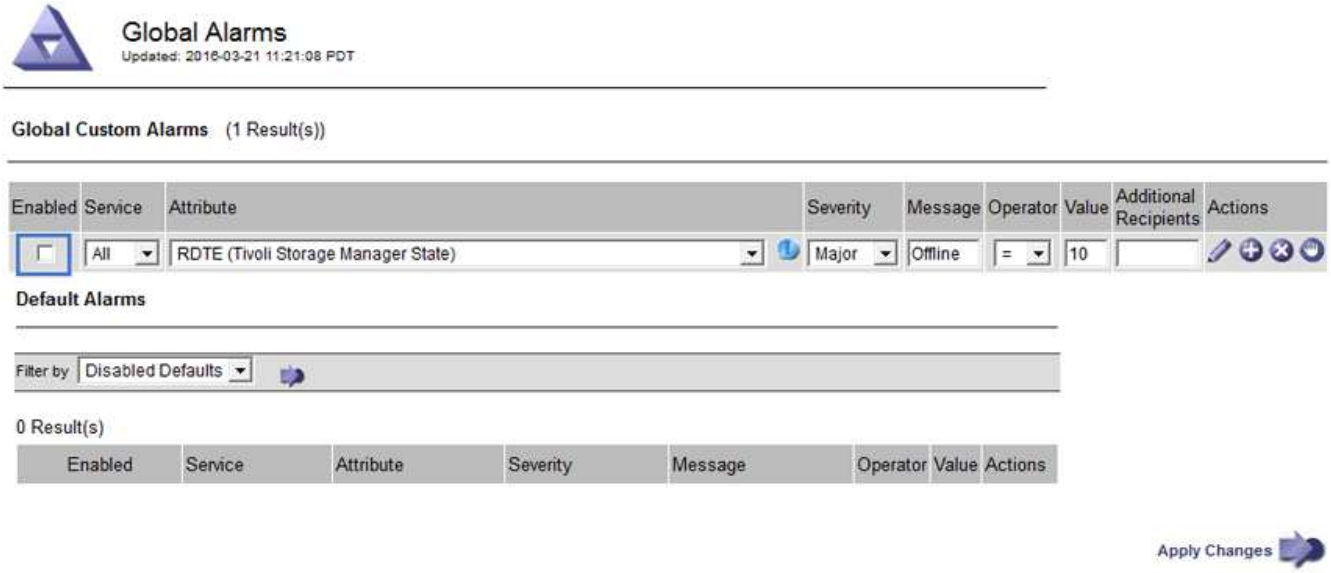

4. Click **Apply Changes**.

   The Global Custom alarm is disabled.

Clear triggered alarms (legacy system)

If a legacy alarm is triggered, you can clear it instead of acknowledging it.

Before you begin

- You must have the `Passwords.txt` file.

Disabling an alarm for an attribute that currently has an alarm triggered against it does not clear the alarm. The alarm will be disabled the next time the attribute changes. You can acknowledge the alarm or, if you want to immediately clear the alarm rather than wait for the attribute value to change (resulting in a change to the alarm state), you can clear the triggered alarm. You might find this helpful if you want to clear an alarm immediately against an attribute whose value does not change often (for example, state attributes).

1. Disable the alarm.
2. Log in to the primary Admin Node:
   a. Enter the following command: *`ssh admin@primary_Admin_Node_IP`*
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.
3. Restart the NMS service: `service nms restart`
4. Log out of the Admin Node: `exit`

   The alarm is cleared.

Configure notifications for alarms (legacy system)

StorageGRID system can automatically send email and SNMP notifications when an alarm is triggered or a service state changes.

By default, alarm email notifications aren't sent. For email notifications, you must configure the email server and specify the email recipients. For SNMP notifications, you must configure the SNMP agent.

Types of alarm notifications (legacy system)

When a legacy alarm is triggered, the StorageGRID system sends out two types of alarm notifications: severity level and service state.

Severity level notifications

An alarm email notification is sent when a legacy alarm is triggered at a selected severity level:

- Notice
- Minor
- Major
- Critical

A mailing list receives all notifications related to the alarm for the selected severity. A notification is also sent when the alarm leaves the alarm level — either by being resolved or by entering a different alarm severity level.

Service state notifications

A service state notification is sent when a service (for example, the LDR service or NMS service) enters the selected service state and when it leaves the selected service state. Service state notifications are send when a service enters or leaves ones of the following service states:

- Unknown
- Administratively Down

A mailing list receives all notifications related to changes in the selected state.

Configure email server settings for alarms (legacy system)

If you want StorageGRID to send email notifications when a legacy alarm is triggered, you must specify the SMTP mail server settings. The StorageGRID system only sends email; it can't receive email.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

Use these settings to define the SMTP server used for legacy alarm email notifications and AutoSupport email messages. These settings aren't used for alert notifications.

> If you use SMTP as the protocol for AutoSupport packages, you might have already configured an SMTP mail server. The same SMTP server is used for alarm email notifications, so you can skip this procedure. See the instructions for administering StorageGRID.

SMTP is the only protocol supported for sending email.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Legacy email setup**.
2. From the Email menu, select **Server**.

   The Email Server page appears. This page is also used to configure the email server for AutoSupport packages.

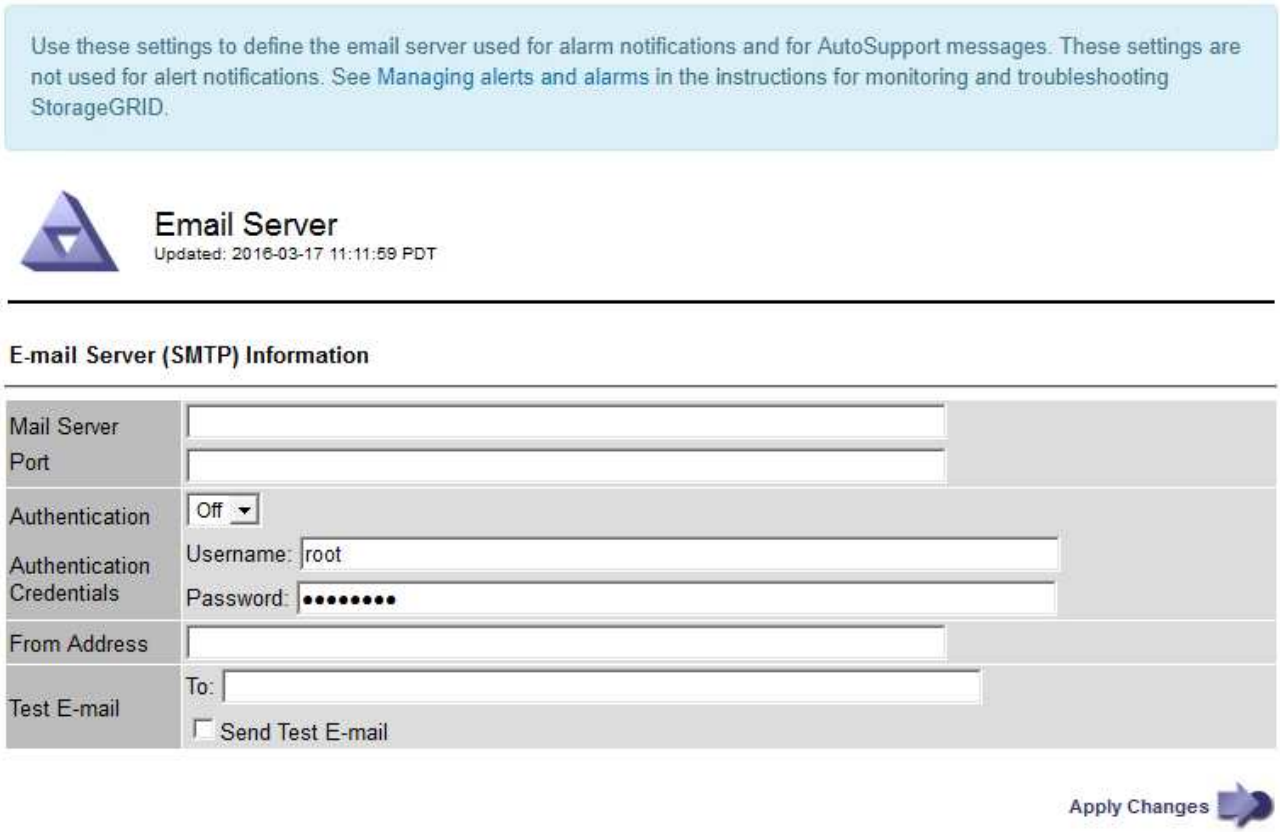

3. Add the following SMTP mail server settings:

Item	Description
Mail Server	IP address of the SMTP mail server. You can enter a hostname rather than an IP address if you have previously configured DNS settings on the Admin Node.
Port	Port number to access the SMTP mail server.
Authentication	Allows for the authentication of the SMTP mail server. By default, authentication is Off.
Authentication Credentials	Username and password of the SMTP mail server. If Authentication is set to On, a username and password to access the SMTP mail server must be provided.

4. Under **From Address**, enter a valid email address that the SMTP server will recognize as the sending email address. This is the official email address from which the email message is sent.
5. Optionally, send a test email to confirm that your SMTP mail server settings are correct.
   a. In the **Test E-mail** > **To** box, add one or more addresses that you can access.

      You can enter a single email address or a comma-delineated list of email addresses. Because the NMS service does not confirm success or failure when a test email is sent, you must be able to check the test recipient's inbox.

   b. Select **Send Test E-mail**.
6. Click **Apply Changes**.

   The SMTP mail server settings are saved. If you entered information for a test email, that email is sent. Test emails are sent to the mail server immediately and aren't sent through the notifications queue. In a system with multiple Admin Nodes, each Admin Node sends an email. Receipt of the test email confirms that your SMTP mail server settings are correct and that the NMS service is successfully connecting to the mail server. A connection problem between the NMS service and the mail server triggers the legacy MINS (NMS Notification Status) alarm at the Minor severity level.

Create alarm email templates (legacy system)

Email templates let you customize the header, footer, and subject line of a legacy alarm email notification. You can use email templates to send unique notifications that contain the same body text to different mailing lists.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

Use these settings to define the email templates used for legacy alarm notifications. These settings aren't used for alert notifications.

Different mailing lists might require different contact information. Templates don't include the body text of the email message.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Legacy email setup**.
2. From the Email menu, select **Templates**.
3. Click **Edit** (or **Insert** if this is not the first template).

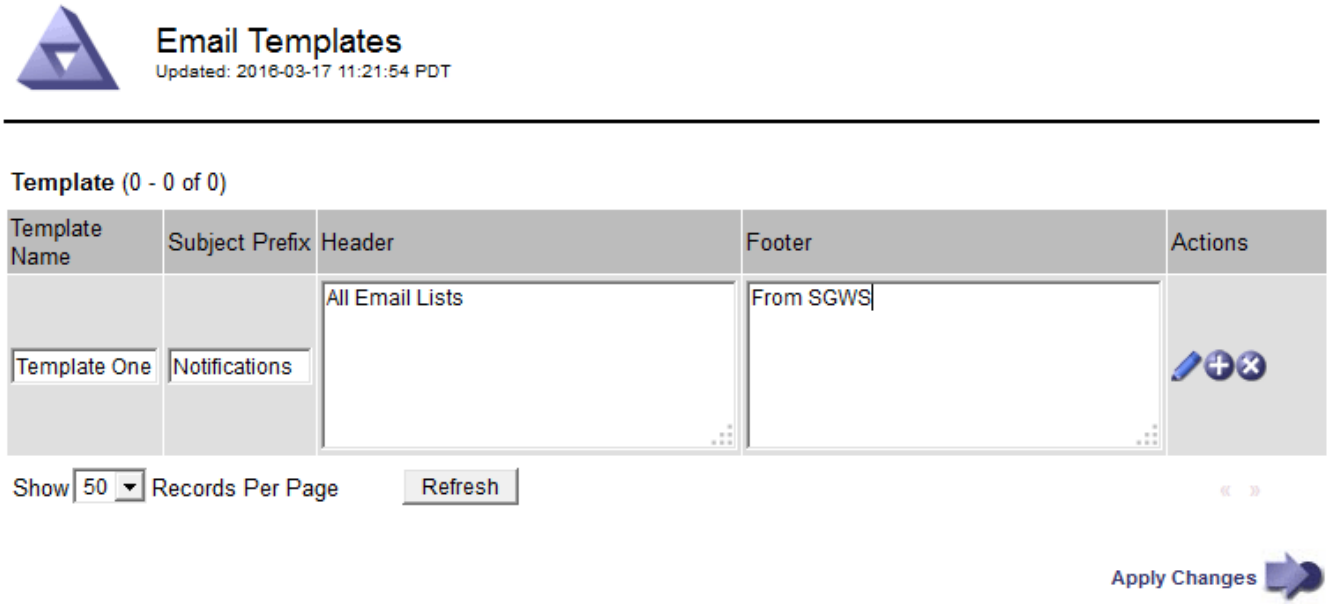

4. In the new row add the following:

Item	Description
Template Name	Unique name used to identify the template. Template names can't be duplicated.
Subject Prefix	Optional. Prefix that will appear at the beginning of an email's subject line. Prefixes can be used to easily configure email filters and organize notifications.
Header	Optional. Header text that appears at the beginning of the email message body. Header text can be used to preface the content of the email message with information such as company name and address.
Footer	Optional. Footer text that appears at the end of the email message body. Footer text can be used to close the email message with reminder information such as a contact phone number or a link to a web site.

5. Click **Apply Changes**.

   A new template for notifications is added.

Create mailing lists for alarm notifications (legacy system)

Mailing lists let you notify recipients when a legacy alarm is triggered or when a service state changes. You must create at least one mailing list before any alarm email notifications can be sent. To send a notification to a single recipient, create a mailing list with one email address.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

- If you want to specify an email template for the mailing list (custom header, footer, and subject line), you must have already created the template.

About this task

Use these settings to define the mailing lists used for legacy alarm email notifications. These settings aren't used for alert notifications.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Legacy email setup**.
2. From the Email menu, select **Lists**.
3. Click **Edit** (or *Insert* if this is not the first mailing list).

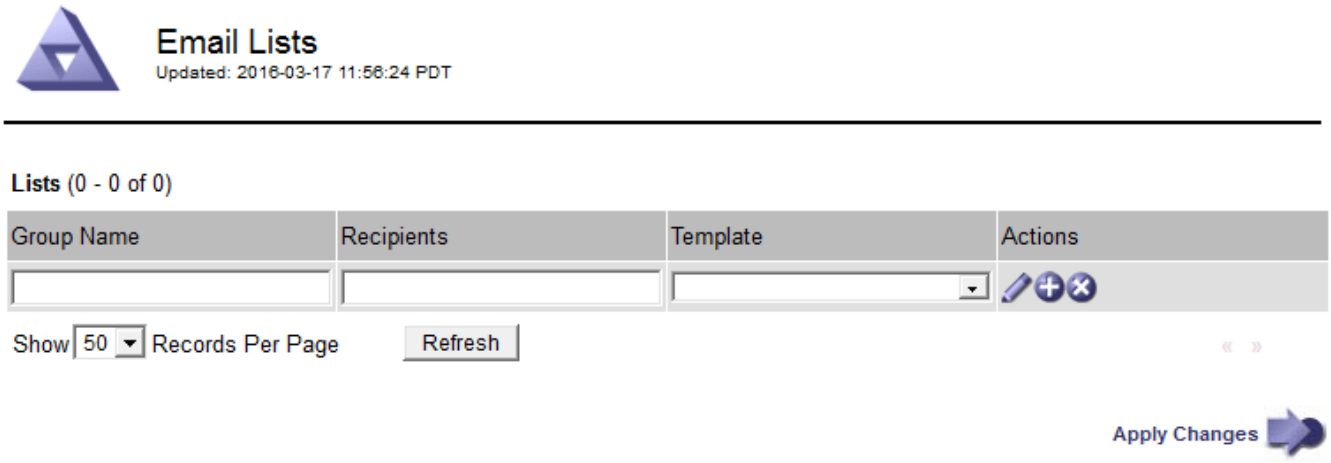

4. In the new row, add the following:

Item	Description
Group Name	Unique name used to identify the mailing list. Mailing list names can't be duplicated. **Note:** If you change the name of a mailing list, the change is not propagated to the other locations that use the mailing list name. You must manually update all configured notifications to use the new mailing list name.
Recipients	Single email address, a previously configured mailing list, or a comma-delineated list of email addresses and mailing lists to which notifications will be sent. **Note:** If an email address belongs to multiple mailing lists, only one email notification is sent when a notification triggering event occurs.
Template	Optionally, select an email template to add a unique header, footer, and subject line to notifications sent to all recipients of this mailing list.

5. Click **Apply Changes**.

   A new mailing list is created.

Configure email notifications for alarms (legacy system)

To receive email notifications for the legacy alarm system, recipients must be a member of a mailing list and that list must be added to the Notifications page. Notifications are configured to send email to recipients only when an alarm with a specified severity level is triggered or when a service state changes. Thus, recipients only receive the notifications they need to receive.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- You must have configured an email list.

About this task

Use these settings to configure notifications for legacy alarms. These settings aren't used for alert notifications.

If an email address (or list) belongs to multiple mailing lists, only one email notification is sent when a notification triggering event occurs. For example, one group of administrators within your organization can be configured to receive notifications for all alarms regardless of severity. Another group might only require notifications for alarms with a severity of critical. You can belong to both lists. If a critical alarm is triggered, you receive only one notification.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Legacy email setup**.
2. From the Email menu, select **Notifications**.
3. Click *Edit* (or *Insert* if this is not the first notification).
4. Under E-mail List, select the mailing list.
5. Select one or more alarm severity levels and service states.
6. Click **Apply Changes**.

   Notifications will be sent to the mailing list when alarms with the selected alarm severity level or service state are triggered or changed.

Suppress alarm notifications for a mailing list (legacy system)

You can suppress alarm notifications for a mailing list when you no longer want the mailing list to receive notifications about alarms. For example, you might want to suppress notifications about legacy alarms after you have transitioned to using alert email notifications.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

Use these settings to suppress email notifications for the legacy alarm system. These settings don't apply to alert email notifications.

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

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Legacy email setup**.
2. From the Email menu, select **Notifications**.
3. Click **Edit** next to the mailing list for which you want to suppress notifications.
4. Under Suppress, select the checkbox next to the mailing list you want to suppress, or select **Suppress** at the top of the column to suppress all mailing lists.
5. Click **Apply Changes**.

   Legacy alarm notifications are suppressed for the selected mailing lists.

View legacy alarms

Alarms (legacy system) are triggered when system attributes reach alarm threshold values. You can view the currently active alarms from the Current Alarms page.

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

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Current alarms**.

The alarm icon indicates the severity of each alarm, as follows:

Icon	Color	Alarm severity	Meaning
	Yellow	Notice	The node is connected to the grid, but an unusual condition exists that does not affect normal operations.
	Light Orange	Minor	The node is connected to the grid, but an abnormal condition exists that could affect operation in the future. You should investigate to prevent escalation.

Icon	Color	Alarm severity	Meaning
	Dark Orange	Major	The node is connected to the grid, but an abnormal condition exists that currently affects operation. This requires prompt attention to prevent escalation.
	Red	Critical	The node is connected to the grid, but an abnormal condition exists that has stopped normal operations. You should address the issue immediately.

2. To learn about the attribute that caused the alarm to be triggered, right click the attribute name in the table.
3. To view additional details about an alarm, click the service name in the table.

   The Alarms tab for the selected service appears (**SUPPORT** > **Tools** > **Grid topology** > ***Grid Node*** > ***Service*** > **Alarms**).

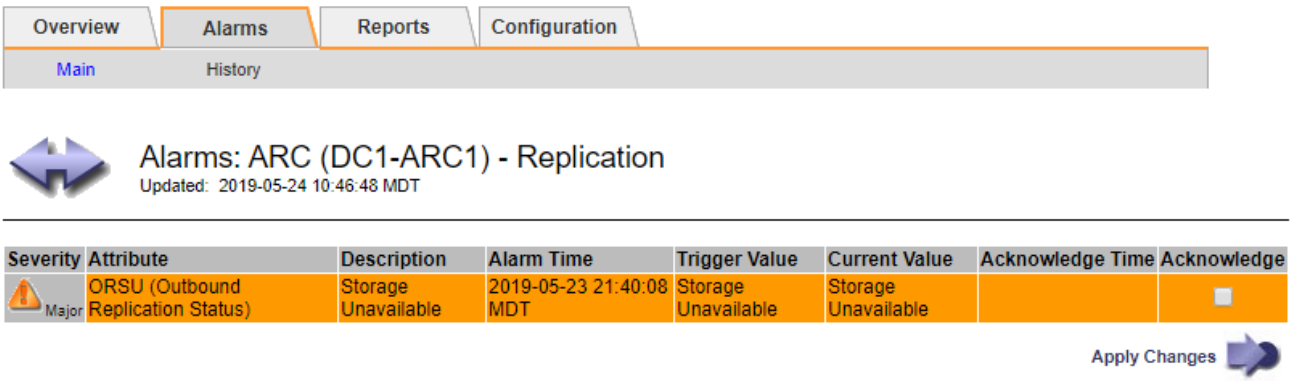

4. If you want to clear the count of current alarms, you can optionally do the following:
   - Acknowledge the alarm. An acknowledged alarm is no longer included in the count of legacy alarms unless it is triggered at the next severity level or it is resolved and occurs again.
   - Disable a particular Default alarm or Global Custom alarm for the entire system to prevent it from being triggered again.

Related information

Alarms reference (legacy system)

Acknowledge current alarms (legacy system)

Disable alarms (legacy system)

Alarms reference (legacy system)

The following table lists all of the legacy Default alarms. If an alarm is triggered, you can look up the alarm code in this table to find the recommended actions.

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

Code	Name	Service	Recommended action
ABRL	Available Attribute Relays	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BSSM, BDDS	Restore connectivity to a service (an ADC service) running an Attribute Relay Service as soon as possible. If there are no connected attribute relays, the grid node can't report attribute values to the NMS service. Thus, the NMS service can no longer monitor the status of the service, or update attributes for the service. If the problem persists, contact technical support.
ACMS	Available Metadata Services	BARC, BLDR, BCMN	An alarm is triggered when an LDR or ARC service loses connection to a DDS service. If this occurs, ingest or retrieve transactions can't be processed. If the unavailability of DDS services is only a brief transient issue, transactions can be delayed. Check and restore connections to a DDS service to clear this alarm and return the service to full functionality.
ACTS	Cloud Tiering Service Status	ARC	Only available for Archive Nodes with a Target Type of Cloud Tiering - Simple Storage Service (S3). If the ACTS attribute for the Archive Node is set to Read-Only Enabled or Read-Write Disabled, you must set the attribute to Read-Write Enabled. If a major alarm is triggered due to an authentication failure, verify the credentials associated with destination bucket and update values, if necessary. If a major alarm is triggered due to any other reason, contact technical support.
ADCA	ADC Status	ADC	If an alarm is triggered, select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **ADC** > **Overview** > **Main** and **ADC** > **Alarms** > **Main** to determine the cause of the alarm. If the problem persists, contact technical support.
ADCE	ADC State	ADC	If the value of ADC State is Standby, continue monitoring the service and if the problem persists, contact technical support. If the value of ADC State is Offline, restart the service. If the problem persists, contact technical support.

Code	Name	Service	Recommended action
AITE	Retrieve State	BARC	Only available for Archive Node's with a Target Type of Tivoli Storage Manager (TSM). If the value of Retrieve State is Waiting for Target, check the TSM middleware server and ensure that it is operating correctly. If the Archive Node has just been added to the StorageGRID system, ensure that the Archive Node's connection to the targeted external archival storage system is configured correctly. If the value of Archive Retrieve State is Offline, attempt to update the state to Online. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **ARC** > **Retrieve** > **Configuration** > **Main**, select **Archive Retrieve State** > **Online**, and click **Apply Changes**. If the problem persists, contact technical support.
AITU	Retrieve Status	BARC	If the value of Retrieve Status is Target Error, check the targeted external archival storage system for errors. If the value of Archive Retrieve Status is Session Lost, check the targeted external archival storage system to ensure it is online and operating correctly. Check the network connection with the target. If the value of Archive Retrieve Status is Unknown Error, contact technical support.
ALIS	Inbound Attribute Sessions	ADC	If the number of inbound attribute sessions on an attribute relay grows too large, it can be an indication that the StorageGRID system has become unbalanced. Under normal conditions, attribute sessions should be evenly distributed amongst ADC services. An imbalance can lead to performance issues. If the problem persists, contact technical support.
ALOS	Outbound Attribute Sessions	ADC	The ADC service has a high number of attribute sessions, and is becoming overloaded. If this alarm is triggered, contact technical support.

Code	Name	Service	Recommended action
ALUR	Unreachable Attribute Repositories	ADC	Check network connectivity with the NMS service to ensure that the service can contact the attribute repository. If this alarm is triggered and network connectivity is good, contact technical support.
AMQS	Audit Messages Queued	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BDDS	If audit messages can't be immediately forwarded to an audit relay or repository, the messages are stored in a disk queue. If the disk queue becomes full, outages can occur. To allow you to respond in time to prevent an outage, AMQS alarms are triggered when the number of messages in the disk queue reaches the following thresholds: • Notice: More than 100,000 messages • Minor: At least 500,000 messages • Major: At least 2,000,000 messages • Critical: At least 5,000,000 messages If an AMQS alarm is triggered, check the load on the system—if there have been a significant number of transactions, the alarm should resolve itself over time. In this case, you can ignore the alarm. If the alarm persists and increases in severity, view a chart of the queue size. If the number is steadily increasing over hours or days, the audit load has likely exceeded the audit capacity of the system. Reduce the client operation rate or decrease the number of audit messages logged by changing the audit level to Error or Off. See Configure audit messages and log destinations.
AOTE	Store State	BARC	Only available for Archive Node's with a Target Type of Tivoli Storage Manager (TSM). If the value of Store State is Waiting for Target, check the external archival storage system and ensure that it is operating correctly. If the Archive Node has just been added to the StorageGRID system, ensure that the Archive Node's connection to the targeted external archival storage system is configured correctly. If the value of Store State is Offline, check the value of Store Status. Correct any problems before moving the Store State back to Online.

Code	Name	Service	Recommended action
AOTU	Store Status	BARC	If the value of Store Status is Session Lost check that the external archival storage system is connected and online. If the value of Target Error, check the external archival storage system for errors. If the value of Store Status is Unknown Error, contact technical support.
APMS	Storage Multipath Connectivity	SSM	If the multipath state alarm appears as "Degraded" (select **SUPPORT** > **Tools** > **Grid topology**, then select ***site*** > ***grid node*** > **SSM** > **Events**), do the following: 1. Plug in or replace the cable that does not display any indicator lights. 2. Wait one to five minutes. Don’t unplug the other cable until at least five minutes after you plug in the first one. Unplugging too early can cause the root volume to become read-only, which requires that the hardware be restarted. 3. Return to the **SSM** > **Resources** page, and verify that the "Degraded" Multipath status has changed to "Nominal" in the Storage Hardware section.
ARCE	ARC State	ARC	The ARC service has a state of Standby until all ARC components (Replication, Store, Retrieve, Target) have started. It then transitions to Online. If the value of ARC State does not transition from Standby to Online, check the status of the ARC components. If the value of ARC State is Offline, restart the service. If the problem persists, contact technical support.
AROQ	Objects Queued	ARC	This alarm can be triggered if the removable storage device is running slowly due to problems with the targeted external archival storage system, or if it encounters multiple read errors. Check the external archival storage system for errors, and ensure that it is operating correctly. In some cases, this error can occur as a result of a high rate of data requests. Monitor the number of objects queued as system activity declines.

Code	Name	Service	Recommended action
ARRF	Request Failures	ARC	If a retrieval from the targeted external archival storage system fails, the Archive Node retries the retrieval as the failure can be due to a transient issue. However, if the object data is corrupt or has been marked as being permanently unavailable, the retrieval does not fail. Instead, the Archive Node continuously retries the retrieval and the value for Request Failures continues to increase. This alarm can indicate that the storage media holding the requested data is corrupt. Check the external archival storage system to further diagnose the problem. If you determine that the object data is no longer in the archive, the object will have to be removed from the StorageGRID system. For more information, contact technical support. Once the problem that triggered this alarm is addressed, reset the failures count. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **ARC** > **Retrieve** > **Configuration** > **Main**, select **Reset Request Failure Count** and click **Apply Changes**.
ARRV	Verification Failures	ARC	To diagnose and correct this problem, contact technical support. After the problem that triggered this alarm is addressed, reset the failures count. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **ARC** > **Retrieve** > **Configuration** > **Main**, select **Reset Verification Failure Count** and click **Apply Changes**.
ARVF	Store Failures	ARC	This alarm can occur as a result of errors with the targeted external archival storage system. Check the external archival storage system for errors, and ensure that it is operating correctly. Once the problem that triggered this alarm is addressed, reset the failures count. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **ARC** > **Retrieve** > **Configuration** > **Main**, select **Reset Store Failure Count**, and click **Apply Changes**.

Code	Name	Service	Recommended action
ASXP	Audit Shares	AMS	An alarm is triggered if the value of Audit Shares is Unknown. This alarm can indicate a problem with the installation or configuration of the Admin Node. If the problem persists, contact technical support.
AUMA	AMS Status	AMS	If the value of AMS Status is DB Connectivity Error, restart the grid node. If the problem persists, contact technical support.
AUME	AMS State	AMS	If the value of AMS State is Standby, continue monitoring the StorageGRID system. If the problem persists, contact technical support. If the value of AMS State is Offline, restart the service. If the problem persists, contact technical support.
AUXS	Audit Export Status	AMS	If an alarm is triggered, correct the underlying problem, and then restart the AMS service. If the problem persists, contact technical support.
BADD	Storage Controller Failed Drive Count	SSM	This alarm is triggered when one or more drives in a StorageGRID appliance has failed or is not optimal. Replace the drives as required.
BASF	Available Object Identifiers	CMN	When a StorageGRID system is provisioned, the CMN service is allocated a fixed number of object identifiers. This alarm is triggered when the StorageGRID system begins to exhaust its supply of object identifiers. To allocate more identifiers, contact technical support.

Code	Name	Service	Recommended action
BASS	Identifier Block Allocation Status	CMN	By default, an alarm is triggered when object identifiers can't be allocated because ADC quorum can't be reached. Identifier block allocation on the CMN service requires a quorum (50% + 1) of the ADC services to be online and connected. If quorum is unavailable, the CMN service is unable to allocate new identifier blocks until ADC quorum is reestablished. If ADC quorum is lost, there is generally no immediate impact on the StorageGRID system (clients can still ingest and retrieve content), as approximately one month's supply of identifiers are cached elsewhere in the grid; however, if the condition continues, the StorageGRID system will lose the ability to ingest new content. If an alarm is triggered, investigate the reason for the loss of ADC quorum (for example, it can be a network or Storage Node failure) and take corrective action. If the problem persists, contact technical support.
BRDT	Compute Controller Chassis Temperature	SSM	An alarm is triggered if the temperature of the compute controller in a StorageGRID appliance exceeds a nominal threshold. Check hardware components and environmental issues for overheated condition. If necessary, replace the component.
BTOF	Offset	BADC, BLDR, BNMS, BAMS, BCLB, BCMN, BARC	An alarm is triggered if the service time (seconds) differs significantly from the operating system time. Under normal conditions, the service should resynchronize itself. If the service time drifts too far from the operating system time, system operations can be affected. Confirm that the StorageGRID system's time source is correct. If the problem persists, contact technical support.
BTSE	Clock State	BADC, BLDR, BNMS, BAMS, BCLB, BCMN, BARC	An alarm is triggered if the service's time is not synchronized with the time tracked by the operating system. Under normal conditions, the service should resynchronize itself. If the time drifts too far from operating system time, system operations can be affected. Confirm that the StorageGRID system's time source is correct. If the problem persists, contact technical support.

Code	Name	Service	Recommended action
CAHP	Java Heap Usage Percent	DDS	An alarm is triggered if Java is unable to perform garbage collection at a rate that allows enough heap space for the system to properly function. An alarm might indicate a user workload that exceeds the resources available across the system for the DDS metadata store. Check the ILM Activity in the dashboard, or select **SUPPORT** > **Tools** > **Grid topology**, then select ***site*** > ***grid node*** > **DDS** > **Resources** > **Overview** > **Main**. If the problem persists, contact technical support.
CASA	Data Store Status	DDS	An alarm is raised if the Cassandra metadata store becomes unavailable. Check the status of Cassandra: 1. At the Storage Node, log in as admin and `su` to root using the password listed in the Passwords.txt file. 2. Enter: `service cassandra status` 3. If Cassandra is not running, restart it: `service cassandra restart` This alarm might also indicate that the metadata store (Cassandra database) for a Storage Node requires rebuilding. See information about troubleshooting the Services: Status - Cassandra (SVST) alarm in Troubleshoot metadata issues. If the problem persists, contact technical support.
CASE	Data Store State	DDS	This alarm is triggered during installation or expansion to indicate a new data store is joining the grid.
CCNA	Compute Hardware	SSM	This alarm is triggered if the status of the compute controller hardware in a StorageGRID appliance is Needs Attention.

Code	Name	Service	Recommended action
CDLP	Metadata Used Space (Percent)	DDS	This alarm is triggered when the Metadata Effective Space (CEMS) reaches 70% full (minor alarm), 90% full (major alarm), and 100% full (critical alarm). If this alarm reaches the 90% threshold, a warning appears on the dashboard in the Grid Manager. You must perform an expansion procedure to add new Storage Nodes as soon as possible. See Expand a grid. If this alarm reaches the 100% threshold, you must stop ingesting objects and add Storage Nodes immediately. Cassandra requires a certain amount of space to perform essential operations such as compaction and repair. These operations will be impacted if object metadata uses more than 100% of the allowed space. Undesirable results can occur. **Note**: Contact technical support if you are unable to add Storage Nodes. After new Storage Nodes are added, the system automatically rebalances object metadata across all Storage Nodes, and the alarm clears. Also see information about troubleshooting the Low metadata storage alert in Troubleshoot metadata issues. If the problem persists, contact technical support.
CMNA	CMN Status	CMN	If the value of CMN Status is Error, select **SUPPORT** > **Tools** > **Grid topology**, then select ***site*** > ***grid node*** > **CMN** > **Overview** > **Main** and **CMN** > **Alarms** > **Main** to determine the cause of the error and to troubleshoot the problem. An alarm is triggered and the value of CMN Status is No Online CMN during a hardware refresh of the primary Admin Node when the CMNs are switched (the value of the old CMN State is Standby and the new is Online). If the problem persists, contact technical support.
CPRC	Remaining Capacity	NMS	An alarm is triggered if the remaining capacity (number of available connections that can be opened to the NMS database) falls below the configured alarm severity. If an alarm is triggered, contact technical support.

Code	Name	Service	Recommended action
CPSA	Compute Controller Power Supply A	SSM	An alarm is triggered if there is an issue with power supply A in the compute controller for a StorageGRID appliance. If necessary, replace the component.
CPSB	Compute Controller Power Supply B	SSM	An alarm is triggered if there is an issue with power supply B in the compute controller for a StorageGRID appliance. If necessary, replace the component.
CPUT	Compute Controller CPU Temperature	SSM	An alarm is triggered if the temperature of the CPU in the compute controller in a StorageGRID appliance exceeds a nominal threshold. If the Storage Node is a StorageGRID appliance, the StorageGRID system indicates that the controller needs attention. Check hardware components and environment issues for overheated condition. If necessary, replace the component.
DNST	DNS Status	SSM	After installation completes, a DNST alarm is triggered in the SSM service. After the DNS is configured and the new server information reaches all grid nodes, the alarm is canceled.
ECCD	Corrupt Fragments Detected	LDR	An alarm is triggered when the background verification process detects a corrupt erasure-coded fragment. If a corrupt fragment is detected, an attempt is made to rebuild the fragment. Reset the Corrupt Fragments Detected and Copies Lost attributes to zero and monitor them to see if counts go up again. If counts do go up, there might be a problem with the Storage Node's underlying storage. A copy of erasure-coded object data is not considered missing until such time that the number of lost or corrupt fragments breaches the erasure code's fault tolerance; therefore, it is possible to have corrupt fragment and to still be able to retrieve the object. If the problem persists, contact technical support.

Code	Name	Service	Recommended action
ECST	Verification Status	LDR	This alarm indicates the current status of the background verification process for erasure-coded object data on this Storage Node. A major alarm is triggered if there is an error in the background verification process.
FOPN	Open File Descriptors	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BSSM, BDDS	FOPN can become large during peak activity. If it does not diminish during periods of slow activity, contact technical support.
HSTE	HTTP State	BLDR	See recommended actions for HSTU.
HSTU	HTTP Status	BLDR	HSTE and HSTU are related to HTTP for all LDR traffic, including S3, Swift, and other internal StorageGRID traffic. An alarm indicates that one of the following situations has occurred: • HTTP has been taken offline manually. • The Auto-Start HTTP attribute has been disabled. • The LDR service is shutting down. The Auto-Start HTTP attribute is enabled by default. If this setting is changed, HTTP could remain offline after a restart. If necessary, wait for the LDR service to restart. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***Storage Node*** > **LDR** > **Configuration**. If HTTP is offline, place it online. Verify that the Auto-Start HTTP attribute is enabled. If HTTP remains offline, contact technical support.
HTAS	Auto-Start HTTP	LDR	Specifies whether to start HTTP services automatically on start-up. This is a user-specified configuration option.
IRSU	Inbound Replication Status	BLDR, BARC	An alarm indicates that inbound replication has been disabled. Confirm configuration settings: Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **LDR** > **Replication** > **Configuration** > **Main**.

Code	Name	Service	Recommended action
LATA	Average Latency	NMS	Check for connectivity issues. Check system activity to confirm that there is an increase in system activity. An increase in system activity will result in an increase to attribute data activity. This increased activity will result in a delay to the processing of attribute data. This can be normal system activity and will subside. Check for multiple alarms. An increase in average latency times can be indicated by an excessive number of triggered alarms. If the problem persists, contact technical support.
LDRE	LDR State	LDR	If the value of LDR State is Standby, continue monitoring the situation and if the problem persists, contact technical support. If the value of LDR State is Offline, restart the service. If the problem persists, contact technical support.
LOST	Lost Objects	DDS, LDR	Triggered when the StorageGRID system fails to retrieve a copy of the requested object from anywhere in the system. Before a LOST (Lost Objects) alarm is triggered, the system attempts to retrieve and replace a missing object from elsewhere in the system. Lost objects represent a loss of data. The Lost Objects attribute is incremented whenever the number of locations for an object drops to zero without the DDS service purposely purging the content to satisfy the ILM policy. Investigate LOST (LOST Object) alarms immediately. If the problem persists, contact technical support. Troubleshoot lost and missing object data
MCEP	Management Interface Certificate Expiry	CMN	Triggered when the certificate used for accessing the management interface is about to expire. 1. From the Grid Manager, select **CONFIGURATION** > **Security** > **Certificates**. 2. On the **Global** tab, select **Management interface certificate**. 3. Upload a new management interface certificate.

Code	Name	Service	Recommended action
MINQ	E-mail Notifications Queued	NMS	Check the network connections of the servers hosting the NMS service and the external mail server. Also confirm that the email server configuration is correct. Configure email server settings for alarms (legacy system)
MINS	E-mail Notifications Status	BNMS	A minor alarm is triggered if the NMS service is unable to connect to the mail server. Check the network connections of the servers hosting the NMS service and the external mail server. Also confirm that the email server configuration is correct. Configure email server settings for alarms (legacy system)
MISS	NMS Interface Engine Status	BNMS	An alarm is triggered if the NMS interface engine on the Admin Node that gathers and generates interface content is disconnected from the system. Check Server Manager to determine if the server individual application is down.
NANG	Network Auto Negotiate Setting	SSM	Check the network adapter configuration. The setting must match preferences of your network routers and switches. An incorrect setting can have a severe impact on system performance.
NDUP	Network Duplex Setting	SSM	Check the network adapter configuration. The setting must match preferences of your network routers and switches. An incorrect setting can have a severe impact on system performance.
NLNK	Network Link Detect	SSM	Check the network cable connections on the port and at the switch. Check the network router, switch, and adapter configurations. Restart the server. If the problem persists, contact technical support.

Code	Name	Service	Recommended action
NRER	Receive Errors	SSM	The following can be causes of NRER alarms: • Forward error correction (FEC) mismatch • Switch port and NIC MTU mismatch • High link error rates • NIC ring buffer overrun See information about troubleshooting the Network Receive Error (NRER) alarm in Troubleshoot network, hardware, and platform issues.
NRLY	Available Audit Relays	BADC, BARC, BCLB, BCMN, BLDR, BNMS, BDDS	If audit relays aren't connected to ADC services, audit events can't be reported. They are queued and unavailable to users until the connection is restored. Restore connectivity to an ADC service as soon as possible. If the problem persists, contact technical support.
NSCA	NMS Status	NMS	If the value of NMS Status is DB Connectivity Error, restart the service. If the problem persists, contact technical support.
NSCE	NMS State	NMS	If the value of NMS State is Standby, continue monitoring and if the problem persists, contact technical support. If the value of NMS State is Offline, restart the service. If the problem persists, contact technical support.
NSPD	Speed	SSM	This can be caused by network connectivity or driver compatibility issues. If the problem persists, contact technical support.

Code	Name	Service	Recommended action
NTBR	Free Tablespace	NMS	If an alarm is triggered, check how fast database usage has been changing. A sudden drop (as opposed to a gradual change over time) indicates an error condition. If the problem persists, contact technical support. Adjusting the alarm threshold allows you to proactively manage when additional storage needs to be allocated. If the available space reaches a low threshold (see alarm threshold), contact technical support to change the database allocation.
NTER	Transmit Errors	SSM	These errors can clear without being manually reset. If they don't clear, check network hardware. Check that the adapter hardware and driver are correctly installed and configured to work with your network routers and switches. When the underlying problem is resolved, reset the counter. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **SSM** > **Resources** > **Configuration** > **Main**, select **Reset Transmit Error Count**, and click **Apply Changes**.
NTFQ	NTP Frequency Offset	SSM	If the frequency offset exceeds the configured threshold, there is likely a hardware problem with the local clock. If the problem persists, contact technical support to arrange a replacement.
NTLK	NTP Lock	SSM	If the NTP daemon is not locked to an external time source, check network connectivity to the designated external time sources, their availability, and their stability.
NTOF	NTP Time Offset	SSM	If the time offset exceeds the configured threshold, there is likely a hardware problem with the oscillator of the local clock. If the problem persists, contact technical support to arrange a replacement.
NTSJ	Chosen Time Source Jitter	SSM	This value indicates the reliability and stability of the time source that NTP on the local server is using as its reference. If an alarm is triggered, it can be an indication that the time source's oscillator is defective, or that there is a problem with the WAN link to the time source.

Code	Name	Service	Recommended action
NTSU	NTP Status	SSM	If the value of NTP Status is Not Running, contact technical support.
OPST	Overall Power Status	SSM	An alarm is triggered if the power of a StorageGRID appliance deviates from the recommended operating voltage. Check the status of Power Supply A or B to determine which power supply is operating abnormally. If necessary, replace the power supply.
OQRT	Objects Quarantined	LDR	After the objects are automatically restored by the StorageGRID system, the quarantined objects can be removed from the quarantine directory. 1. Select **SUPPORT** > **Tools** > **Grid topology**. 2. Select **site** > **Storage Node** > **LDR** > **Verification** > **Configuration** > **Main**. 3. Select **Delete Quarantined Objects**. 4. Click **Apply Changes**. The quarantined objects are removed, and the count is reset to zero.
ORSU	Outbound Replication Status	BLDR, BARC	An alarm indicates that outbound replication is not possible: storage is in a state where objects can't be retrieved. An alarm is triggered if outbound replication is disabled manually. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **LDR** > **Replication** > **Configuration**. An alarm is triggered if the LDR service is unavailable for replication. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **LDR** > **Storage**.
OSLF	Shelf Status	SSM	An alarm is triggered if the status of one of the components in the storage shelf for a storage appliance is degraded. Storage shelf components include the IOMs, fans, power supplies, and drive drawers.If this alarm is triggered, see the maintenance instructions for your appliance.

Code	Name	Service	Recommended action
PMEM	Service Memory Usage (Percent)	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BSSM, BDDS	Can have a value of Over Y% RAM, where Y represents the percentage of memory being used by the server. Figures under 80% are normal. Over 90% is considered a problem. If memory usage is high for a single service, monitor the situation and investigate. If the problem persists, contact technical support.
PSAS	Power Supply A Status	SSM	An alarm is triggered if power supply A in a StorageGRID appliance deviates from the recommended operating voltage. If necessary, replace power supply A.
PSBS	Power Supply B Status	SSM	An alarm is triggered if power supply B in a StorageGRID appliance deviates from the recommended operating voltage. If necessary, replace the power supply B.
RDTE	Tivoli Storage Manager State	BARC	Only available for Archive Nodes with a Target Type of Tivoli Storage Manager (TSM). If the value of Tivoli Storage Manager State is Offline, check Tivoli Storage Manager Status and resolve any problems. Bring the component back online. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **ARC** > **Target** > **Configuration** > **Main**, select **Tivoli Storage Manager State** > **Online**, and click **Apply Changes**.

Code	Name	Service	Recommended action
RDTU	Tivoli Storage Manager Status	BARC	Only available for Archive Nodes with a Target Type of Tivoli Storage Manager (TSM). If the value of Tivoli Storage Manager Status is Configuration Error and the Archive Node has just been added to the StorageGRID system, ensure that the TSM middleware server is correctly configured. If the value of Tivoli Storage Manager Status is Connection Failure, or Connection Failure, Retrying, check the network configuration on the TSM middleware server, and the network connection between the TSM middleware server and the StorageGRID system. If the value of Tivoli Storage Manager Status is Authentication Failure, or Authentication Failure, Reconnecting, the StorageGRID system can connect to the TSM middleware server, but can't authenticate the connection. Check that the TSM middleware server is configured with the correct user, password, and permissions, and restart the service. If the value of Tivoli Storage Manager Status is Session Failure, an established session has been lost unexpectedly. Check the network connection between the TSM middleware server and the StorageGRID system. Check the middleware server for errors. If the value of Tivoli Storage Manager Status is Unknown Error, contact technical support.
RIRF	Inbound Replications — Failed	BLDR, BARC	An Inbound Replications — Failed alarm can occur during periods of high load or temporary network disruptions. After system activity reduces, this alarm should clear. If the count of failed replications continues to increase, look for network problems and verify that the source and destination LDR and ARC services are online and available. To reset the count, select **SUPPORT** > **Tools** > **Grid topology**, then select ***site*** > ***grid node*** > **LDR** > **Replication** > **Configuration** > **Main**. Select **Reset Inbound Replication Failure Count**, and click **Apply Changes**.

Code	Name	Service	Recommended action
RIRQ	Inbound Replications — Queued	BLDR, BARC	Alarms can occur during periods of high load or temporary network disruption. After system activity reduces, this alarm should clear. If the count for queued replications continues to increase, look for network problems and verify that the source and destination LDR and ARC services are online and available.
RORQ	Outbound Replications — Queued	BLDR, BARC	The outbound replication queue contains object data being copied to satisfy ILM rules and objects requested by clients. An alarm can occur as a result of a system overload. Wait to see if the alarm clears when system activity declines. If the alarm recurs, add capacity by adding Storage Nodes.
SAVP	Total Usable Space (Percent)	LDR	If usable space reaches a low threshold, options include expanding the StorageGRID system or move object data to archive through an Archive Node.
SCAS	Status	CMN	If the value of Status for the active grid task is Error, look up the grid task message. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **CMN** > **Grid Tasks** > **Overview** > **Main**. The grid task message displays information about the error (for example, "check failed on node 12130011"). After you have investigated and corrected the problem, restart the grid task. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **CMN** > **Grid Tasks** > **Configuration** > **Main**, and select **Actions** > **Run**. If the value of Status for a grid task being stopped is Error, retry ending the grid task. If the problem persists, contact technical support.
SCEP	Storage API Service Endpoints Certificate Expiry	CMN	Triggered when the certificate used for accessing storage API endpoints is about to expire. 1. Select **CONFIGURATION** > **Security** > **Certificates**. 2. On the **Global** tab, select **S3 and Swift API certificate**. 3. Upload a new S3 and Swift API certificate.

Code	Name	Service	Recommended action
SCHR	Status	CMN	If the value of Status for the historical grid task is Aborted, investigate the reason and run the task again if required. If the problem persists, contact technical support.
SCSA	Storage Controller A	SSM	An alarm is triggered if there is an issue with storage controller A in a StorageGRID appliance. If necessary, replace the component.
SCSB	Storage Controller B	SSM	An alarm is triggered if there is an issue with storage controller B in a StorageGRID appliance. If necessary, replace the component. Some appliance models don’t have a storage controller B.
SHLH	Health	LDR	If the value of Health for an object store is Error, check and correct: • problems with the volume being mounted • file system errors
SLSA	CPU Load Average	SSM	The higher the value the busier the system. If the CPU Load Average persists at a high value, the number of transactions in the system should be investigated to determine whether this is due to heavy load at the time. View a chart of the CPU load average: Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **SSM** > **Resources** > **Reports** > **Charts**. If the load on the system is not heavy and the problem persists, contact technical support.
SMST	Log Monitor State	SSM	If the value of Log Monitor State is not Connected for a persistent period of time, contact technical support.

Code	Name	Service	Recommended action
SMTT	Total Events	SSM	If the value of Total Events is greater than zero, check if there are known events (such as network failures) that can be the cause. Unless these errors have been cleared (that is, the count has been reset to 0), Total Events alarms can be triggered. When an issue is resolved, reset the counter to clear the alarm. Select **NODES** > ***site*** > ***grid node*** > **Events** > **Reset event counts**. To reset event counts, you must have the Grid topology page configuration permission. If the value of Total Events is zero, or the number increases and the problem persists, contact technical support.
SNST	Status	CMN	An alarm indicates that there is a problem storing the grid task bundles. If the value of Status is Checkpoint Error or Quorum Not Reached, confirm that a majority of ADC services are connected to the StorageGRID system (50 percent plus one) and then wait a few minutes. If the problem persists, contact technical support.
SOSS	Storage Operating System Status	SSM	An alarm is triggered if SANtricity OS indicates that there is a "Needs attention" issue with a component in a StorageGRID appliance. Select **NODES**. Then select **appliance Storage Node** > **Hardware**. Scroll down to view the status of each component. In SANtricity OS, check other appliance components to isolate the issue.
SSMA	SSM Status	SSM	If the value of SSM Status is Error, select **SUPPORT** > **Tools** > **Grid topology**, then select ***site*** > ***grid node*** > **SSM** > **Overview** > **Main** and **SSM** > **Overview** > **Alarms** to determine the cause of the alarm. If the problem persists, contact technical support.
SSME	SSM State	SSM	If the value of SSM State is Standby, continue monitoring, and if the problem persists, contact technical support. If the value of SSM State is Offline, restart the service. If the problem persists, contact technical support.

Code	Name	Service	Recommended action
SSTS	Storage Status	BLDR	If the value of Storage Status is Insufficient Usable Space, there is no more available storage on the Storage Node and data ingests are redirected to other available Storage Node. Retrieval requests can continue to be delivered from this grid node. Additional storage should be added. It is not impacting end user functionality, but the alarm persists until additional storage is added. If the value of Storage Status is Volume(s) Unavailable, a part of the storage is unavailable. Storage and retrieval from these volumes is not possible. Check the volume's Health for more information: Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **LDR** > **Storage** > **Overview** > **Main**. The volume's Health is listed under Object Stores. If the value of Storage Status is Error, contact technical support. Troubleshoot the Storage Status (SSTS) alarm

Code	Name	Service	Recommended action
SVST	Status	SSM	This alarm clears when other alarms related to a non-running service are resolved. Track the source service alarms to restore operation. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **SSM** > **Services** > **Overview** > **Main**. When the status of a service is shown as Not Running, its state is Administratively Down. The service's status can be listed as Not Running for the following reasons: • The service has been manually stopped (`/etc/init.d/<service\> stop`). • There is an issue with the MySQL database and Server Manager shuts down the MI service. • A grid node has been added, but not started. • During installation, a grid node has not yet connected to the Admin Node. If a service is listed as Not Running, restart the service (`/etc/init.d/<service\> restart`). This alarm might also indicate that the metadata store (Cassandra database) for a Storage Node requires rebuilding. If the problem persists, contact technical support. Troubleshoot the Services: Status - Cassandra (SVST) alarm
TMEM	Installed Memory	SSM	Nodes running with less than 24 GiB of installed memory can lead to performance problems and system instability. The amount of memory installed on the system should be increased to at least 24 GiB.
TPOP	Pending Operations	ADC	A queue of messages can indicate that the ADC service is overloaded. Too few ADC services can be connected to the StorageGRID system. In a large deployment, the ADC service can require adding computational resources, or the system can require additional ADC services.
UMEM	Available Memory	SSM	If the available RAM gets low, determine whether this is a hardware or software issue. If it is not a hardware issue, or if available memory falls below 50 MB (the default alarm threshold), contact technical support.

Code	Name	Service	Recommended action
VMFI	Entries Available	SSM	This is an indication that additional storage is required. Contact technical support.
VMFR	Space Available	SSM	If the value of Space Available gets too low (see alarm thresholds), it needs to be investigated as to whether there are log files growing out of proportion, or objects taking up too much disk space (see alarm thresholds) that need to be reduced or deleted. If the problem persists, contact technical support.
VMST	Status	SSM	An alarm is triggered if the value of Status for the mounted volume is Unknown. A value of Unknown or Offline can indicate that the volume can't be mounted or accessed due to a problem with the underlying storage device.
VPRI	Verification Priority	BLDR, BARC	By default, the value of Verification Priority is Adaptive. If Verification Priority is set to High, an alarm is triggered because storage verification can slow normal operations of the service.
VSTU	Object Verification Status	BLDR	Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***site*** > ***grid node*** > **LDR** > **Storage** > **Overview** > **Main**. Check the operating system for any signs of block-device or file system errors. If the value of Object Verification Status is Unknown Error, it usually indicates a low-level file system or hardware problem (I/O error) that prevents the Storage Verification task from accessing stored content. Contact technical support.
XAMS	Unreachable Audit Repositories	BADC, BARC, BCLB, BCMN, BLDR, BNMS	Check network connectivity to the server hosting the Admin Node. If the problem persists, contact technical support.

Log files reference

Log files reference: Overview

StorageGRID provides logs that are used to capture events, diagnostic messages, and error conditions. You might be asked to collect log files and forward them to technical support to assist with troubleshooting.

The logs are categorized as follows:

- StorageGRID software logs
- Deployment and maintenance logs
- Logs for third-party software
- About the bycast.log

> The details provided for each log type are for reference only. The logs are intended for advanced troubleshooting by technical support. Advanced techniques that involve reconstructing the problem history using the audit logs and the application log files are beyond the scope of these instructions.

Access the logs

To access the logs, you can collect log files and system data from one or more nodes as a single log file archive. Or, if the primary Admin Node is unavailable or unable to reach a specific node, you can access individual log files for each grid node as follows:

1. Enter the following command: `ssh admin@grid_node_IP`
2. Enter the password listed in the `Passwords.txt` file.
3. Enter the following command to switch to root: `su -`
4. Enter the password listed in the `Passwords.txt` file.

Log file categories

The StorageGRID log file archive contains the logs described for each category and additional files that contain metrics and debug command output.

Archive location	Description
`audit`	Audit messages generated during normal system operation.
`base-os-logs`	Base operating system information, including StorageGRID image versions.
`bundles`	Global configuration information (bundles).
`cassandra`	Cassandra database information and Reaper repair logs.
`ec`	VCSs information about the current node and EC group information by profile ID.
`grid`	General grid logs including debug (`bycast.log`) and `servermanager` logs.
`grid.xml`	Grid configuration file shared across all nodes.
`hagroups`	High availability groups metrics and logs.
`install`	`Gdu-server` and install logs.

Archive location	Description
`lumberjack.log`	Debug messages related to log collection.
`Lambda-arbitrator`	Logs related to the S3 Select proxy request.
`Metrics`	Service logs for Grafana, Jaeger, node exporter, and Prometheus.
`miscd`	Miscd access and error logs.
`mysql`	The mariaDB database configuration and related logs.
`net`	Logs generated by networking-related scripts and the Dynip service.
`nginx`	Load balancer and grid federation configuration files and logs. Also includes Grid Manager and Tenant Manager traffic logs.
`nginx-gw`	Load balancer and grid federation configuration files and logs.
`ntp`	NTP configuration file and logs.
`os`	Node and grid state file, including services `pid`.
`other`	Log files under `/var/local/log` that aren't collected in other folders.
`perf`	Peformance information for CPU, networking, and disk I/O.
`prometheus-data`	Current Prometheus metrics, if the log collection includes Prometheus data.
`provisioning`	Logs related to grid provisioning process.
`raft`	Logs from Raft cluster used in platform services.
`ssh`	Logs related to SSH configuration and service.
`snmp`	SNMP agent configuration and alarm allow/deny lists used for sending SNMP notifications.
`sockets-data`	Sockets data for network debug.
`system-commands.txt`	Output of StorageGRID container commands. Contains system information, such as networking and disk usage.

StorageGRID software logs

You can use StorageGRID logs to troubleshoot issues.

If you want to send your logs to an external syslog server or change the destination of audit information such as the `bycast.log` and `nms.log`, see Configure audit messages and log destinations.

General StorageGRID logs

File name	Notes	Found on
`/var/local/log/bycast.log`	The primary StorageGRID troubleshooting file. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***Site*** > ***Node*** > **SSM** > **Events**.	All nodes
`/var/local/log/bycast-err.log`	Contains a subset of `bycast.log` (messages with severity ERROR and CRITICAL). CRITICAL messages are also displayed in the system. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***Site*** > ***Node*** > **SSM** > **Events**.	All nodes
`/var/local/core/`	Contains any core dump files created if the program terminates abnormally. Possible causes include assertion failures, violations, or thread timeouts. **Note**: The file `` `/var/local/core/kexec_cmd `` usually exists on appliance nodes and does not indicate an error.	All nodes

Cipher-related logs

File name	Notes	Found on
`/var/local/log/ssh-config-generation.log`	Contains logs related to generating SSH configurations and reloading SSH services.	All nodes
`/var/local/log/nginx/config-generation.log`	Contains logs related to generating nginx configurations and reloading nginx services.	All nodes
`/var/local/log/nginx-gw/config-generation.log`	Contains logs related to generating nginx-gw configurations (and reloading nginx-gw services).	Admin and Gateway Nodes
`/var/local/log/update-cipher-configurations.log`	Contains logs related to configuring TLS and SSH policies.	All nodes

Grid federation logs

File name	Notes	Found on
`/var/local/log/update_grid_federation_config.log`	Contains logs related to generating nginx and nginx-gw configurations for grid federation connections.	All nodes

NMS logs

File name	Notes	Found on
`/var/local/log/nms.log`	• Captures notifications from the Grid Manager and the Tenant Manager. • Captures events related to the operation of the NMS service, for example, alarm processing, email notifications, and configuration changes. • Contains XML bundle updates resulting from configuration changes made in the system. • Contains error messages related to the attribute downsampling done once a day. • Contains Java web server error messages, for example, page generation errors and HTTP Status 500 errors.	Admin Nodes
`/var/local/log/nms.errlog`	Contains error messages related to MySQL database upgrades. Contains the Standard Error (stderr) stream of the corresponding services. There is one log file per service. These files are generally empty unless there are problems with the service.	Admin Nodes
`/var/local/log/nms.requestlog`	Contains information about outgoing connections from the Management API to internal StorageGRID services.	Admin Nodes

Server Manager logs

File name	Notes	Found on
`/var/local/log/servermanager.log`	Log file for the Server Manager application running on the server.	All nodes

File name	Notes	Found on
`/var/local/log/GridstatBackend.errlog`	Log file for the Server Manager GUI backend application.	All nodes
`/var/local/log/gridstat.errlog`	Log file for the Server Manager GUI.	All nodes

StorageGRID services logs

File name	Notes	Found on
`/var/local/log/acct.errlog`		Storage Nodes running the ADC service
`/var/local/log/adc.errlog`	Contains the Standard Error (stderr) stream of the corresponding services. There is one log file per service. These files are generally empty unless there are problems with the service.	Storage Nodes running the ADC service
`/var/local/log/ams.errlog`		Admin Nodes
`/var/local/log/arc.errlog`		Archive Nodes
`/var/local/log/cassandra/system.log`	Information for the metadata store (Cassandra database) that can be used if problems occur when adding new Storage Nodes, or if the nodetool repair task stalls.	Storage Nodes
`/var/local/log/cassandra-reaper.log`	Information for the Cassandra Reaper service, which performs repairs of the data in the Cassandra database.	Storage Nodes
`/var/local/log/cassandra-reaper.errlog`	Error information for the Cassandra Reaper service.	Storage Nodes
`/var/local/log/chunk.errlog`		Storage Nodes
`/var/local/log/cmn.errlog`		Admin Nodes
`/var/local/log/cms.errlog`	This log file might be present on systems that have been upgraded from an older version of StorageGRID. It contains legacy information.	Storage Nodes

File name	Notes	Found on
`/var/local/log/cts.errlog`	This log file is only created if the Target Type is **Cloud Tiering - Simple Storage Service (S3).**	Archive Nodes
`/var/local/log/dds.errlog`		Storage Nodes
`/var/local/log/dmv.errlog`		Storage Nodes
`/var/local/log/dynip*`	Contains logs related to the dynip service, which monitors the grid for dynamic IP changes and updates local configuration.	All nodes
`/var/local/log/grafana.log`	The log associated with the Grafana service, which is used for metrics visualization in the Grid Manager.	Admin Nodes
`/var/local/log/hagroups.log`	The log associated with high availability groups.	Admin Nodes and Gateway Nodes
`/var/local/log/hagroups_events.log`	Tracks state changes, such as transition from BACKUP to MASTER or FAULT.	Admin Nodes and Gateway Nodes
`/var/local/log/idnt.errlog`		Storage Nodes running the ADC service
`/var/local/log/jaeger.log`	The log associated with the jaeger service, which is used for trace collection.	All nodes
`/var/local/log/kstn.errlog`		Storage Nodes running the ADC service
`/var/local/log/lambda*`	Contains logs for the S3 Select service.	Admin and Gateway Nodes Only certain Admin and Gateway Nodes contain this log. See the S3 Select requirements and limitations for Admin and Gateway Nodes.
`/var/local/log/ldr.errlog`		Storage Nodes

File name	Notes	Found on
`/var/local/log/miscd/*.log`	Contains logs for the MISCd service (Information Service Control Daemon), which provides an interface for querying and managing services on other nodes and for managing environmental configurations on the node such as querying the state of services running on other nodes.	All nodes
`/var/local/log/nginx/*.log`	Contains logs for the nginx service, which acts as an authentication and secure communication mechanism for various grid services (such as Prometheus and Dynip) to be able to talk to services on other nodes over HTTPS APIs.	All nodes
`/var/local/log/nginx-gw/*.log`	Contains general logs related to the nginx-gw service, including error logs, and logs for the restricted admin ports on Admin Nodes.	Admin Nodes and Gateway Nodes
`/var/local/log/nginx-gw/cgr-access.log.gz`	Contains access logs related to cross-grid replication traffic.	Admin Nodes, Gateway Nodes, or both, based on the grid federation configuration. Only found on the destination grid for cross-grid replication.
`/var/local/log/nginx-gw/endpoint-access.log.gz`	Contains access logs for the Load Balancer service, which provides load balancing of S3 and Swift traffic from clients to Storage Nodes.	Admin Nodes and Gateway Nodes
`/var/local/log/persistence*`	Contains logs for the Persistence service, which manages files on the root disk that need to persist across a reboot.	All nodes
`/var/local/log/prometheus.log`	For all nodes, contains the node exporter service log and the ade-exporter metrics service log. For Admin Nodes, also contains logs for the Prometheus and Alert Manager services.	All nodes

File name	Notes	Found on
`/var/local/log/raft.log`	Contains the output of the library used by the RSM service for the Raft protocol.	Storage Nodes with RSM service
`/var/local/log/rms.errlog`	Contains logs for the Replicated State Machine Service (RSM) service, which is used for S3 platform services.	Storage Nodes with RSM service
`/var/local/log/ssm.errlog`		All nodes
`/var/local/log/update-s3vs-domains.log`	Contains logs related to processing updates for the S3 virtual hosted domain names configuration.See the instructions for implementing S3 client applications.	Admin and Gateway Nodes
`/var/local/log/update-snmp-firewall.*`	Contain logs related to the firewall ports being managed for SNMP.	All nodes
`/var/local/log/update-sysl.log`	Contains logs related to changes made to the system syslog configuration.	All nodes
`/var/local/log/update-traffic-classes.log`	Contains logs related to changes to the traffic classifiers configuration.	Admin and Gateway Nodes
`/var/local/log/update-utcn.log`	Contains logs related to Untrusted Client Network mode on this node.	All nodes

Related information

About the bycast.log

Use S3 REST API

Deployment and maintenance logs

You can use the deployment and maintenance logs to troubleshoot issues.

File name	Notes	Found on
`/var/local/log/install.log`	Created during software installation. Contains a record of the installation events.	All nodes
`/var/local/log/expansion-progress.log`	Created during expansion operations. Contains a record of the expansion events.	Storage Nodes
`/var/local/log/pa-move.log`	Created while running the `pa-move.sh` script.	Primary Admin Node

File name	Notes	Found on
`/var/local/log/pa-move-new_pa.log`	Created while running the `pa-move.sh` script.	Primary Admin Node
`/var/local/log/pa-move-old_pa.log`	Created while running the `pa-move.sh` script.	Primary Admin Node
`/var/local/log/gdu-server.log`	Created by the GDU service. Contains events related to provisioning and maintenance procedures managed by the primary Admin Node.	Primary Admin Node
`/var/local/log/send_admin_hw.log`	Created during installation. Contains debugging information related to a node's communications with the primary Admin Node.	All nodes
`/var/local/log/upgrade.log`	Created during software upgrade. Contains a record of the software update events.	All nodes

Logs for third-party software

You can use the third-party software logs to troubleshoot issues.

Category	File name	Notes	Found on
Archiving	`/var/local/log/dsierror.log`	Error information for TSM Client APIs.	Archive Nodes

MySQL	`/var/local/log/mysql.err`
	
	`/var/local/log/mysql-slow.log`
		Log files generated by MySQL.
		
		`mysql.err` captures database errors and events such as startups and shutdowns.
		
		`mysql-slow.log` (the slow query log) captures the SQL statements that took more than 10 seconds to execute.
			Admin Nodes

Operating system	`/var/local/log/messages`	This directory contains log files for the operating system. The errors contained in these logs are also displayed in the Grid Manager. Select **SUPPORT** > **Tools** > **Grid topology**. Then select **Topology** > ***Site*** > ***Node*** > **SSM** > **Events**.	All nodes

Category	File name	Notes	Found on

NTP	`/var/local/log/ntp.log`
	
	`/var/lib/ntp/var/log/ntpstats/`
		`/var/local/log/ntp.log` contains the log file for NTP error messages.
		
		`/var/lib/ntp/var/log/ntpstats/` directory contains NTP timing statistics.
		
		`loopstats` records loop filter statistics information.
		
		`peerstats` records peer statistics information.
			All nodes

About the bycast.log

The file `/var/local/log/bycast.log` is the primary troubleshooting file for the StorageGRID software. There is a `bycast.log` file for every grid node. The file contains messages specific to that grid node.

The file `/var/local/log/bycast-err.log` is a subset of `bycast.log`. It contains messages of severity ERROR and CRITICAL.

Optionally, you can change the destination of audit logs and send audit information to an external syslog server. Local logs of audit records continue to be generated and stored when an external syslog server is configured. See Configure audit messages and log destinations.

File rotation for bycast.log

When the `bycast.log` file reaches 1 GB, the existing file is saved, and a new log file is started.

The saved file is renamed `bycast.log.1`, and the new file is named `bycast.log`. When the new `bycast.log` reaches 1 GB, `bycast.log.1` is renamed and compressed to become `bycast.log.2.gz`, and `bycast.log` is renamed `bycast.log.1`.

The rotation limit for `bycast.log` is 21 files. When the 22nd version of the `bycast.log` file is created, the oldest file is deleted.

The rotation limit for `bycast-err.log` is seven files.

> If a log file has been compressed, you must not uncompress it to the same location in which it was written. Uncompressing the file to the same location can interfere with the log rotation scripts.

Optionally, you can change the destination of audit logs and send audit information to an external syslog server. Local logs of audit records continue to be generated and stored when an external syslog server is configured. See Configure audit messages and log destinations.

Related information

Collect log files and system data

Messages in bycast.log

Messages in `bycast.log` are written by the ADE (Asynchronous Distributed Environment). ADE is the runtime environment used by each grid node's services.

Example ADE message:

```
May 15 14:07:11 um-sec-rg1-agn3 ADE: |12455685    0357819531
SVMR EVHR 2019-05-05T27T17:10:29.784677| ERROR 0906 SVMR: Health
check on volume 3 has failed with reason 'TOUT'
```

ADE messages contain the following information:

Message segment	Value in example
Node ID	`12455685`
ADE process ID	`0357819531`
Module name	`SVMR`
Message identifier	`EVHR`
UTC system time	`2019-05-05T27T17:10:29.784677 (YYYY-MM-DDTHH:MM:SS.uuuuuu)`
Severity level	`ERROR`
Internal tracking number	`0906`
Message	`SVMR: Health check on volume 3 has failed with reason 'TOUT'`

Message severities in bycast.log

The messages in `bycast.log` are assigned severity levels.

For example:

- **NOTICE** — An event that should be recorded has occurred. Most log messages are at this level.
- **WARNING** — An unexpected condition has occurred.
- **ERROR** — A major error has occurred that will impact operations.
- **CRITICAL** — An abnormal condition has occurred that has stopped normal operations. You should address the underlying condition immediately. Critical messages are also displayed in the Grid Manager. Select **SUPPORT** > **Tools** > **Grid topology**. Then select **Site** > **Node** > **SSM** > **Events**.

Error codes in `bycast.log`

Most of the error messages in `bycast.log` contain error codes.

The following table lists common non-numerical codes in `bycast.log`. The exact meaning of a non-numerical code depends on the context in which it is reported.

Error code	Meaning
SUCS	No error
GERR	Unknown
CANC	Canceled
ABRT	Aborted
TOUT	Timeout
INVL	Invalid
NFND	Not found
VERS	Version
CONF	Configuration
FAIL	Failed
ICPL	Incomplete
DONE	Done
SUNV	Service unavailable

The following table lists the numerical error codes in `bycast.log`.

Error number	Error code	Meaning
001	EPERM	Operation not permitted
002	ENOENT	No such file or directory
003	ESRCH	No such process
004	EINTR	Interrupted system call
005	EIO	I/O error
006	ENXIO	No such device or address

Error number	Error code	Meaning
007	E2BIG	Argument list too long
008	ENOEXEC	Exec format error
009	EBADF	Bad file number
010	ECHILD	No child processes
011	EAGAIN	Try again
012	ENOMEM	Out of memory
013	EACCES	Permission denied
014	EFAULT	Bad address
015	ENOTBLK	Block device required
016	EBUSY	Device or resource busy
017	EEXIST	File exists
018	EXDEV	Cross-device link
019	ENODEV	No such device
020	ENOTDIR	Not a directory
021	EISDIR	Is a directory
022	EINVAL	Invalid argument
023	ENFILE	File table overflow
024	EMFILE	Too many open files
025	ENOTTY	Not a typewriter
026	ETXTBSY	Text file busy
027	EFBIG	File too large
028	ENOSPC	No space left on device

Error number	Error code	Meaning
029	ESPIPE	Illegal seek
030	EROFS	Read-only file system
031	EMLINK	Too many links
032	EPIPE	Broken pipe
033	EDOM	Math argument out of domain of func
034	ERANGE	Math result not representable
035	EDEADLK	Resource deadlock would occur
036	ENAMETOOLONG	File name too long
037	ENOLCK	No record locks available
038	ENOSYS	Function not implemented
039	ENOTEMPTY	Directory not empty
040	ELOOP	Too many symbolic links encountered
041		
042	ENOMSG	No message of desired type
043	EIDRM	Identifier removed
044	ECHRNG	Channel number out of range
045	EL2NSYNC	Level 2 not synchronized
046	EL3HLT	Level 3 halted
047	EL3RST	Level 3 reset
048	ELNRNG	Link number out of range
049	EUNATCH	Protocol driver not attached
050	ENOCSI	No CSI structure available

Error number	Error code	Meaning
051	EL2HLT	Level 2 halted
052	EBADE	Invalid exchange
053	EBADR	Invalid request descriptor
054	EXFULL	Exchange full
055	ENOANO	No anode
056	EBADRQC	Invalid request code
057	EBADSLT	Invalid slot
058		
059	EBFONT	Bad font file format
060	ENOSTR	Device not a stream
061	ENODATA	No data available
062	ETIME	Timer expired
063	ENOSR	Out of streams resources
064	ENONET	Machine is not on the network
065	ENOPKG	Package not installed
066	EREMOTE	Object is remote
067	ENOLINK	Link has been severed
068	EADV	Advertise error
069	ESRMNT	Srmount error
070	ECOMM	Communication error on send
071	EPROTO	Protocol error
072	EMULTIHOP	Multihop attempted

Error number	Error code	Meaning
073	EDOTDOT	RFS specific error
074	EBADMSG	Not a data message
075	EOVERFLOW	Value too large for defined data type
076	ENOTUNIQ	Name not unique on network
077	EBADFD	File descriptor in bad state
078	EREMCHG	Remote address changed
079	ELIBACC	Can’t access a needed shared library
080	ELIBBAD	Accessing a corrupted shared library
081	ELIBSCN	
082	ELIBMAX	Attempting to link in too many shared libraries
083	ELIBEXEC	Can’t exec a shared library directly
084	EILSEQ	Illegal byte sequence
085	ERESTART	Interrupted system call should be restarted
086	ESTRPIPE	Streams pipe error
087	EUSERS	Too many users
088	ENOTSOCK	Socket operation on non-socket
089	EDESTADDRREQ	Destination address required
090	EMSGSIZE	Message too long
091	EPROTOTYPE	Protocol wrong type for socket
092	ENOPROTOOPT	Protocol not available
093	EPROTONOSUPPORT	Protocol not supported
094	ESOCKTNOSUPPORT	Socket type not supported

Error number	Error code	Meaning
095	EOPNOTSUPP	Operation not supported on transport endpoint
096	EPFNOSUPPORT	Protocol family not supported
097	EAFNOSUPPORT	Address family not supported by protocol
098	EADDRINUSE	Address already in use
099	EADDRNOTAVAIL	Can’t assign requested address
100	ENETDOWN	Network is down
101	ENETUNREACH	Network is unreachable
102	ENETRESET	Network dropped connection because of reset
103	ECONNABORTED	Software caused connection to terminate
104	ECONNRESET	Connection reset by peer
105	ENOBUFS	No buffer space available
106	EISCONN	Transport endpoint is already connected
107	ENOTCONN	Transport endpoint is not connected
108	ESHUTDOWN	Can’t send after transport endpoint shutdown
109	ETOOMANYREFS	Too many references: can’t splice
110	ETIMEDOUT	Connection timed out
111	ECONNREFUSED	Connection refused
112	EHOSTDOWN	Host is down
113	EHOSTUNREACH	No route to host
114	EALREADY	Operation already in progress
115	EINPROGRESS	Operation now in progress
116		

Error number	Error code	Meaning
117	EUCLEAN	Structure needs cleaning
118	ENOTNAM	Not a XENIX named type file
119	ENAVAIL	No XENIX semaphores available
120	EISNAM	Is a named type file
121	EREMOTEIO	Remote I/O error
122	EDQUOT	Quota exceeded
123	ENOMEDIUM	No medium found
124	EMEDIUMTYPE	Wrong medium type
125	ECANCELED	Operation Canceled
126	ENOKEY	Required key not available
127	EKEYEXPIRED	Key has expired
128	EKEYREVOKED	Key has been revoked
129	EKEYREJECTED	Key was rejected by service
130	EOWNERDEAD	For robust mutexes: Owner died
131	ENOTRECOVERABLE	For robust mutexes: State not recoverable

Configure audit message and log destinations

Considerations for using an external syslog server

An external syslog server is a server outside of StorageGRID you can use to collect system audit information in a single location. Using an external syslog server enables you to reduce network traffic on your Admin Nodes and manage the information more efficiently. For StorageGRID, the outbound syslog message packet format is compliant with RFC 3164.

The types of audit information you can send to the external syslog server include:

- Audit logs containing the audit messages generated during normal system operation
- Security-related events such as logins and escalations to root

- Application logs that might be requested if it is necessary to open a support case to troubleshoot an issue you have encountered

When to use an external syslog server

An external syslog server is especially useful if you have a large grid, use multiple types of S3 applications, or want to retain all audit data. Sending audit information to an external syslog server enables you to:

- Collect and manage audit information such as audit messages, application logs, and security events more efficiently.
- Reduce network traffic on your Admin Nodes because audit information is transferred directly from the various Storage Nodes to the external syslog server, without having to go through an Admin Node.

> When logs are sent to an external syslog server, single logs greater than 8,192 bytes are truncated at the end of the message to conform with common limitations in external syslog server implementations.

> To maximize the options for full data recovery in the event of a failure of the external syslog server, up to 20 GB of local logs of audit records (`localaudit.log`) are maintained on each node.

How to configure an external syslog server

To learn how to configure an external syslog server, see Configure audit messages and external syslog server.

If you plan to configure use the TLS or RELP/TLS protocol, you must have the following certificates:

- **Server CA certificates**: One or more trusted CA certificates for verifying the external syslog server in PEM encoding. If omitted, the default Grid CA certificate will be used.
- **Client certificate**: The client certificate for authentication to the external syslog server in PEM encoding.
- **Client private key**: Private key for the client certificate in PEM encoding.

> If you use a client certificate you must also use a client private key. If you provide an encrypted private key, you must also provide the passphrase. There is no significant security benefit from using an encrypted private key because the key and passphrase must be stored; using an unencrypted private key, if available, is recommended for simplicity.

How to estimate the size of the external syslog server

Normally, your grid is sized to achieve a required throughput, defined in terms of S3 operations per second or bytes per second. For example, you might have a requirement that your grid handle 1,000 S3 operations per second, or 2,000 MB per second, of object ingests and retrievals. You should size your external syslog server according to your grid's data requirements.

This section provides some heuristic formulas that help you estimate the rate and average size of log messages of various types that your external syslog server needs to be capable of handling, expressed in terms of the known or desired performance characteristics of the grid (S3 operations per second).

Use S3 operations per second in estimation formulas

If your grid was sized for a throughput expressed in bytes per second, you must convert this sizing into S3 operations per second to use the estimation formulas. To convert grid throughput, you must first determine your average object size, which you can do using the information in existing audit logs and metrics (if any), or by using your knowledge of the applications that will use StorageGRID. For example, if your grid was sized to achieve a throughput of 2,000 MB/second, and your average object size is 2 MB, then your grid was sized to be able to handle 1,000 S3 operations per second (2,000 MB / 2 MB).

> The formulas for external syslog server sizing in the following sections provide common-case estimates (rather than worst-case estimates). Depending on your configuration and workload, you might see a higher or lower rate of syslog messages or volume of syslog data than the formulas predict. The formulas are meant to be used as guidelines only.

Estimation formulas for audit logs

If you have no information about your S3 workload other than number of S3 operations per second your grid is expected to support, you can estimate the volume of audit logs your external syslog server will need to handle using the following formulas, under the assumption that you leave the Audit Levels set to the default values (all categories set to Normal, except Storage, which is set to Error):

```
Audit Log Rate = 2 x S3 Operations Rate
Audit Log Average Size = 800 bytes
```

For example, if your grid is sized for 1,000 S3 operations per second, your external syslog server should be sized to support 2,000 syslog messages per second and should be able to receive (and typically store) audit log data at a rate of 1.6 MB per second.

If you know more about your workload, more accurate estimations are possible. For audit logs, the most important additional variables are the percentage of S3 operations that are PUTs (vs. GETS), and the average size, in bytes, of the following S3 fields (4-character abbreviations used in the table are audit log field names):

Code	Field	Description
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
S3BK	S3 bucket	The S3 bucket name.
S3KY	S3 key	The S3 key name, not including the bucket name. Operations on buckets don’t include this field.

Let’s use P to represent the percentage of S3 operations that are PUTs, where $0 \leq P \leq 1$ (so for a 100% PUT

workload, P = 1, and for a 100% GET workload, P = 0).

Let's use K to represent the average size of the sum of the S3 account names, S3 bucket, and S3 key. Suppose the S3 account name is always my-s3-account (13 bytes), buckets have fixed-length names like /my/application/bucket-12345 (28 bytes), and objects have fixed-length keys like 5733a5d7-f069-41ef-8fbd-13247494c69c (36 bytes). Then the value of K is 90 (13+13+28+36).

If you can determine values for P and K, you can estimate the volume of audit logs your external syslog server will need to handle using the following formulas, under the assumption that you leave the Audit Levels set to the defaults (all categories set to Normal, except Storage, which is set to Error):

```
Audit Log Rate = ((2 x P) + (1 - P)) x S3 Operations Rate
Audit Log Average Size = (570 + K) bytes
```

For example, if your grid is sized for 1,000 S3 operations per second, your workload is 50% PUTs, and your S3 account names, bucket names, and object names average 90 bytes, your external syslog server should be sized to support 1,500 syslog messages per second and should be able to receive (and typically store) audit log data at a rate of approximately 1 MB per second.

Estimation formulas for non-default audit levels

The formulas provided for audit logs assume the use of default audit level settings (all categories set to Normal, except Storage, which is set to Error). Detailed formulas for estimating the rate and average size of audit messages for non-default audit level settings aren't available. However, the following table can be used to make a rough estimate of the rate; you can use the average size formula provided for audit logs, but be aware that it is likely to result in an over-estimate because the "extra" audit messages are, on average, smaller than the default audit messages.

Condition	Formula
Replication: Audit levels all set to Debug or Normal	Audit log rate = 8 x S3 Operations Rate
Erasure coding: audit levels all set to Debug or Normal	Use same formula as for default settings

Estimation formulas for security events

Security events aren't correlated with S3 operations and typically produce a negligible volume of logs and data. For these reasons, no estimation formulas are provided.

Estimation formulas for application logs

If you have no information about your S3 workload other than the number of S3 operations per second your grid is expected to support, you can estimate the volume of applications logs your external syslog server will need to handle using the following formulas:

```
Application Log Rate = 3.3 x S3 Operations Rate
Application Log Average Size = 350 bytes
```

So, for example, if your grid is sized for 1,000 S3 operations per second, your external syslog server should be sized to support 3,300 application logs per second and be able to receive (and store) application log data at a rate of about 1.2 MB per second.

If you know more about your workload, more accurate estimations are possible. For application logs, the most important additional variables are the data protection strategy (replication vs. erasure coding), the percentage of S3 operations that are PUTs (vs. GETs/other), and the average size, in bytes, of the following S3 fields (4-character abbreviations used in table are audit log field names):

Code	Field	Description
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
S3BK	S3 bucket	The S3 bucket name.
S3KY	S3 key	The S3 key name, not including the bucket name. Operations on buckets don’t include this field.

Example sizing estimations

This section explains example cases of how to use the estimation formulas for grids with the following methods of data protection:

- Replication
- Erasure coding

If you use replication for data protection

Let P represent the percentage of S3 operations that are PUTs, where 0 ≤ P ≤ 1 (so for a 100% PUT workload, P = 1, and for a 100% GET workload, P = 0).

Let K represent the average size of the sum of the S3 account names, S3 bucket, and S3 key. Suppose the S3 account name is always my-s3-account (13 bytes), buckets have fixed-length names like /my/application/bucket-12345 (28 bytes), and objects have fixed-length keys like 5733a5d7-f069-41ef-8fbd-13247494c69c (36 bytes). Then K has a value of 90 (13+13+28+36).

If you can determine values for P and K, you can estimate the volume of application logs your external syslog server will have to be able to handle using the following formulas.

```
Application Log Rate = ((1.1 x P) + (2.5 x (1 - P))) x S3 Operations Rate
Application Log Average Size = (P x (220 + K)) + ((1 - P) x (240 + (0.2 x
K))) Bytes
```

So, for example, if your grid is sized for 1,000 S3 operations per second, your workload is 50% PUTs, and your S3 account names, bucket names, and object names average 90 bytes, your external syslog server should be sized to support 1800 application logs per second, and will be receiving (and typically storing) application data at a rate of 0.5 MB per second.

If you use erasure coding for data protection

Let P represent the percentage of S3 operations that are PUTs, where 0 ≤ P ≤ 1 (so for a 100% PUT workload, P = 1, and for a 100% GET workload, P = 0).

Let K represent the average size of the sum of the S3 account names, S3 bucket, and S3 key. Suppose the S3 account name is always my-s3-account (13 bytes), buckets have fixed-length names like /my/application/bucket-12345 (28 bytes), and objects have fixed-length keys like 5733a5d7-f069-41ef-8fbd-13247494c69c (36 bytes). Then K has a value of 90 (13+13+28+36).

If you can determine values for P and K, you can estimate the volume of application logs your external syslog server will have to be able to handle using the following formulas.

```
Application Log Rate = ((3.2 x P) + (1.3 x (1 - P))) x S3 Operations Rate
Application Log Average Size = (P x (240 + (0.4 x K))) + ((1 - P) x (185 +
(0.9 x K))) Bytes
```

So, for example, if your grid is sized for 1,000 S3 operations per second, your workload is 50% PUTs, and your S3 account names, bucket names, and object names average 90 bytes, your external syslog server should be sized to support 2,250 application logs per second and should be able to receive (and typically store) application data at a rate of 0.6 MB per second.

Configure audit messages and external syslog server

You can configure a number of settings related to audit messages. You can adjust the number of audit messages recorded; define any HTTP request headers you want to include in client read and write audit messages; configure an external syslog server; and specify where audit logs, security event logs, and StorageGRID software logs are sent.

Audit messages and logs record system activities and security events, and are essential tools for monitoring and troubleshooting. All StorageGRID nodes generate audit messages and logs to track system activity and events.

Optionally, you can configure an external syslog server to save audit information remotely. Using an external server minimizes the performance impact of audit message logging without reducing the completeness of audit data. An external syslog server is especially useful if you have a large grid, use multiple types of S3 applications, or want to retain all audit data. See Considerations for external syslog server for details.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- If you plan to configure an external syslog server, you have reviewed the considerations for using an external syslog server and ensured that the server has enough capacity to receive and store the log files.
- If you plan to configure an external syslog server using TLS or RELP/TLS protocol, you have the required server CA and client certificates and the client private key.

Change audit message levels

You can set a different audit level for each of the following categories of messages in the audit log:

Audit category	Default setting	More information
System	Normal	System audit messages
Storage	Error	Object storage audit messages
Management	Normal	Management audit message
Client reads	Normal	Client read audit messages
Client writes	Normal	Client write audit messages
ILM	Normal	ILM audit messages
Cross-grid replication	Error	CGRR: Cross-Grid Replication Request

These defaults apply if you initially installed StorageGRID using version 10.3 or later. If you initially used an earlier version of StorageGRID, the default for all categories is set to Normal.

During upgrades, audit level configurations will not be effective immediately.

Steps

1. Select **CONFIGURATION** > **Monitoring** > **Audit and syslog server**.
2. For each category of audit message, select an audit level from the drop-down list:

Audit level	Description
Off	No audit messages from the category are logged.
Error	Only error messages are logged—audit messages for which the result code was not "successful" (SUCS).
Normal	Standard transactional messages are logged—the messages listed in these instructions for the category.
Debug	Deprecated. This level behaves the same as the Normal audit level.

The messages included for any particular level include those that would be logged at the higher levels. For example, the Normal level includes all of the Error messages.

> If you don't require a detailed record of client read operations for your S3 applications, optionally change the **Client Reads** setting to **Error** to decrease the number of audit messages recorded in the audit log.

3. Select **Save**.

   A green banner indicates your configuration has been saved.

Define HTTP request headers

You can optionally define any HTTP request headers you want to include in client read and write audit messages. These protocol headers apply to S3 and Swift requests only.

Steps

1. In the **Audit protocol headers** section, define the HTTP request headers you want to include in client read and write audit messages.

   Use an asterisk (*) as a wildcard to match zero or more characters. Use the escape sequence (*) to match a literal asterisk.

2. Select **Add another header** to create additional headers, if needed.

   When HTTP headers are found in a request, they are included in the audit message under the field HTRH.

   > Audit protocol request headers are logged only if the audit level for **Client Reads** or **Client Writes** is not **Off**.

3. Select **Save**

   A green banner indicates your configuration has been saved.

Use an external syslog server

You can optionally configure an external syslog server to save audit logs, application logs, and security event logs to a location outside of your grid.

> If you don't want to use an external syslog server, skip this step and go to Select audit information destinations.

> If the configuration options available in this procedure aren't flexible enough to meet your requirements, additional configuration options can be applied using the `audit-destinations` endpoints, which are in the private API section of the Grid Management API. For example, you can use the API if you want to use different syslog servers for different groups of nodes.

Enter syslog information

Access the Configure external syslog server wizard and provide the information StorageGRID needs to access the external syslog server.

Steps

1. From the Audit and syslog server page, select **Configure external syslog server**. Or, if you have

previously configured an external syslog server, select **Edit external syslog server**.

The Configure external syslog server wizard appears.

2. For the **Enter syslog info** step of the wizard, enter a valid fully qualified domain name or an IPv4 or IPv6 address for the external syslog server in the **Host** field.
3. Enter the destination port on the external syslog server (must be an integer between 1 and 65535). The default port is 514.
4. Select the protocol used to send audit information to the external syslog server.

   Using **TLS** or **RELP/TLS** is recommended. You must upload a server certificate to use either of these options. Using certificates helps secure the connections between your grid and the external syslog server. For more information, see Manage security certificates.

   All protocol options require support by, and configuration of, the external syslog server. You must choose an option that is compatible with the external syslog server.

   > Reliable Event Logging Protocol (RELP) extends the functionality of the syslog protocol to provide reliable delivery of event messages. Using RELP can help prevent the loss of audit information if your external syslog server has to restart.

5. Select **Continue**.
6. If you selected **TLS** or **RELP/TLS**, upload the server CA certificates, client certificate, and client private key.
   a. Select **Browse** for the certificate or key you want to use.
   b. Select the certificate or key file.
   c. Select **Open** to upload the file.

      A green check appears next to the certificate or key file name, notifying you that it has been uploaded successfully.

7. Select **Continue**.

Manage syslog content

You can select which information to send to the external syslog server.

Steps

1. For the **Manage syslog content** step of the wizard, select each type of audit information you want to send to the external syslog server.
   - **Send audit logs**: Sends StorageGRID events and system activities
   - **Send security events**: Sends security events such as when an unauthorized user attempts to sign in or a user signs in as root
   - **Send application logs**: Sends log files useful for troubleshooting including:
     - `bycast-err.log`
     - `bycast.log`
     - `jaeger.log`

- `nms.log` (Admin Nodes only)
- `prometheus.log`
- `raft.log`
- `hagroups.log`

For information about StorageGRID software logs, see StorageGRID software logs.

2. Use the drop-down menus to select the severity and facility (type of message) for each category of audit information you want to send.

   Setting severity and facility values can help you aggregate the logs in customizable ways for easier analysis.

   a. For **Severity**, select **Passthrough**, or select a severity value between 0 and 7.

      If you select a value, the selected value will be applied to all messages of this type. Information about different severities will be lost if you override severity with a fixed value.

Severity	Description
Passthrough	Each message sent to the external syslog to have the same severity value as when it was logged locally onto the node: • For audit logs, the severity is "info." • For security events, the severity values are generated by the Linux distribution on the nodes. • For application logs, the severities vary between "info" and "notice," depending on what the issue is. For example, adding an NTP server and configuring an HA group gives a value of "info," while intentionally stopping the SSM or RSM service gives a value of "notice."
0	Emergency: System is unusable
1	Alert: Action must be taken immediately
2	Critical: Critical conditions
3	Error: Error conditions
4	Warning: Warning conditions
5	Notice: Normal but significant condition
6	Informational: Informational messages
7	Debug: Debug-level messages

   b. For **Facilty**, select **Passthrough**, or select a facility value between 0 and 23.

If you select a value, it will be applied to all messages of this type. Information about different facilities will be lost if you override facility with a fixed value.

Facility	Description
Passthrough	Each message sent to the external syslog to have the same facility value as when it was logged locally onto the node: • For audit logs, the facility sent to the external syslog server is "local7." • For security events, the facility values are generated by the linux distribution on the nodes. • For application logs, the application logs sent to the external syslog server have the following facility values: ◦ `bycast.log`: user or daemon ◦ `bycast-err.log`: user, daemon, local3, or local4 ◦ `jaeger.log`: local2 ◦ `nms.log`: local3 ◦ `prometheus.log`: local4 ◦ `raft.log`: local5 ◦ `hagroups.log`: local6
0	kern (kernel messages)
1	user (user-level messages)
2	mail
3	daemon (system daemons)
4	auth (security/authorization messages)
5	syslog (messages generated internally by syslogd)
6	lpr (line printer subsystem)
7	news (network news subsystem)
8	UUCP
9	cron (clock daemon)
10	security (security/authorization messages)

Facility	Description
11	FTP
12	NTP
13	logaudit (log audit)
14	logalert (log alert)
15	clock (clock daemon)
16	local0
17	local1
18	local2
19	local3
20	local4
21	local5
22	local6
23	local7

3. Select **Continue**.

Send test messages

Before starting to use an external syslog server, you should request that all nodes in your grid send test messages to the external syslog server. You should use these test messages to help you validate your entire log collection infrastructure before you commit to sending data to the external syslog server.

> Don’t use the external syslog server configuration until you confirm that the external syslog server received a test message from each node in your grid and that the message was processed as expected.

Steps

1. If you don’t want to send test messages because you are certain your external syslog server is configured properly and can receive audit information from all the nodes in your grid, select **Skip and finish**.

   A green banner indicates that the configuration has been saved.

2. Otherwise, select **Send test messages** (recommended).

Test results continuously appear on the page until you stop the test. While the test is in progress, your audit messages continue to be sent to your previously configured destinations.

3. If you receive any errors, correct them and select **Send test messages** again.

   See Troubleshoot an external syslog server to help you resolve any errors.

4. Wait until you see a green banner indicating all nodes have passed testing.
5. Check your syslog server to determine if test messages are being received and processed as expected.

   > If you are using UDP, check your entire log collection infrastructure. The UDP protocol does not allow for as rigorous error detection as the other protocols.

6. Select **Stop and finish**.

   You are returned to the **Audit and syslog server** page. A green banner indicates that the syslog server configuration has been saved.

   > StorageGRID audit information is not sent to the external syslog server until you select a destination that includes the external syslog server.

Select audit information destinations

You can specify where audit logs, security event logs, and StorageGRID software logs are sent.

> Some destination are available only if you have configured an external syslog server.

Steps

1. On the Audit and syslog server page, select the destination for audit information.

   > **Local nodes only** and **External syslog server** typically provide better performance.

Option	Description
Local nodes only	Audit messages, security event logs, and application logs are not sent to Admin Nodes. Instead, they are saved only on the nodes that generated them ("the local node"). The audit information generated on every local node is stored in `/var/local/log/localaudit.log` **Note**: StorageGRID periodically removes local logs in a rotation to free up space. When the log file for a node reaches 1 GB, the existing file is saved, and a new log file is started. The rotation limit for the log is 21 files. When the 22nd version of the log file is created, the oldest log file is deleted. On average about 20 GB of log data is stored on each node.

Option	Description
Admin Nodes/local nodes	Audit messages are sent to the audit log (`/var/local/log/audit.log`) on Admin Nodes, and security event logs and application logs are stored on the nodes that generated them.
External syslog server	Audit information is sent to an external syslog server and saved on the local nodes. The type of information sent depends upon how you configured the external syslog server. This option is enabled only after you have configured an external syslog server.
Admin Node and external syslog server	Audit messages are sent to the audit log (`/var/local/log/audit.log`) on Admin Nodes, and audit information is sent to the external syslog server and saved on the local node. The type of information sent depends upon how you configured the external syslog server. This option is enabled only after you have configured an external syslog server.

2. Select **Save**.

   A warning message appears.

3. Select **OK** to confirm that you want to change the destination for audit information.

   A green banner indicates that the audit configuration has been saved.

   New logs are sent to the destinations you selected. Existing logs remain in their current location.

Use SNMP monitoring

Use SNMP monitoring: Overview

If you want to monitor StorageGRID using the Simple Network Management Protocol (SNMP), you must configure the SNMP agent that is included with StorageGRID.

- Configure the SNMP agent
- Update the SNMP agent

Capabilities

Each StorageGRID node runs an SNMP agent, or daemon, that provides a MIB. The StorageGRID MIB contains table and notification definitions for alerts and alarms. The MIB also contains system description information such as platform and model number for each node. Each StorageGRID node also supports a subset of MIB-II objects.

> See Access MIB files if you want to download the MIB files on your grid nodes.

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. It provides read-only MIB access for queries, and it can send two types of event-driven notifications to a management system:

Traps

Traps are notifications sent by the SNMP agent that don't 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

Informs are similar to traps, but they require acknowledgment by the management system. If the SNMP agent doesn't 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 the preferred sender Admin Node.

  Each alert is mapped to one of three trap types based on the severity level of the alert: activeMinorAlert, activeMajorAlert, and activeCriticalAlert. For a list of the alerts that can trigger these traps, see the Alerts reference.

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

  SNMP notifications aren't sent for every alarm or every alarm severity.

SNMP version support

The table provides a high-level summary of what is supported for each SNMP version.

	SNMPv1	SNMPv2c	SNMPv3
Queries (GET and GETNEXT)	Read-only MIB queries	Read-only MIB queries	Read-only MIB queries
Query authentication	Community string	Community string	User-based Security Model (USM) user
Notifications (TRAP and INFORM)	Traps only	Traps and informs	Traps and informs
Notification authentication	Default trap community or a custom community string for each trap destination	Default trap community or a custom community string for each trap destination	USM user for each trap destination

Limitations

- StorageGRID supports read-only MIB access. Read-write access is not supported.
- All nodes in the grid receive the same configuration.
- SNMPv3: StorageGRID does not support the Transport Support Mode (TSM).
- SNMPv3: The only authentication protocol supported is SHA (HMAC-SHA-96).
- SNMPv3: The only privacy protocol supported is AES.

Configure the SNMP agent

You can configure the StorageGRID SNMP agent to use a third-party SNMP management system for read-only MIB access and notifications.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

About this task

The StorageGRID SNMP agent supports SNMPv1, SNMPv2c, and SNMPv3. You can configure the agent for one or more versions.
For SNMPv3, only User Security Model (USM) authentication is supported.

All nodes in the grid use the same SNMP configuration.

Specify basic configuration

As a first step, enable the StorageGRID SMNP agent and provide basic information.

Steps

1. Select **CONFIGURATION** > **Monitoring** > **SNMP agent**.

   The SNMP agent page appears.

2. To enable the SNMP agent on all grid nodes, select the **Enable SNMP** checkbox.
3. Enter the following information in the Basic configuration section.

Field	Description
System contact	Optional. The primary contact for the StorageGRID system, which is returned in SNMP messages as sysContact. The System contact is typically an email address. This value applies to all nodes in the StorageGRID system. **System contact** can be a maximum of 255 characters.

Field	Description
System location	Optional. The location of the StorageGRID system, which is returned in SNMP messages as sysLocation. The System location can be any information that is useful for identifying where your StorageGRID system is located. For example, you might use the street address of a facility. This value applies to all nodes in the StorageGRID system. **System location** can be a maximum of 255 characters.
Enable SNMP agent notifications	• If selected, the StorageGRID SNMP agent sends trap and inform notifications. • If not selected, the SNMP agent supports read-only MIB access, but it doesn't send any SNMP notifications.
Enable authentication traps	If selected, the StorageGRID SNMP agent sends authentication traps if it receives improperly authenticated protocol messages.

Enter community strings

If you use SNMPv1 or SNMPv2c, complete the Community strings section.

When the management system queries the StorageGRID MIB, it sends a community string. If the community string matches one of the values specified here, the SNMP agent sends a response to the management system.

Steps

1. For **Read-only community**, optionally enter a community string to allow read-only MIB access on IPv4 and IPv6 agent addresses.

   > To ensure the security of your StorageGRID system, don't use "public" as the community string. If you leave this field blank, the SNMP agent uses the grid ID of your StorageGRID system as the community string.

   Each community string can be a maximum of 32 characters and can't contain whitespace characters.

2. Select **Add another community string** to add additional strings.

   Up to five strings are allowed.

Create trap destinations

Use the Trap destinations tab in the Other configurations section to define one or more destinations for StorageGRID trap or inform notifications. When you enable the SNMP agent and select **Save**, StorageGRID sends notifications to each defined destination when alerts are triggered. Standard notifications are also sent for the supported MIB-II entities (for example, ifDown and coldStart).

Steps

1. For the **Default trap community** field, optionally enter the default community string you want to use for SNMPv1 or SNMPv2 trap destinations.

As required, you can provide a different ("custom") community string when you define a specific trap destination.

Default trap community can be a maximum of 32 characters and can’t contain whitespace characters.

2. To add a trap destination, select **Create**.
3. Select which SNMP version will be used for this trap destination.
4. Complete the Create trap destination form for the version you selected.

SNMPv1

If you selected SNMPv1 as the version, complete these fields.

Field	Description
Type	Must be Trap for SNMPv1.
Host	An IPv4 or IPv6 address or a fully-qualified domain name (FQDN) to receive the trap.
Port	Use 162, which the standard port for SNMP traps unless you must use another value.
Protocol	Use UDP, which is the standard SNMP trap protocol unless you need to use TCP.
Community string	Use the default trap community, if one was specified, or enter a custom community string for this trap destination. The custom community string can be a maximum of 32 characters and can’t contain whitespace.

SNMPv2c

If you selected SNMPv2c as the version, complete these fields.

Field	Description
Type	Whether the destination will be used for traps or informs.
Host	An IPv4 or IPv6 address or FQDN to receive the trap.
Port	Use 162, which is the standard port for SNMP traps unless you must use another value.
Protocol	Use UDP, which is the standard SNMP trap protocol unless you need to use TCP.
Community string	Use the default trap community, if one was specified, or enter a custom community string for this trap destination. The custom community string can be a maximum of 32 characters and can’t contain whitespace.

SNMPv3

If you selected SNMPv3 as the version, complete these fields.

Field	Description
Type	Whether the destination will be used for traps or informs.
Host	An IPv4 or IPv6 address or FQDN to receive the trap.
Port	Use 162, which is the standard port for SNMP traps unless you must use another value.
Protocol	Use UDP, which is the standard SNMP trap protocol unless you need to use TCP.
USM user	The USM user that will be used for authentication. • If you selected **Trap**, only USM users without authoritative engine IDs are shown. • If you selected **Inform**, only USM users with authoritative engine IDs are shown. • If no users are shown: 1. Create and save the trap destination. 2. Go to Create USM users and create the user. 3. Return to the Trap destinations tab, select the saved destination from the table, and select **Edit**. 4. Select the user.

5. Select **Create**.

   The trap destination is created and added to the table.

Create agent addresses

Optionally, use the Agent addresses tab in the Other configurations section to specify one or more "listening addresses." These are the StorageGRID addresses on which the SNMP agent can receive queries.

If you don’t configure an agent address, the default listening address is UDP port 161 on all StorageGRID networks.

Steps

1. Select **Create**.
2. Enter the following information.

Field	Description
Internet protocol	Whether this address will use IPv4 or IPv6. By default, SNMP uses IPv4.

Field	Description
Transport protocol	Whether this address will use UDP or TCP. By default, SNMP uses UDP.
StorageGRID network	Which StorageGRID network the agent will listen on. • Grid, Admin, and Client Networks: The SNMP agent will listen for queries on all three networks. • Grid Network • Admin Network • Client Network **Note**: If you use the Client Network for insecure data and you create an agent address for the Client Network, be aware that SNMP traffic will also be insecure.
Port	Optionally, the port number that the SNMP agent should listen on. The default UDP port for an SNMP agent is 161, but you can enter any unused port number. **Note**: When you save the SNMP agent, StorageGRID automatically opens the agent address ports on the internal firewall. You must ensure that any external firewalls allow access to these ports.

3. Select **Create**.

   The agent address is created and added to the table.

Create USM users

If you are using SNMPv3, use the USM users tab in the Other configurations section to define the USM users who are authorized to query the MIB or to receive traps and informs.

> SNMPv3 *inform* destinations must have users with engine IDs. SNMPv3 *trap* destination can't have users with engine IDs.

These steps don't apply if you are only using SNMPv1 or SNMPv2c.

Steps

1. Select **Create**.
2. Enter the following information.

Field	Description
Username	A unique name for this USM user. Usernames can have a maximum of 32 characters and can't contain whitespace characters. The username can't be changed after the user is created.
Read-only MIB access	If selected, this user should have read-only access to the MIB.
Authoritative engine ID	If this user will be used in an inform destination, the authoritative engine ID for this user. Enter 10 to 64 hex characters (5 to 32 bytes) with no spaces. This value is required for USM users that will be selected in trap destinations for informs. This value is not allowed for USM users that will be selected in trap destinations for traps. **Note**: This field is not shown if you selected **Read-only MIB access** because USM users who have read-only MIB access can't have engine IDs.
Security level	The security level for the USM user: • **authPriv**: This user communicates with authentication and privacy (encryption). You must specify an authentication protocol and password and a privacy protocol and password. • **authNoPriv**: This user communicates with authentication and without privacy (no encryption). You must specify an authentication protocol and password.
Authentication protocol	Always set to SHA, which is the only protocol supported (HMAC-SHA-96).
Password	The password this user will use for authentication.
Privacy protocol	Shown only if you selected **authPriv** and always set to AES, which is the only privacy protocol supported.
Password	Shown only if you selected **authPriv**. The password this user will use for privacy.

3. Select **Create**.

   The USM user is created and added to the table.

4. When you have completed the SNMP agent configuration, select **Save**.

   The new SNMP agent configuration becomes active.

Update the SNMP agent

You can disable SNMP notifications, update community strings, or add or remove agent addresses, USM users, and trap destinations.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

About this task

See Configure the SNMP agent for details about each field on the SNMP agent page. You must select **Save** at the bottom of the page to commit any changes you make on each tab.

Steps

1. Select **CONFIGURATION** > **Monitoring** > **SNMP agent**.

   The SNMP agent page appears.

2. To disable the SNMP agent on all grid nodes, clear the **Enable SNMP** checkbox, and select **Save**.

   If you re-enable the SNMP agent, any previous SNMP configuration settings are retained.

3. Optionally, update the information in the Basic configuration section:
   a. As required, update the **System contact** and **System location**.
   b. Optionally, select or clear the **Enable SNMP agent notifications** checkbox to control whether the StorageGRID SNMP agent sends trap and inform notifications.

      When this checkbox is cleared, the SNMP agent supports read-only MIB access, but it doesn't send SNMP notifications.

   c. Optionally, select or clear the **Enable authentication traps** checkbox to control whether the StorageGRID SNMP agent sends authentication traps when it receives improperly authenticated protocol messages.
4. If you use SNMPv1 or SNMPv2c, optionally update or add a **Read-only community** in the Community strings section.
5. To update trap destinations, select the Trap destinations tab in the Other configurations section.

   Use this tab to define one or more destinations for StorageGRID trap or inform notifications. When you enable the SNMP agent and select **Save**, StorageGRID sends notifications to each defined destination when alerts are triggered. Standard notifications are also sent for the supported MIB-II entities (for example, ifDown and coldStart).

   For details about what to enter, see Create trap destinations.

   - Optionally, update or remove the default trap community.

     If you remove the default trap community, you must first ensure that any existing trap destinations use a custom community string.

   - To add a trap destination, select **Create**.
   - To edit a trap destination, select the radio button, and select **Edit**.

- To remove a trap destination, select the radio button, and select **Remove**.
- To commit your changes, select **Save** at the bottom of the page.

6. To update agent addresses, select the Agent addresses tab in the Other configurations section.

   Use this tab to specify one or more "listening addresses." These are the StorageGRID addresses on which the SNMP agent can receive queries.

   For details about what to enter, see Create agent addresses.

   - To add an agent address, select **Create**.
   - To edit an agent address, select the radio button, and select **Edit**.
   - To remove an agent address, select the radio button, and select **Remove**.
   - To commit your changes, select **Save** at the bottom of the page.

7. To update USM users, select the USM users tab in the Other configurations section.

   Use this tab to define the USM users who are authorized to query the MIB or to receive traps and informs.

   For details about what to enter, see Create USM users.

   - To add a USM user, select **Create**.
   - To edit a USM user, select the radio button, and select **Edit**.

     The username for an existing USM user can't be changed. If you need to change a username, you must remove the user and create a new one.

     > If you add or remove a user's authoritative engine ID and that user is currently selected for a destination, you must edit or remove the destination. Otherwise, a validation error occurs when you save the SNMP agent configuration.

   - To remove a USM user, select the radio button, and select **Remove**.

     > If the user you removed is currently selected for a trap destination, you must edit or remove the destination. Otherwise, a validation error occurs when you save the SNMP agent configuration.

   - To commit your changes, select **Save** at the bottom of the page.

8. When you have updated the SNMP agent configuration, select **Save**.

Access MIB files

MIB files contain definitions and information about the properties of managed resources and services for the nodes in your grid. You can access MIB files that define the objects and notifications for StorageGRID. These files can be useful for monitoring your grid.

See Use SNMP monitoring for more information about SNMP and MIB files.

Access MIB files

Follow these steps to access the MIB files.

Steps

1. Select **CONFIGURATION** > **Monitoring** > **SNMP agent**.
2. On the SNMP agent page, select the file you want to download:
   - **NETAPP-STORAGEGRID-MIB.txt**: Defines the alert table and notifications (traps) accessible on all Admin Nodes.
   - **ES-NETAPP-06-MIB.mib**: Defines objects and notifications for E-Series-based appliances.
   - **MIB_1_10.zip**: Defines objects and notifications for appliances with a BMC interface.

   > You can also access MIB files at the following location on any StorageGRID node: `/usr/share/snmp/mibs`

3. To extract the StorageGRID OIDs from the MIB file:
   a. Get the OID of the root of the StorageGRID MIB:

   ```
   root@user-adm1:~ # snmptranslate -On -IR storagegrid
   ```

   Result: `.1.3.6.1.4.1.789.28669` (`28669` is always the OID for StorageGRID)

   b. Grep for the StorageGRID OID in the entire tree (using `paste` to join lines):

   ```
   root@user-adm1:~ # snmptranslate -Tso | paste -d " " - - | grep 28669
   ```

> The `snmptranslate` command has many options that are useful for exploring the MIB. This command is available on any StorageGRID node.

MIB file contents

All objects are under the StorageGRID OID.

Object name	Object ID (OID)	Description
`.iso.org.dod.internet.private.enterprises.netapp.storagegrid`	`.1.3.6.1.4.1.789.28669`	The MIB module for NetApp StorageGRID entities.

MIB objects

Object name	Object ID (OID)	Description
`activeAlertCount`	`.1.3.6.1.4.1.789.28669.1.3`	The number of active alerts in the activeAlertTable.
`activeAlertTable`	`.1.3.6.1.4.1.789.28669.1.4`	A table of active alerts in StorageGRID.

Object name	Object ID (OID)	Description
`activeAlertId`	`.1.3.6.1.4.1.789.28669.1.4.1.1`	The ID of the alert. Only unique in the current set of active alerts.
`activeAlertName`	`.1.3.6.1.4.1.789.28669.1.4.1.2`	The name of the alert.
`activeAlertInstance`	`.1.3.6.1.4.1.789.28669.1.4.1.3`	The name of the entity that generated the alert, typically the node name.
`activeAlertSeverity`	`.1.3.6.1.4.1.789.28669.1.4.1.4`	The severity of the alert.
`activeAlertStartTime`	`.1.3.6.1.4.1.789.28669.1.4.1.5`	The date and time the alert was triggered.

Notification types (Traps)

All notifications include the following variables as varbinds:

- activeAlertId
- activeAlertName
- activeAlertInstance
- activeAlertSeverity
- activeAlertStartTime

Notification type	Object ID (OID)	Description
`activeMinorAlert`	`.1.3.6.1.4.1.789.28669.0.6`	An alert with minor severity
`activeMajorAlert`	`.1.3.6.1.4.1.789.28669.0.7`	An alert with major severity
`activeCriticalAlert`	`.1.3.6.1.4.1.789.28669.0.8`	An alert with critical severity

Collect additional StorageGRID data

Use charts and graphs

You can use charts and reports to monitor the state of the StorageGRID system and troubleshoot problems.

> The Grid Manager is updated with each release and might not match the example screenshots on this page.

Types of charts

Charts and graphs summarize the values of specific StorageGRID metrics and attributes.

The Grid Manager dashboard includes cards that summarize available storage for the grid and each site.

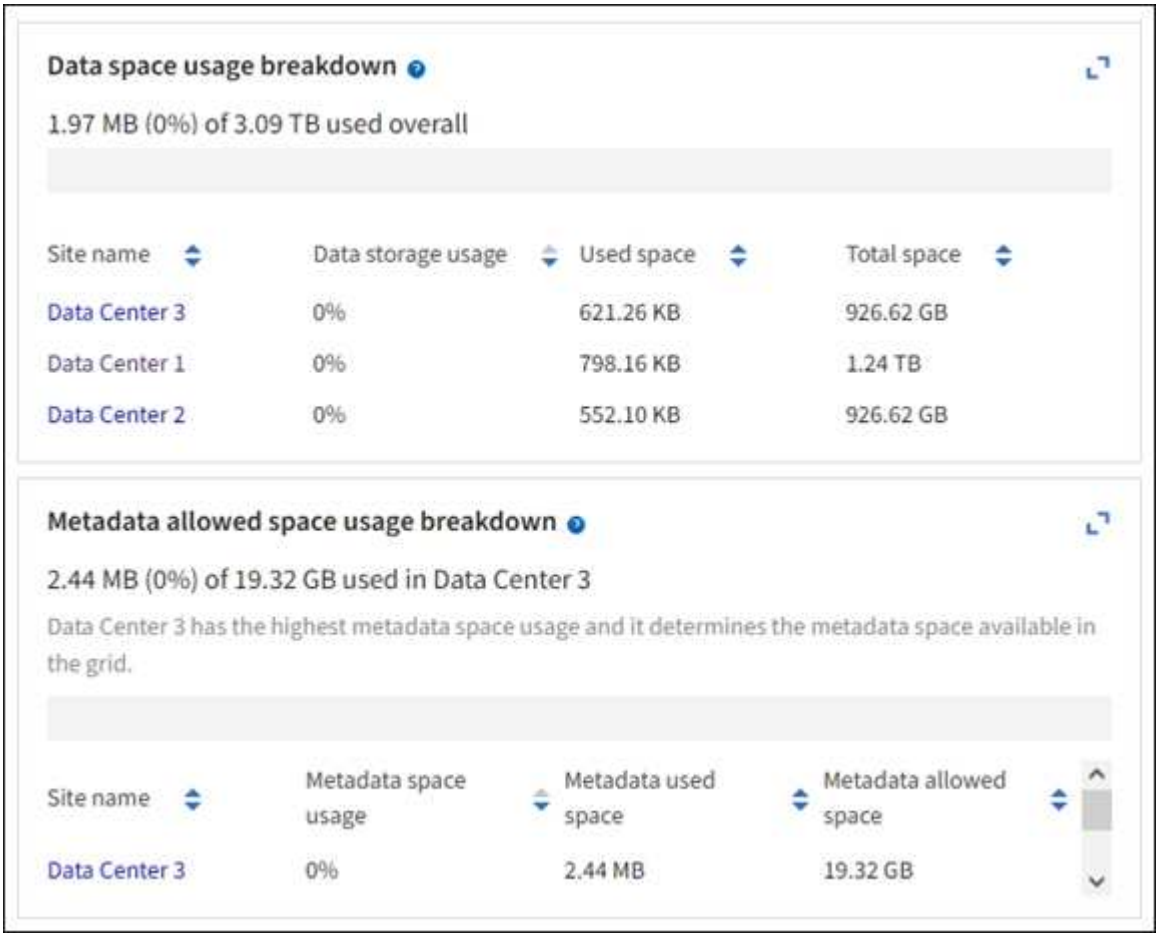

The Storage usage panel on the Tenant Manager dashboard displays the following:

- A list of the largest buckets (S3) or containers (Swift) for the tenant
- A bar chart that represents the relative sizes of the largest buckets or containers
- The total amount of space used and, if a quota is set, the amount and percentage of space remaining

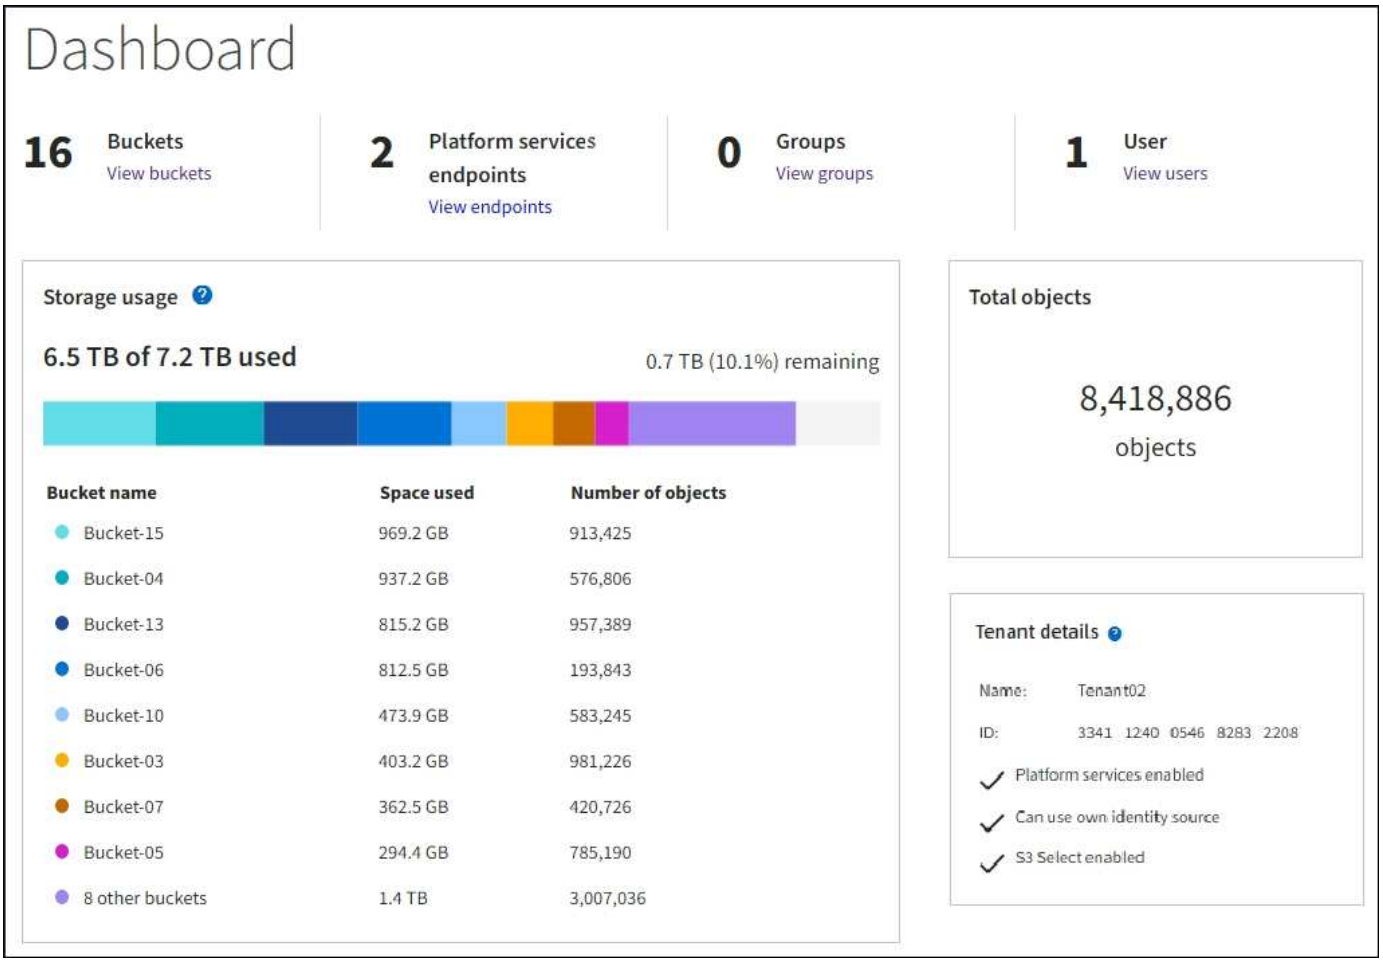

In addition, graphs that show how StorageGRID metrics and attributes change over time are available from the Nodes page and from the **SUPPORT** > **Tools** > **Grid topology** page.

There are four types of graphs:

- **Grafana charts**: Shown on the Nodes page, Grafana charts are used to plot the values of Prometheus metrics over time. For example, the **NODES** > **Network** tab for a Storage Node includes a Grafana chart for network traffic.

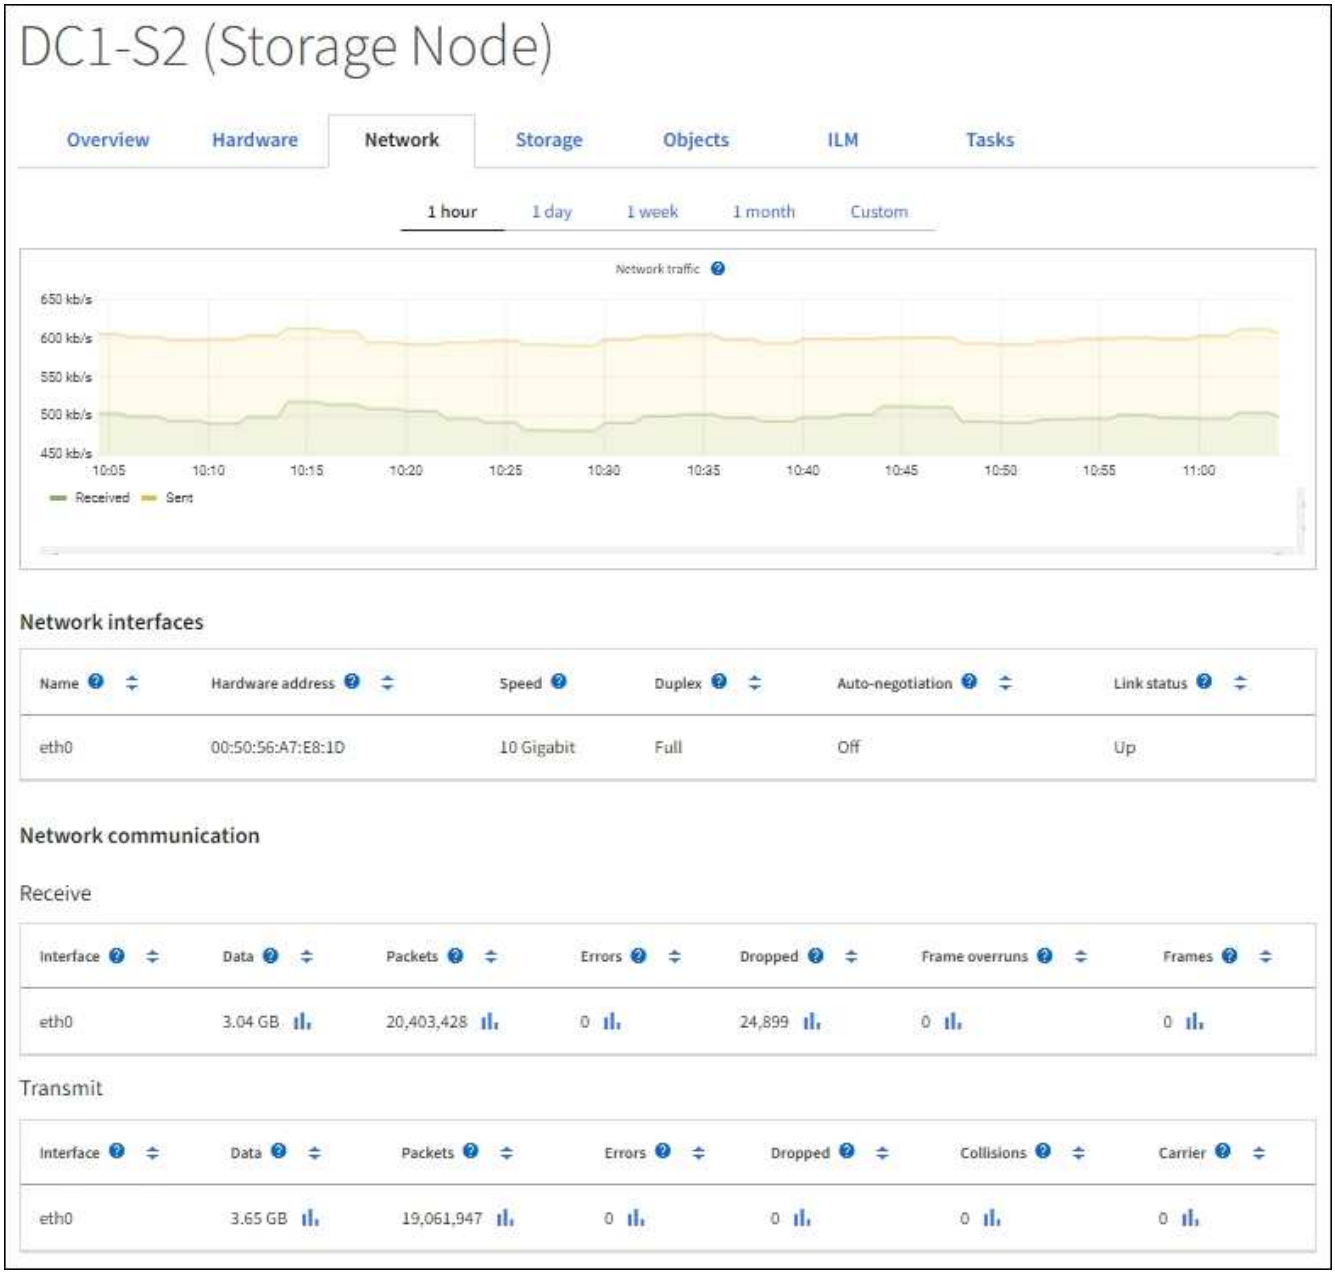

Grafana charts are also included on the pre-constructed dashboards available from the **SUPPORT** > **Tools** > **Metrics** page.

- **Line graphs**: Available from the Nodes page and from the **SUPPORT** > **Tools** > **Grid topology** page (select the chart icon after a data value), line graphs are used to plot the values of StorageGRID attributes that have a unit value (such as NTP Frequency Offset, in ppm). The changes in the value are plotted in regular data intervals (bins) over time.

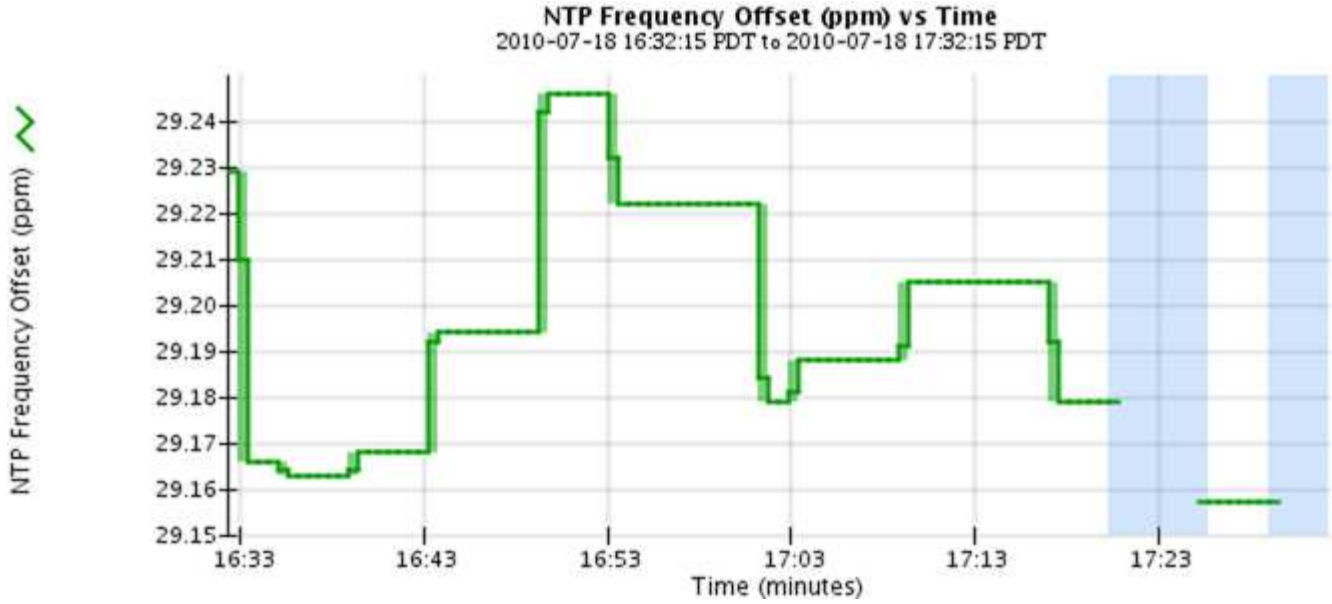

- **Area graphs**: Available from the Nodes page and from the **SUPPORT** > **Tools** > **Grid topology** page (select the chart icon after a data value), area graphs are used to plot volumetric attribute quantities, such as object counts or service load values. Area graphs are similar to line graphs, but include a light brown shading below the line. The changes in the value are plotted in regular data intervals (bins) over time.

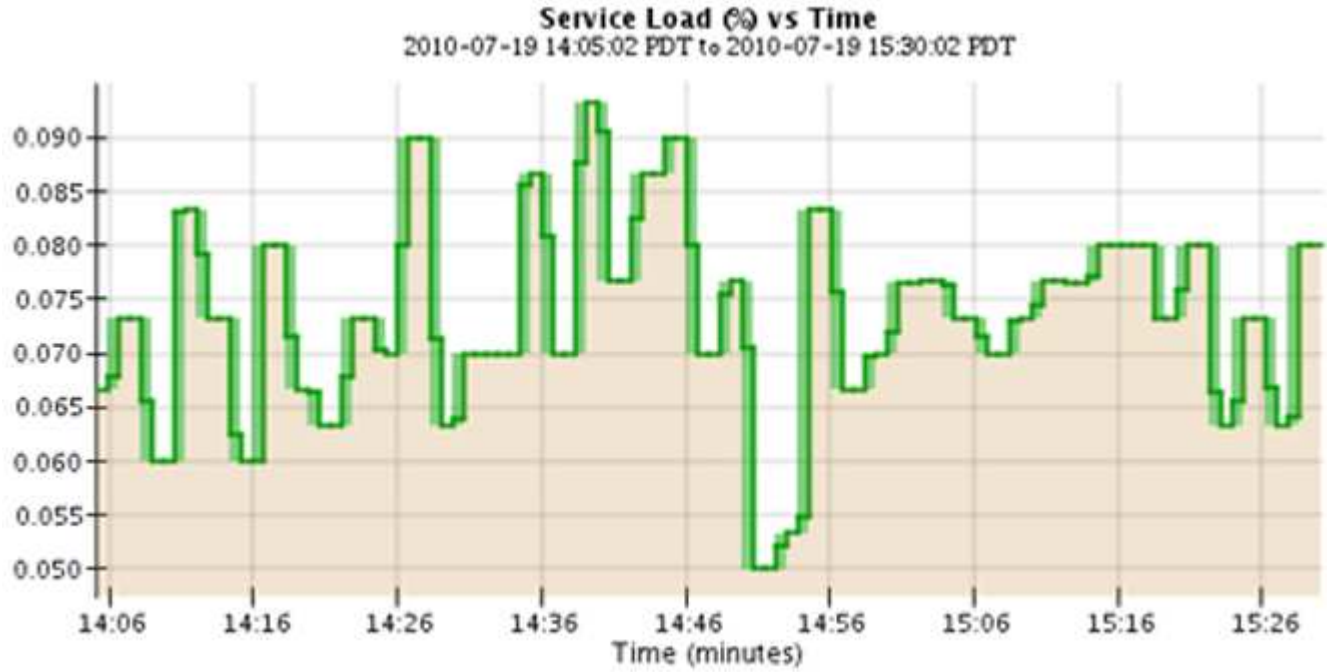

- Some graphs are denoted with a different type of chart icon and have a different format:

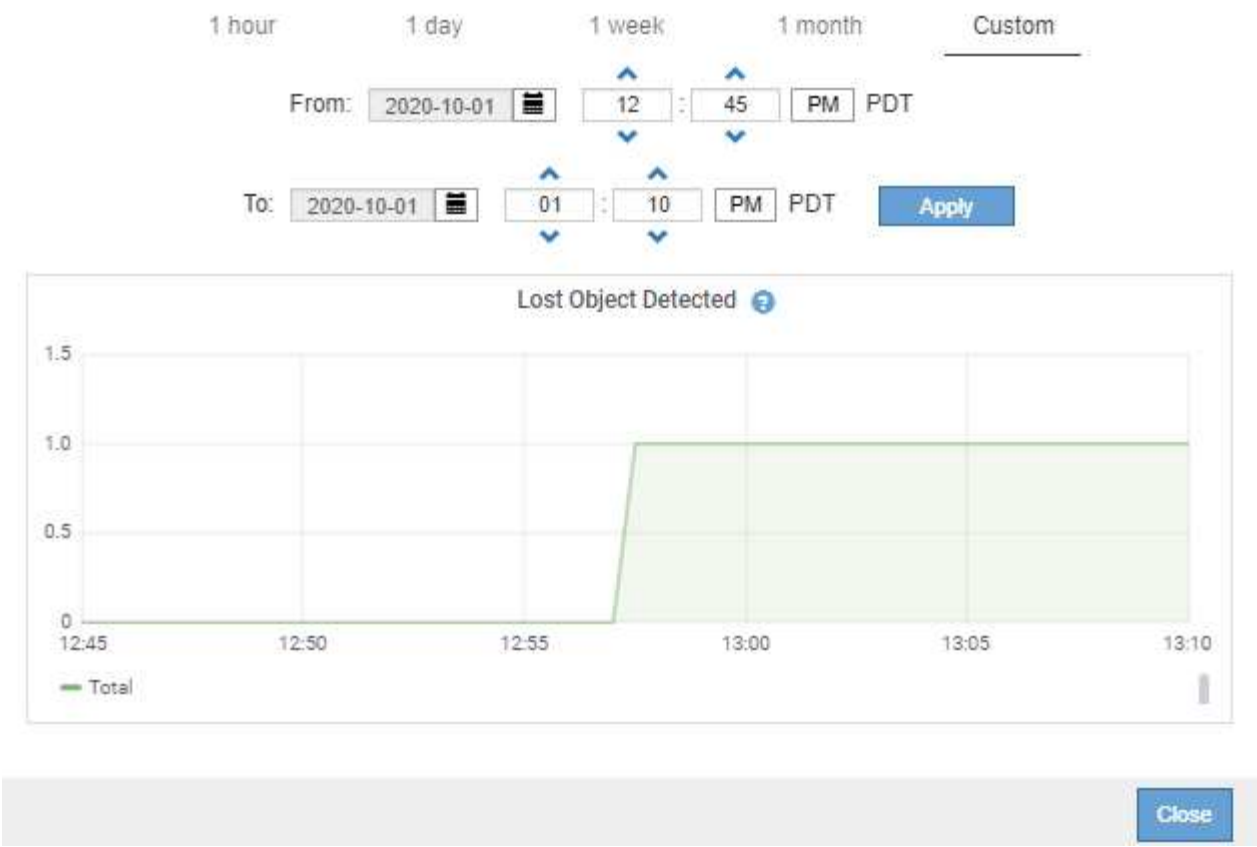

- **State graph**: Available from the **SUPPORT** > **Tools** > **Grid topology** page (select the chart icon after a data value), state graphs are used to plot attribute values that represent distinct states such as a service state that can be online, standby, or offline. State graphs are similar to line graphs, but the transition is discontinuous; that is, the value jumps from one state value to another.

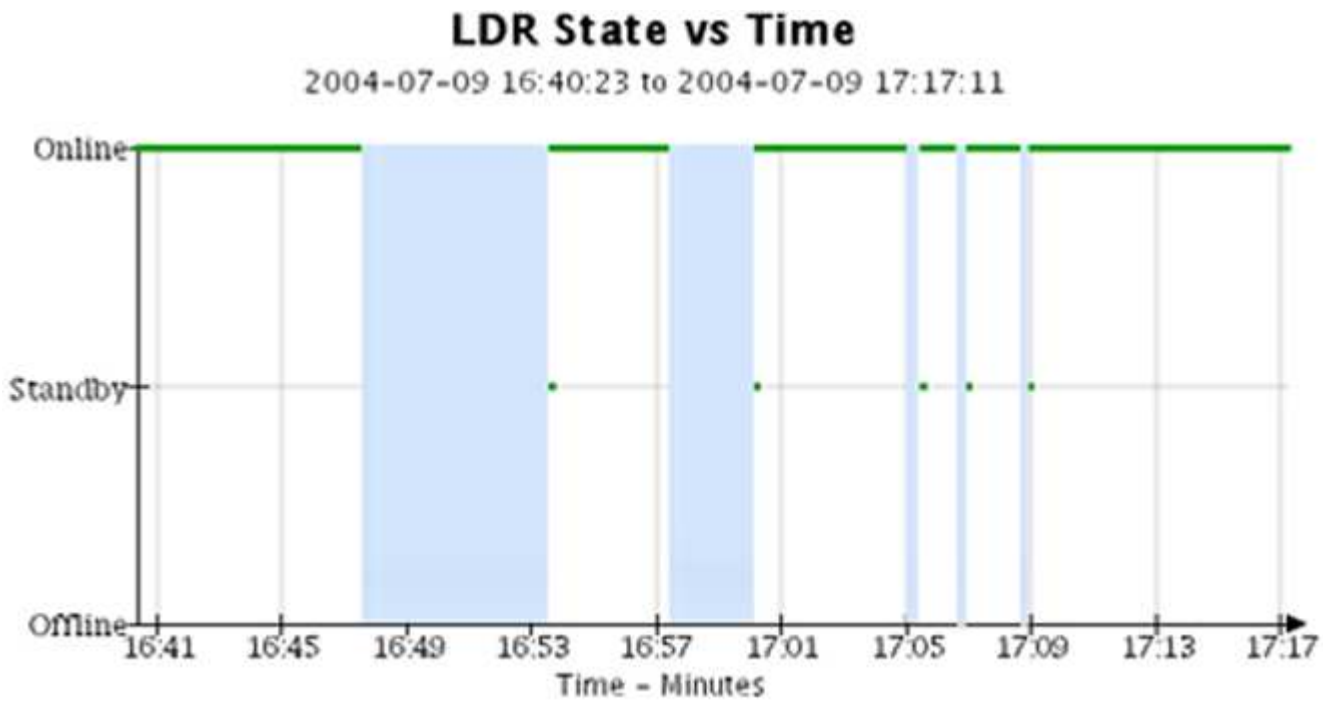

Related information

View the Nodes page

View the Grid Topology tree

Review support metrics

Chart legend

The lines and colors used to draw charts have specific meaning.

Example	Meaning
	Reported attribute values are plotted using dark green lines.
	Light green shading around dark green lines indicates that the actual values in that time range vary and have been "binned" for faster plotting. The dark line represents the weighted average. The range in light green indicates the maximum and minimum values within the bin. Light brown shading is used for area graphs to indicate volumetric data.
	Blank areas (no data plotted) indicate that the attribute values were unavailable. The background can be blue, gray, or a mixture of gray and blue, depending on the state of the service reporting the attribute.
	Light blue shading indicates that some or all of the attribute values at that time were indeterminate; the attribute was not reporting values because the service was in an unknown state.
	Gray shading indicates that some or all of the attribute values at that time were not known because the service reporting the attributes was administratively down.
	A mixture of gray and blue shading indicates that some of the attribute values at the time were indeterminate (because the service was in an unknown state), while others were not known because the service reporting the attributes was administratively down.

Display charts and graphs

The Nodes page contains the charts and graphs you should access regularly to monitor attributes such as storage capacity and throughput. In some cases, especially when working with technical support, you can use the **SUPPORT** > **Tools** > **Grid topology** page to access additional charts.

Before you begin

You must be signed in to the Grid Manager using a supported web browser.

Steps

1. Select **NODES**. Then, select a node, a site, or the entire grid.
2. Select the tab for which you want to view information.

   Some tabs include one or more Grafana charts, which are used to plot the values of Prometheus metrics over time. For example, the **NODES** > **Hardware** tab for a node includes two Grafana charts.

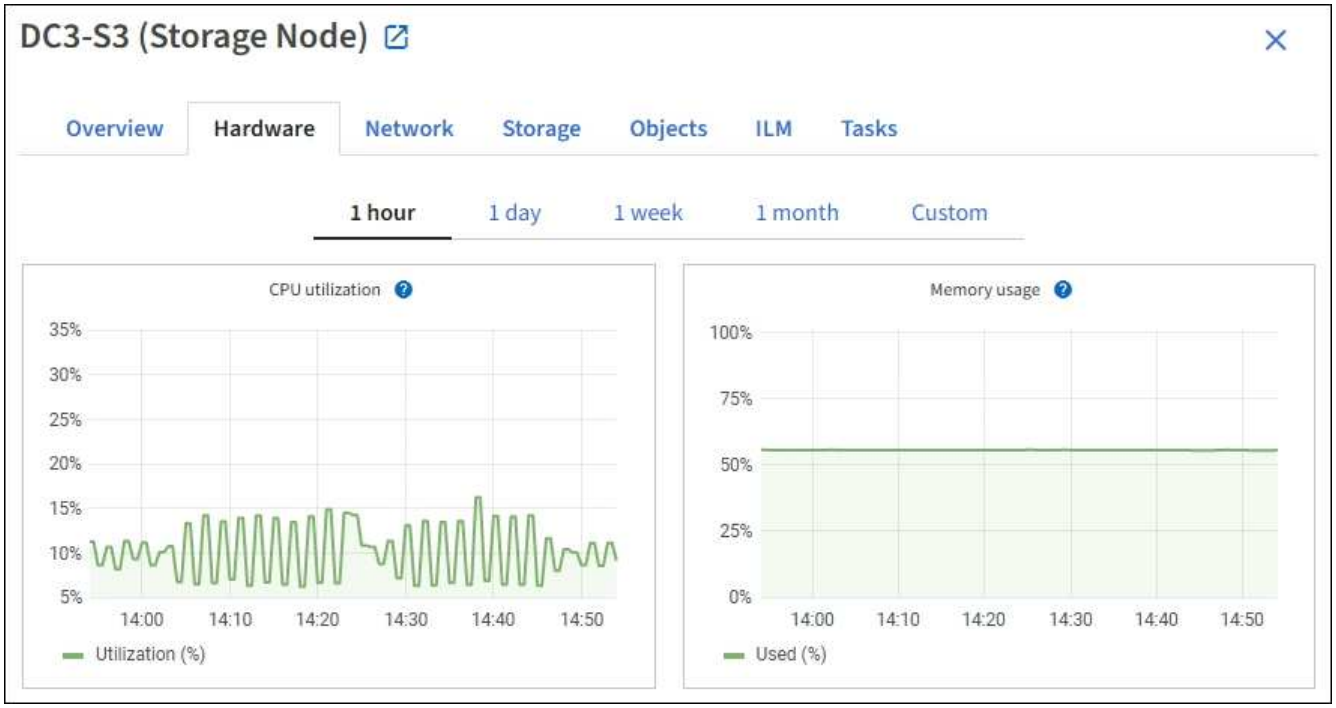

3. Optionally, position your cursor over the chart to see more detailed values for a particular point in time.

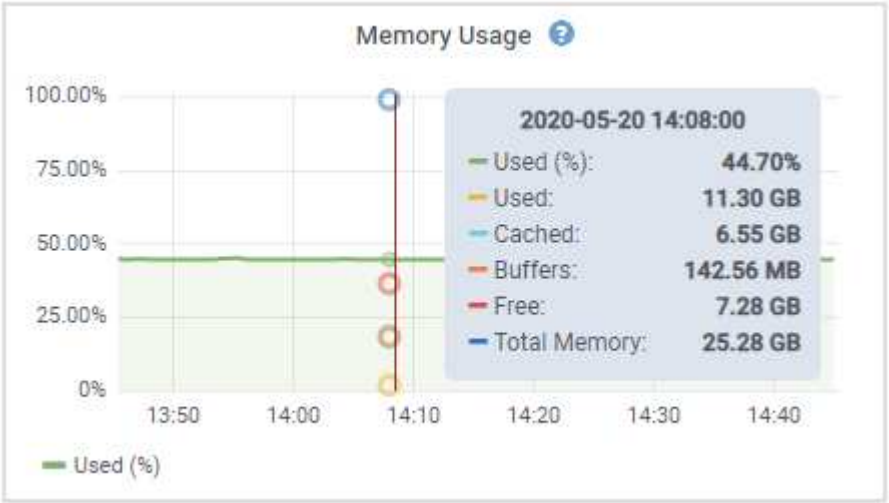

4. As required, you can often display a chart for a specific attribute or metric. From the table on the Nodes page, select the chart icon to the right of the attribute name.

> Charts aren't available for all metrics and attributes.

Example 1: From the Objects tab for a Storage Node, you can select the chart icon to see the total number of successful metadata store queries for the Storage Node.


Reports (Charts): DDS (DC1-S1) - Data Store
Attribute:
Total Queries - Successful
Quick Query:
Last Hour
Update
Vertical Scaling:
Raw Data:
YYYY/MM/DD HH:MM:SS
Start Date:
2021/09/24 13:24:01
End Date:
2021/09/24 14:24:01
Total Queries - Successful vs Time
2021-09-24 13:24:01 MDT to 2021-09-24 14:24:01 MDT
Total Queries - Successful
32,000
31,000
30,000
29,000
13:25
13:35
13:45
13:55
14:05
14:15
Time (minutes)
Close

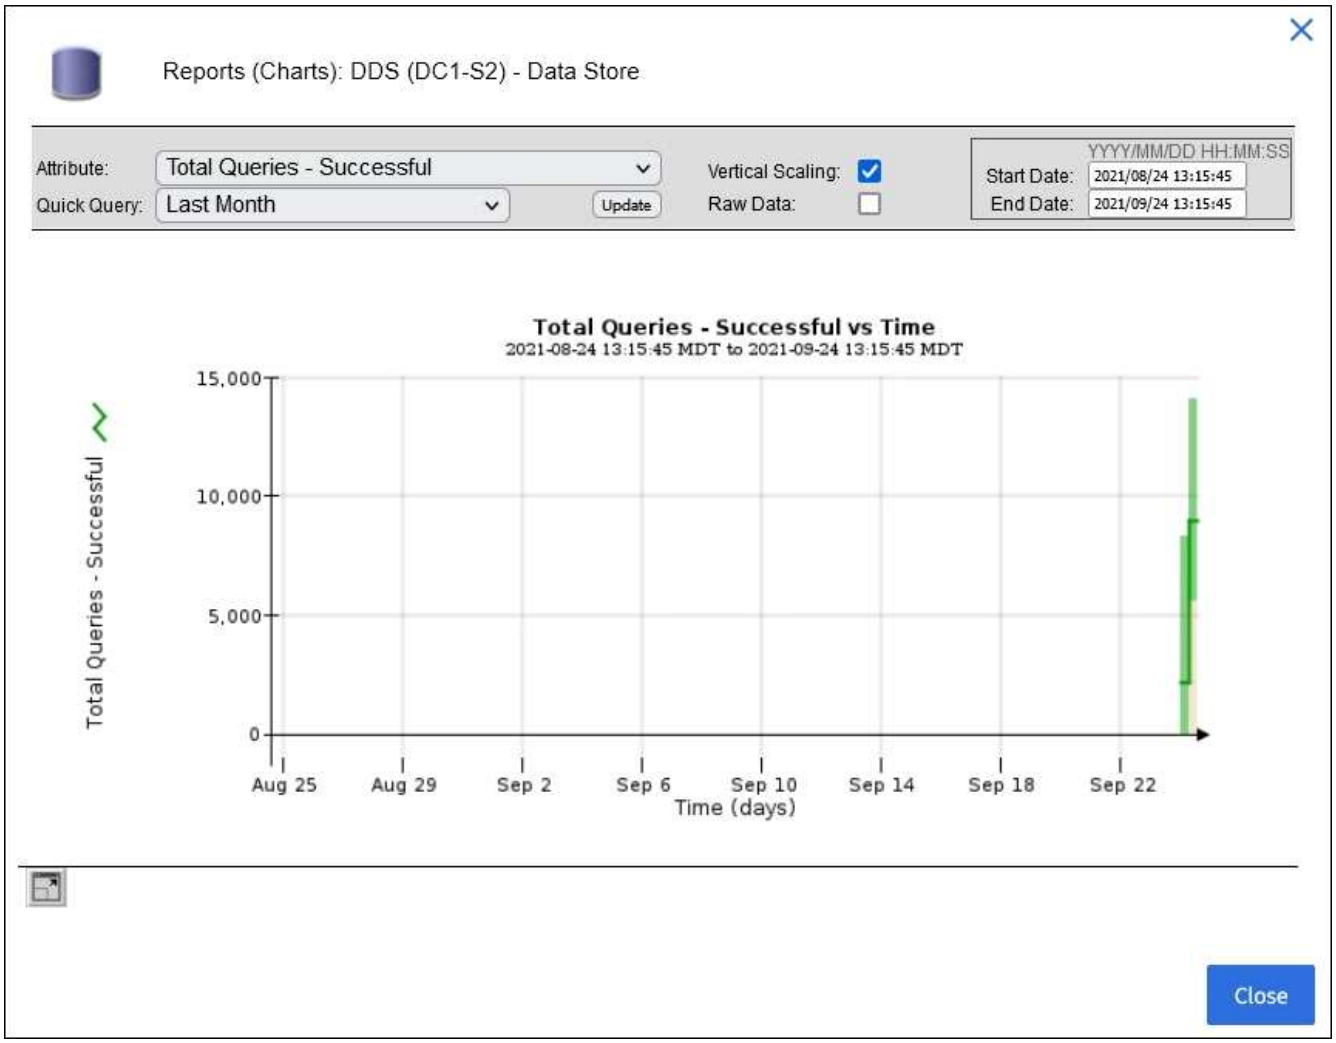

Example 2: From the Objects tab for a Storage Node, you can select the chart icon to see the Grafana graph of the count of lost objects detected over time.

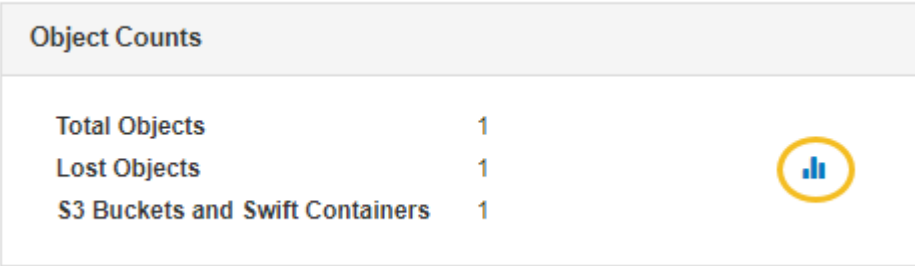

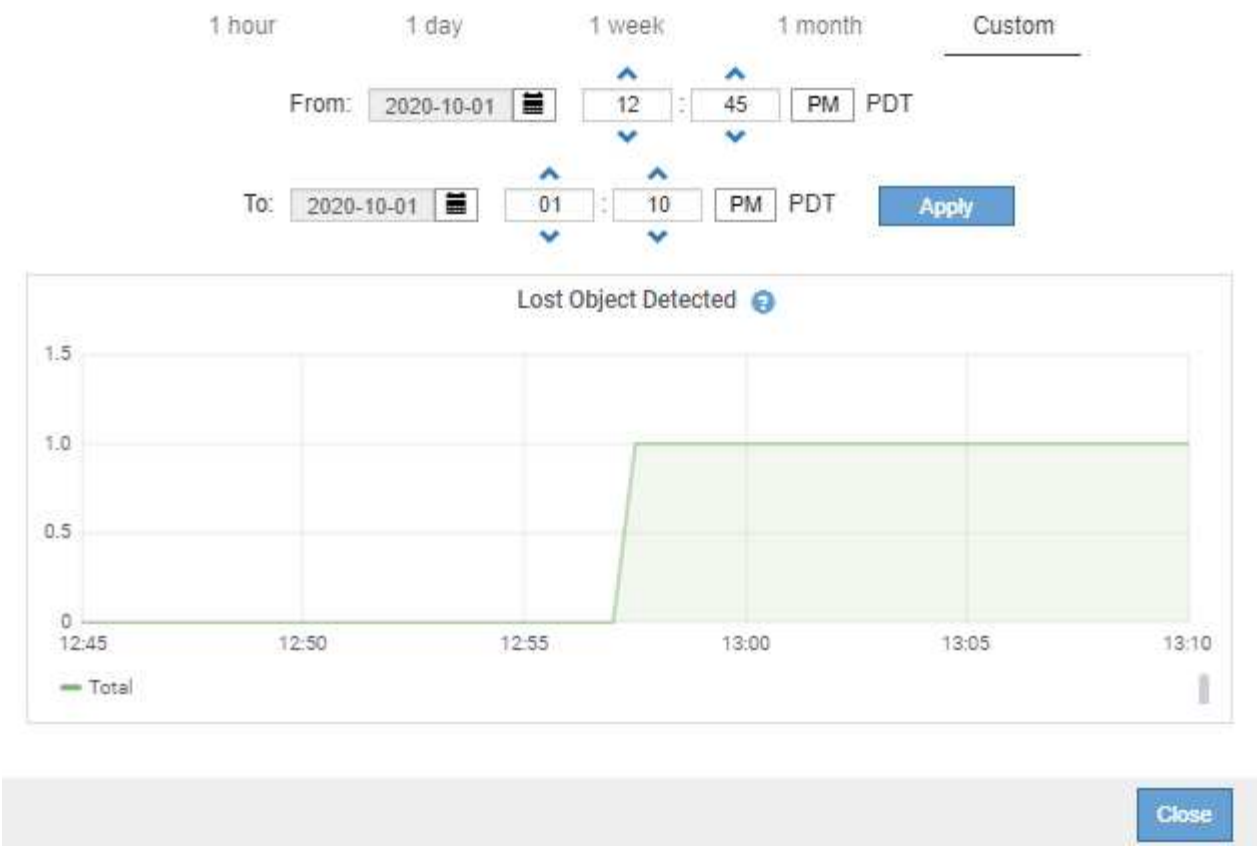

5. To display charts for attributes that aren't shown on the Node page, select **SUPPORT** > **Tools** > **Grid topology**.
6. Select ***grid node*** > ***component or service*** > **Overview** > **Main**.

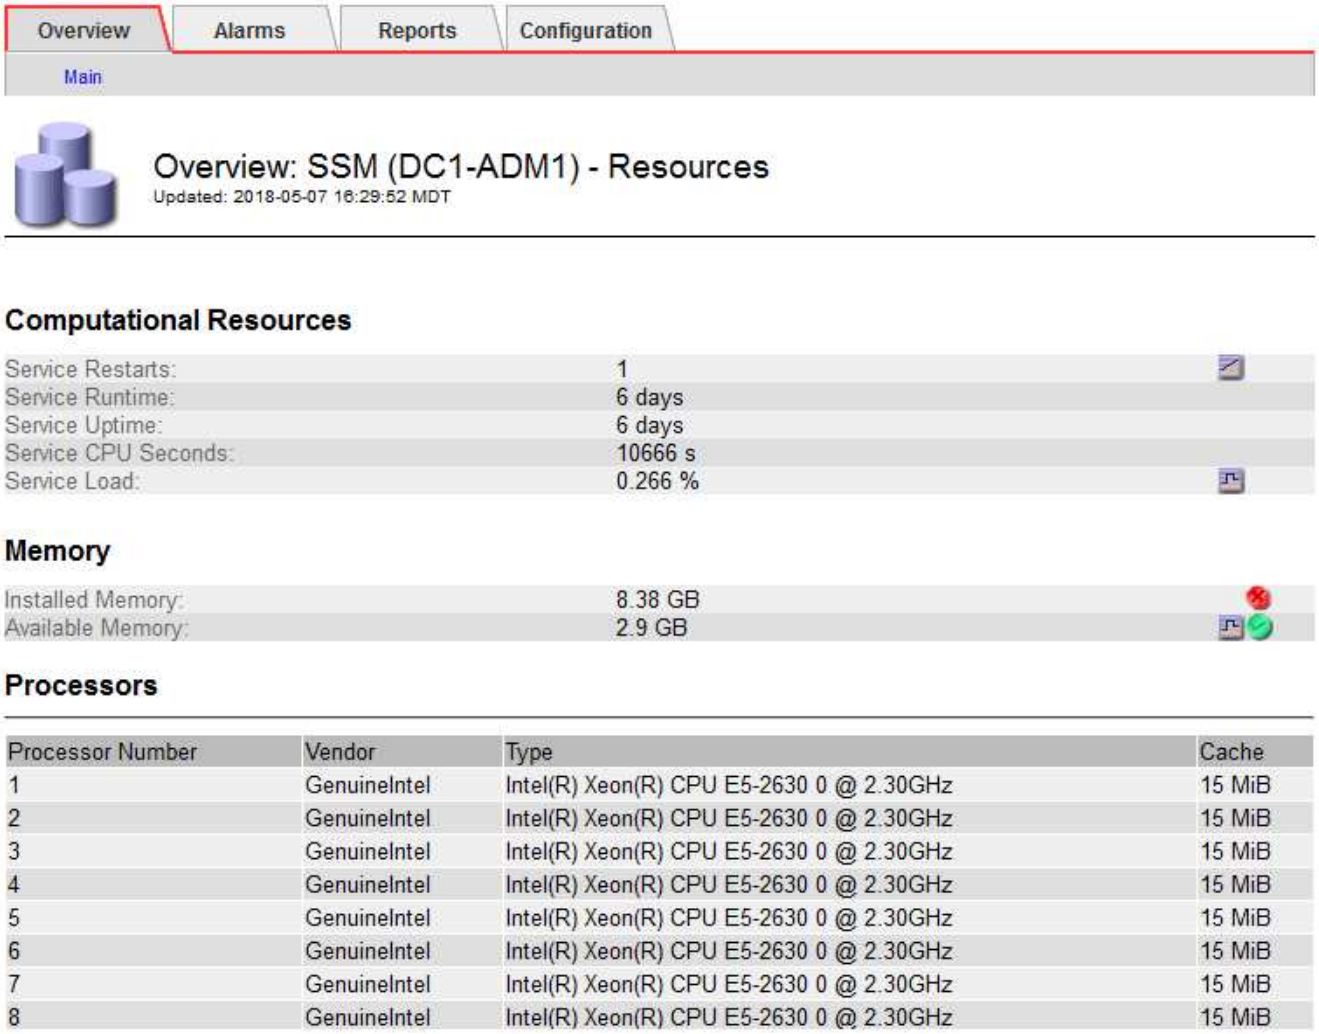

7. Select the chart icon next to the attribute.

   The display automatically changes to the **Reports** > **Charts** page. The chart displays the attribute's data over the past day.

Generate charts

Charts display a graphical representation of attribute data values. You can report on a data center site, grid node, component, or service.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select ***grid node*** > ***component or service*** > **Reports** > **Charts**.
3. Select the attribute to report on from the **Attribute** drop-down list.
4. To force the Y-axis to start at zero, clear the **Vertical Scaling** checkbox.
5. To show values at full precision, select the **Raw Data** checkbox, or to round values to a maximum of three

decimal places (for example, for attributes reported as percentages), clear the **Raw Data** checkbox.

6. Select the time period to report on from the **Quick Query** drop-down list.

   Select the Custom Query option to select a specific time range.

   The chart appears after a few moments. Allow several minutes for tabulation of long time ranges.

7. If you selected Custom Query, customize the time period for the chart by entering the **Start Date** and **End Date**.

   Use the format *YYYY/MM/DDHH:MM:SS* in local time. Leading zeros are required to match the format. For example, 2017/4/6 7:30:00 fails validation. The correct format is: 2017/04/06 07:30:00.

8. Select **Update**.

   A chart is generated after a few seconds. Allow several minutes for tabulation of long time ranges. Depending on the length of time set for the query, either a raw text report or aggregate text report is displayed.

Use text reports

Text reports display a textual representation of attribute data values that have been processed by the NMS service. There are two types of reports generated depending on the time period you are reporting on: raw text reports for periods less than a week, and aggregate text reports for time periods greater than a week.

Raw text reports

A raw text report displays details about the selected attribute:

- Time Received: Local date and time that a sample value of an attribute's data was processed by the NMS service.
- Sample Time: Local date and time that an attribute value was sampled or changed at the source.
- Value: Attribute value at sample time.

Text Results for Services: Load - System Logging

2010-07-18 15:58:39 PDT To 2010-07-19 15:58:39 PDT

Time Received	Sample Time	Value
2010-07-19 15:58:09	2010-07-19 15:58:09	0.016 %
2010-07-19 15:56:06	2010-07-19 15:56:06	0.024 %
2010-07-19 15:54:02	2010-07-19 15:54:02	0.033 %
2010-07-19 15:52:00	2010-07-19 15:52:00	0.016 %
2010-07-19 15:49:57	2010-07-19 15:49:57	0.008 %
2010-07-19 15:47:54	2010-07-19 15:47:54	0.024 %
2010-07-19 15:45:50	2010-07-19 15:45:50	0.016 %
2010-07-19 15:43:47	2010-07-19 15:43:47	0.024 %
2010-07-19 15:41:43	2010-07-19 15:41:43	0.032 %
2010-07-19 15:39:40	2010-07-19 15:39:40	0.024 %
2010-07-19 15:37:37	2010-07-19 15:37:37	0.008 %
2010-07-19 15:35:34	2010-07-19 15:35:34	0.016 %
2010-07-19 15:33:31	2010-07-19 15:33:31	0.024 %
2010-07-19 15:31:27	2010-07-19 15:31:27	0.032 %
2010-07-19 15:29:24	2010-07-19 15:29:24	0.032 %
2010-07-19 15:27:21	2010-07-19 15:27:21	0.049 %
2010-07-19 15:25:18	2010-07-19 15:25:18	0.024 %
2010-07-19 15:21:12	2010-07-19 15:21:12	0.016 %
2010-07-19 15:19:09	2010-07-19 15:19:09	0.008 %
2010-07-19 15:17:07	2010-07-19 15:17:07	0.016 %

Aggregate text reports

An aggregate text report displays data over a longer period of time (usually a week) than a raw text report. Each entry is the result of summarizing multiple attribute values (an aggregate of attribute values) by the NMS service over time into a single entry with average, maximum, and minimum values that are derived from the aggregation.

Each entry displays the following information:

- Aggregate Time: Last local date and time that the NMS service aggregated (collected) a set of changed attribute values.
- Average Value: The average of the attribute's value over the aggregated time period.
- Minimum Value: The minimum value over the aggregated time period.
- Maximum Value: The maximum value over the aggregated time period.

Text Results for Attribute Send to Relay Rate

2010-07-11 16:02:46 PDT To 2010-07-19 16:02:46 PDT

Aggregate Time	Average Value	Minimum Value	Maximum Value
2010-07-19 15:59:52	0.271072196 Messages/s	0.266649743 Messages/s	0.274983464 Messages/s
2010-07-19 15:53:52	0.275585378 Messages/s	0.266562352 Messages/s	0.283302736 Messages/s
2010-07-19 15:49:52	0.279315709 Messages/s	0.233318712 Messages/s	0.333313579 Messages/s
2010-07-19 15:43:52	0.28181323 Messages/s	0.241651024 Messages/s	0.374976601 Messages/s
2010-07-19 15:39:52	0.284233141 Messages/s	0.249982001 Messages/s	0.324971987 Messages/s
2010-07-19 15:33:52	0.325752083 Messages/s	0.266641993 Messages/s	0.358306197 Messages/s
2010-07-19 15:29:52	0.278531507 Messages/s	0.274984766 Messages/s	0.283320999 Messages/s
2010-07-19 15:23:52	0.281437642 Messages/s	0.274981961 Messages/s	0.291577735 Messages/s
2010-07-19 15:17:52	0.261563307 Messages/s	0.258318006 Messages/s	0.266655787 Messages/s
2010-07-19 15:13:52	0.265159147 Messages/s	0.258318557 Messages/s	0.26663986 Messages/s

Generate text reports

Text reports display a textual representation of attribute data values that have been processed by the NMS service. You can report on a data center site, grid node, component, or service.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

For attribute data that is expected to be continuously changing, this attribute data is sampled by the NMS service (at the source) at regular intervals. For attribute data that changes infrequently (for example, data based on events such as state or status changes), an attribute value is sent to the NMS service when the value changes.

The type of report displayed depends on the configured time period. By default, aggregate text reports are generated for time periods longer than one week.

Gray text indicates the service was administratively down during the time it was sampled. Blue text indicates the service was in an unknown state.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select ***grid node*** > ***component or service*** > **Reports** > **Text**.
3. Select the attribute to report on from the **Attribute** drop-down list.
4. Select the number of results per page from the **Results per Page** drop-down list.
5. To round values to a maximum of three decimal places (for example, for attributes reported as percentages), clear the **Raw Data** checkbox.
6. Select the time period to report on from the **Quick Query** drop-down list.

   Select the Custom Query option to select a specific time range.

   The report appears after a few moments. Allow several minutes for tabulation of long time ranges.

7. If you selected Custom Query, you need to customize the time period to report on by entering the **Start Date** and **End Date**.

   Use the format `YYYY/MM/DDHH:MM:SS` in local time. Leading zeros are required to match the format. For example, 2017/4/6 7:30:00 fails validation. The correct format is: 2017/04/06 07:30:00.

8. Click **Update**.

   A text report is generated after a few moments. Allow several minutes for tabulation of long time ranges. Depending on the length of time set for the query, either a raw text report or aggregate text report is displayed.

Export text reports

Exported text reports open a new browser tab, which enables you to select and copy the data.

About this task

The copied data can then be saved into a new document (for example, a spreadsheet) and used to analyze the performance of the StorageGRID system.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Create a text report.
3. Click *Export*.

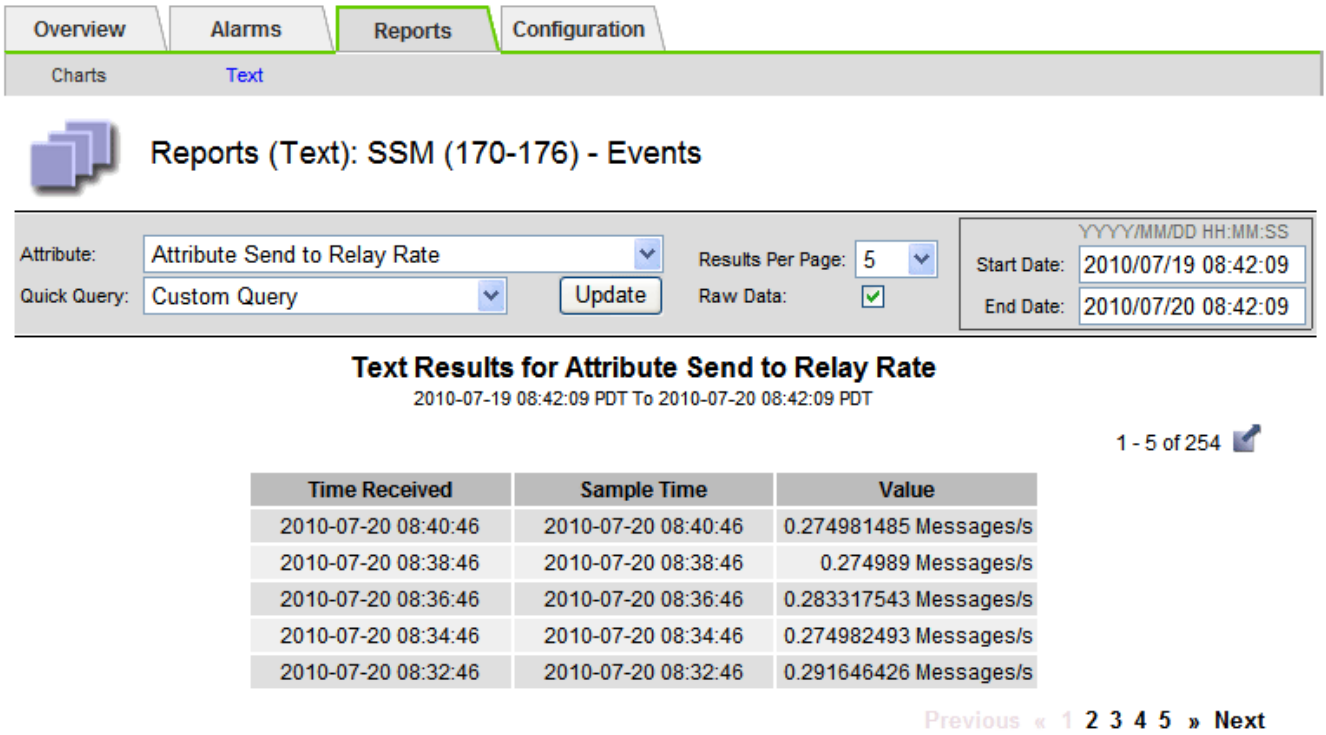

   The Export Text Report window opens displaying the report.

```
Grid ID: 000 000
OID: 2.16.124.113590.2.1.400019.1.1.1.1.16996732.200
Node Path: Site/170-176/SSM/Events
Attribute: Attribute Send to Relay Rate (ABSR)
Query Start Date: 2010-07-19 08:42:09 PDT
Query End Date: 2010-07-20 08:42:09 PDT
Time Received,Time Received (Epoch),Sample Time,Sample Time (Epoch),Value,Type
2010-07-20 08:40:46,1279640446559000,2010-07-20 08:40:46,1279640446537209,0.274981485 Messages/s,U
2010-07-20 08:38:46,1279640326561000,2010-07-20 08:38:46,1279640326529124,0.274989 Messages/s,U
2010-07-20 08:36:46,1279640206556000,2010-07-20 08:36:46,1279640206524330,0.283317543 Messages/s,U
2010-07-20 08:34:46,1279640086540000,2010-07-20 08:34:46,1279640086517645,0.274982493 Messages/s,U
2010-07-20 08:32:46,1279639966543000,2010-07-20 08:32:46,1279639966510022,0.291646426 Messages/s,U
2010-07-20 08:30:46,1279639846561000,2010-07-20 08:30:46,1279639846501672,0.308315369 Messages/s,U
2010-07-20 08:28:46,1279639726527000,2010-07-20 08:28:46,1279639726494673,0.291657509 Messages/s,U
2010-07-20 08:26:46,1279639606526000,2010-07-20 08:26:46,1279639606490890,0.266627739 Messages/s,U
2010-07-20 08:24:46,1279639486495000,2010-07-20 08:24:46,1279639486473368,0.258318523 Messages/s,U
2010-07-20 08:22:46,1279639366480000,2010-07-20 08:22:46,1279639366466497,0.274985902 Messages/s,U
2010-07-20 08:20:46,1279639246469000,2010-07-20 08:20:46,1279639246460346,0.283253871 Messages/s,U
2010-07-20 08:18:46,1279639126469000,2010-07-20 08:18:46,1279639126426669,0.274982804 Messages/s,U
2010-07-20 08:16:46,1279639006437000,2010-07-20 08:16:46,1279639006419168,0.283315503 Messages/s,U
```

4. Select and copy the contents of the Export Text Report window.

   This data can now be pasted into a third-party document such as a spreadsheet.

Monitor PUT and GET performance

You can monitor the performance of certain operations, such as object store and retrieve, to help identify changes that might require further investigation.

About this task

To monitor PUT and GET performance, you can run S3 and Swift commands directly from a workstation or by using the open-source S3tester application. Using these methods allows you to assess performance independently of factors that are external to StorageGRID, such as issues with a client application or issues with an external network.

When performing tests of PUT and GET operations, use the following guidelines:

- Use object sizes comparable to the objects that you typically ingest into your grid.
- Perform operations against both local and remote sites.

Messages in the audit log indicate the total time required to run certain operations. For example, to determine the total processing time for an S3 GET request, you can review the value of the TIME attribute in the SGET audit message. You can also find the TIME attribute in the audit messages for the following operations:

- **S3**: DELETE, GET, HEAD, Metadata Updated, POST, PUT
- **Swift**: DELETE, GET, HEAD, PUT

When analyzing results, look at the average time required to satisfy a request, as well as the overall throughput that you can achieve. Repeat the same tests regularly and record the results, so that you can identify trends that might require investigation.

- You can download S3tester from github.

Monitor object verification operations

The StorageGRID system can verify the integrity of object data on Storage Nodes, checking for both corrupt and missing objects.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.

About this task

Two verification processes work together to ensure data integrity:

- **Background verification** runs automatically, continuously checking the correctness of object data.

  Background verification automatically and continuously checks all Storage Nodes to determine if there are corrupt copies of replicated and erasure-coded object data. If problems are found, the StorageGRID system automatically attempts to replace the corrupt object data from copies stored elsewhere in the system. Background verification does not run on Archive Nodes or on objects in a Cloud Storage Pool.

  > The **Unidentified corrupt object detected** alert is triggered if the system detects a corrupt object that can't be corrected automatically.

- **Object existence check** can be triggered by a user to more quickly verify the existence (although not the correctness) of object data.

  Object existence check verifies whether all expected replicated copies of objects and erasure-coded fragments exist on a Storage Node. Object existence check provides a way to verify the integrity of storage devices, especially if a recent hardware issue could have affected data integrity.

You should review the results from background verifications and object existence checks regularly. Investigate any instances of corrupt or missing object data immediately to determine the root cause.

Steps

1. Review the results from background verifications:
   a. Select **NODES** > ***Storage Node*** > **Objects**.
   b. Check the verification results:
      - To check replicated object data verification, look at the attributes in the Verification section.

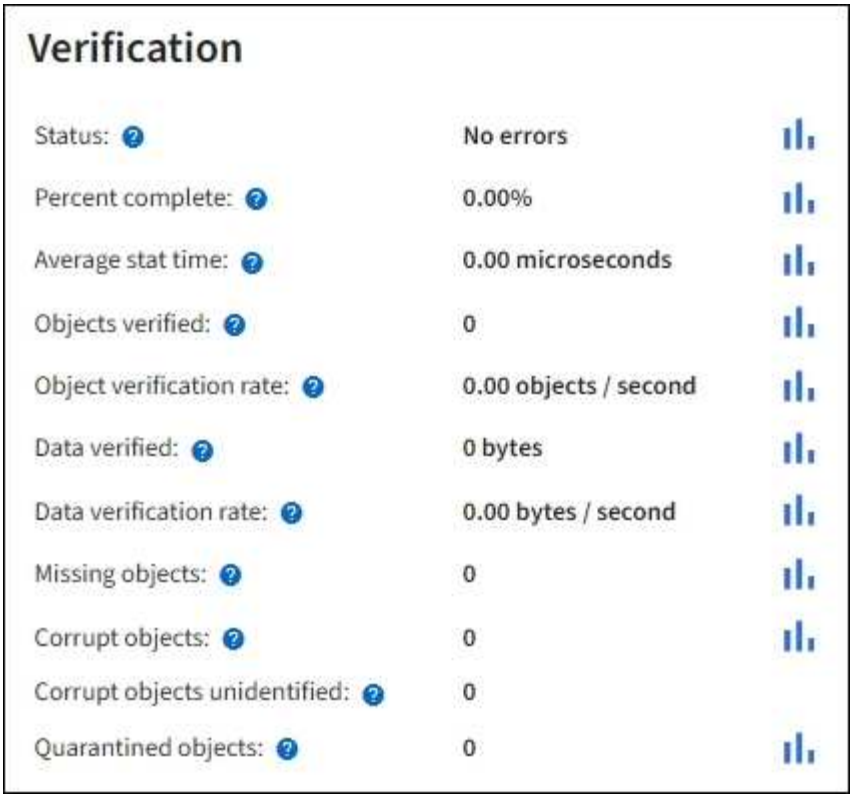

- To check erasure-coded fragment verification, select ***Storage Node*** > **ILM** and look at the attributes in the Erasure coding verification section.

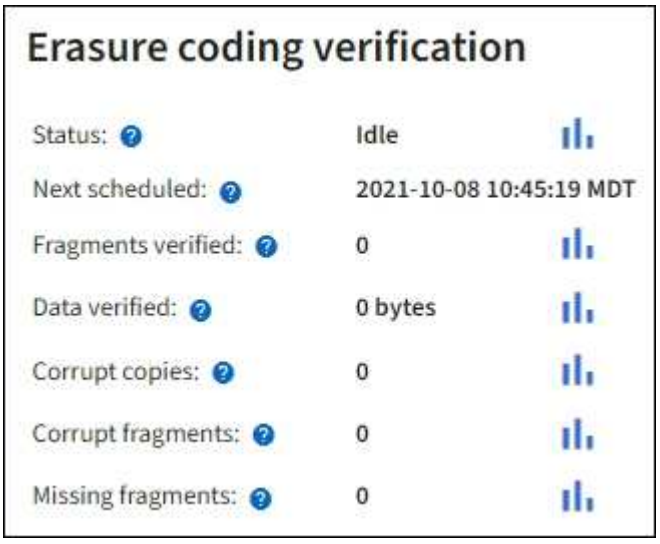

Select the question mark ? next to an attribute's name to display help text.

2. Review the results from object existence check jobs:
   a. Select **MAINTENANCE** > **Object existence check** > **Job history**.
   b. Scan the Missing object copies detected column. If any jobs resulted in 100 or more missing object copies and the **Objects lost** alert has been triggered, contact technical support.

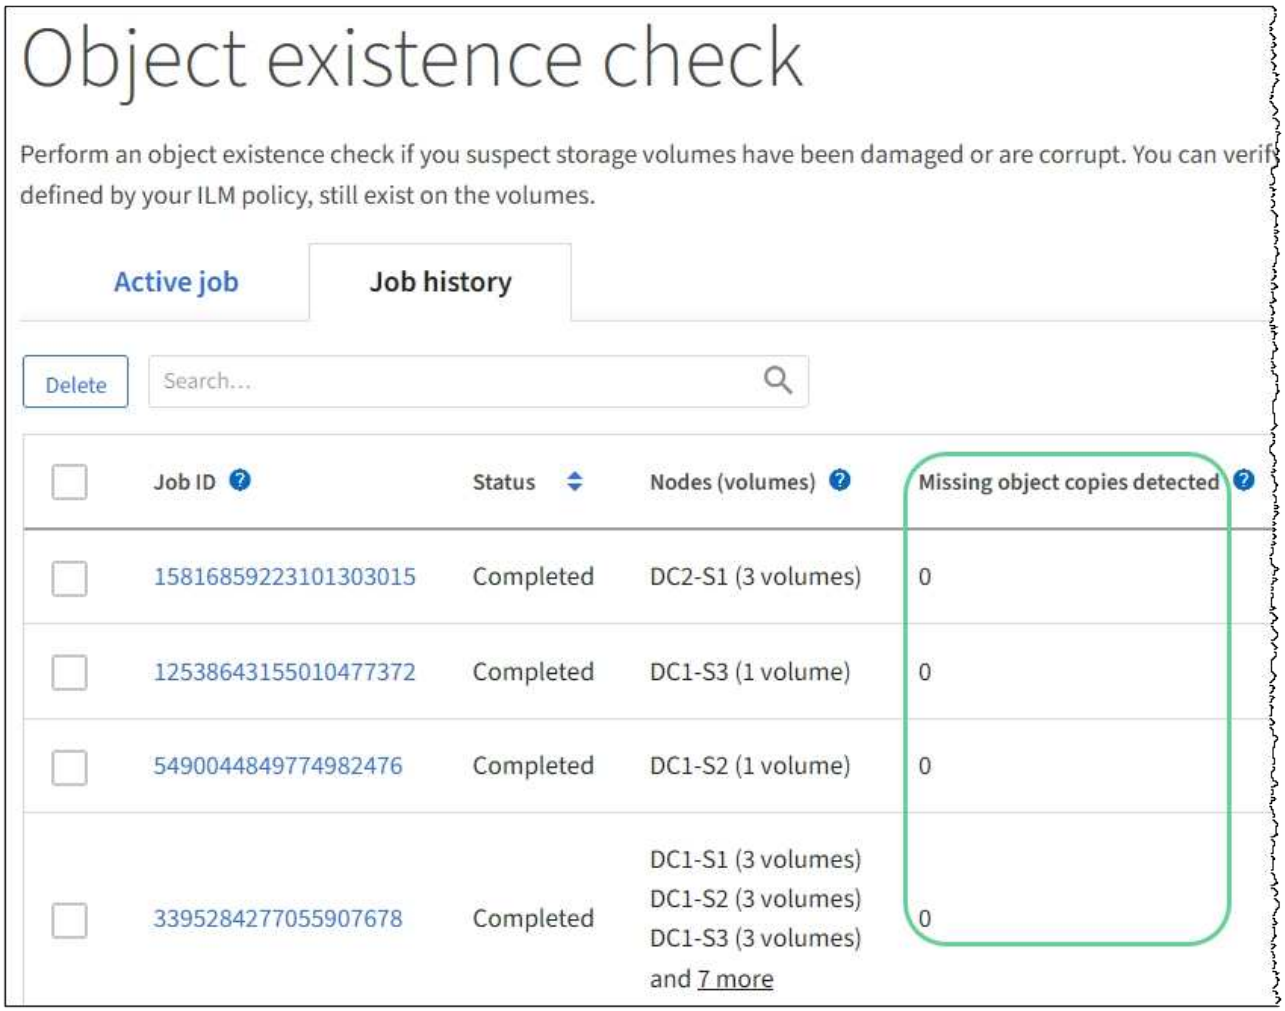

Monitor events

You can monitor events that are detected by a grid node, including custom events that you have created to track events that are logged to the syslog server. The Last Event message shown in the Grid Manager provides more information about the most recent event.

Event messages are also listed in the `/var/local/log/bycast-err.log` log file. See the Log files reference.

The SMTT (Total events) alarm can be repeatedly triggered by issues such as network problems, power outages or upgrades. This section has information about investigating events so that you can better understand why these alarms have occurred. If an event occurred because of a known issue, it is safe to reset the event counters.

Steps

1. Review the system events for each grid node:
   a. Select **SUPPORT** > **Tools** > **Grid topology**.
   b. Select ***site*** > ***grid node*** > **SSM** > **Events** > **Overview** > **Main**.
2. Generate a list of previous event messages to help isolate issues that occurred in the past:

a. Select **SUPPORT** > **Tools** > **Grid topology**.

b. Select ***site*** > ***grid node*** > **SSM** > **Events** > **Reports**.

c. Select **Text**.

The **Last Event** attribute is not shown in the charts view. To view it:

d. Change **Attribute** to **Last Event**.

e. Optionally, select a time period for **Quick Query**.

f. Select **Update**.

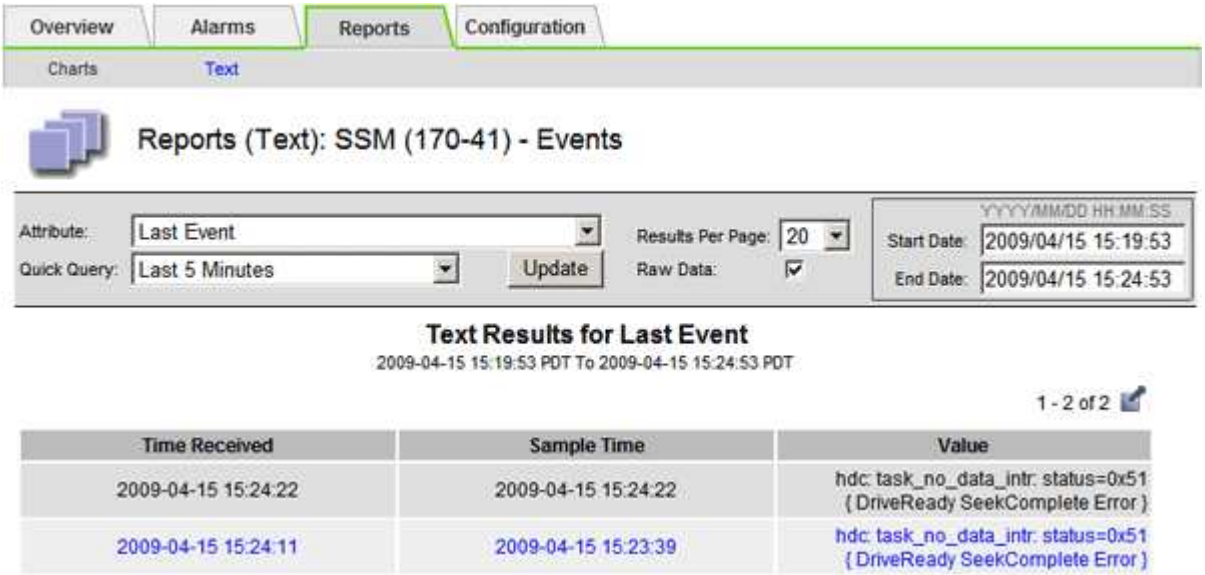

Create custom syslog events

Custom events allow you to track all kernel, daemon, error and critical level user events logged to the syslog server. A custom event can be useful for monitoring the occurrence of system log messages (and thus network security events and hardware faults).

About this task

Consider creating custom events to monitor recurring problems. The following considerations apply to custom events.

- After a custom event is created, every occurrence of it is monitored.
- To create a custom event based on keywords in the `/var/local/log/messages` files, the logs in those files must be:
  - Generated by the kernel
  - Generated by daemon or user program at the error or critical level

Note: Not all entries in the `/var/local/log/messages` files will be matched unless they satisfy the requirements stated above.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Custom events**.
2. Click **Edit** (or **Insert** if this is not the first event).

3. Enter a custom event string, for example, shutdown

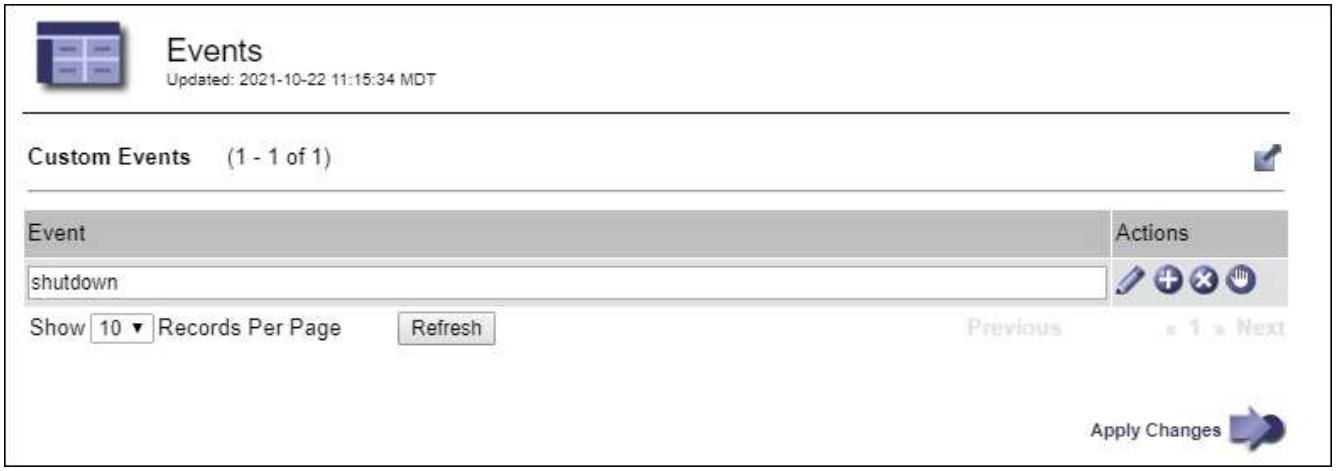

4. Select **Apply Changes**.
5. Select **SUPPORT** > **Tools** > **Grid topology**.
6. Select ***grid node*** > **SSM** > **Events**.
7. Locate the entry for Custom Events in the Events table, and monitor the value for **Count**.

   If the count increases, a custom event you are monitoring is being triggered on that grid node.

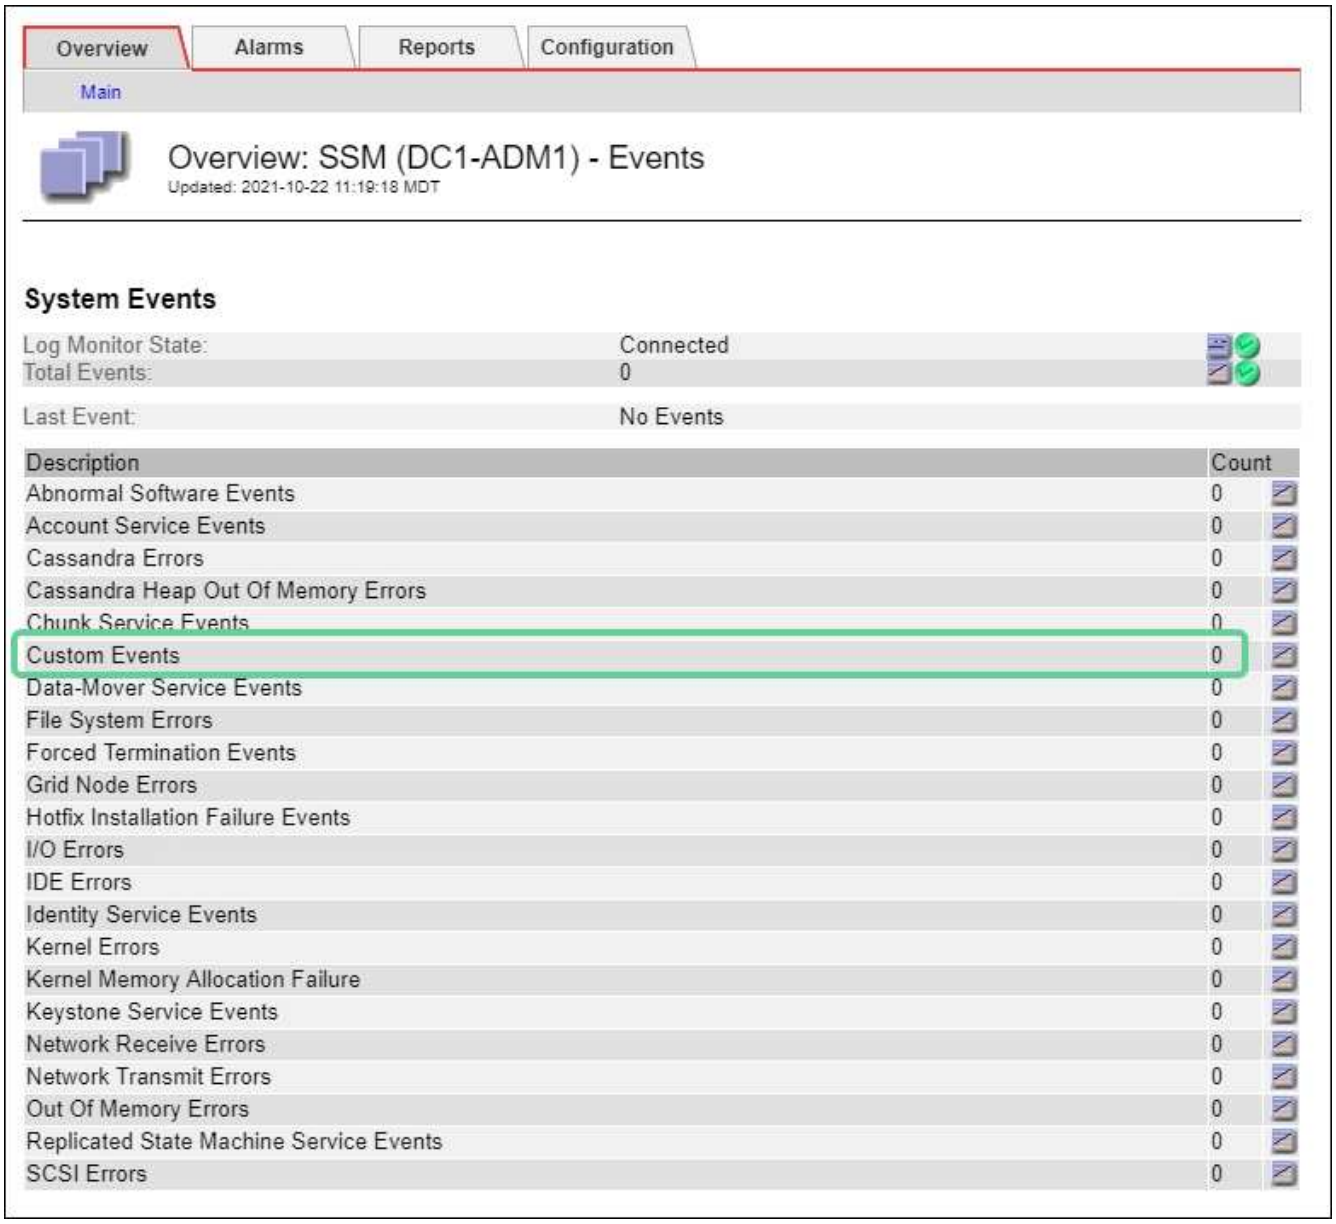

Reset the count of custom events to zero

If you want to reset the counter only for custom events, you must use the Grid Topology page in the Support menu.

Resetting a counter causes the alarm to be triggered by the next event. In contrast, when you acknowledge an alarm, that alarm is only re-triggered if the next threshold level is reached.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select ***grid node*** > **SSM** > **Events** > **Configuration** > **Main**.
3. Select the **Reset** checkbox for Custom Events.

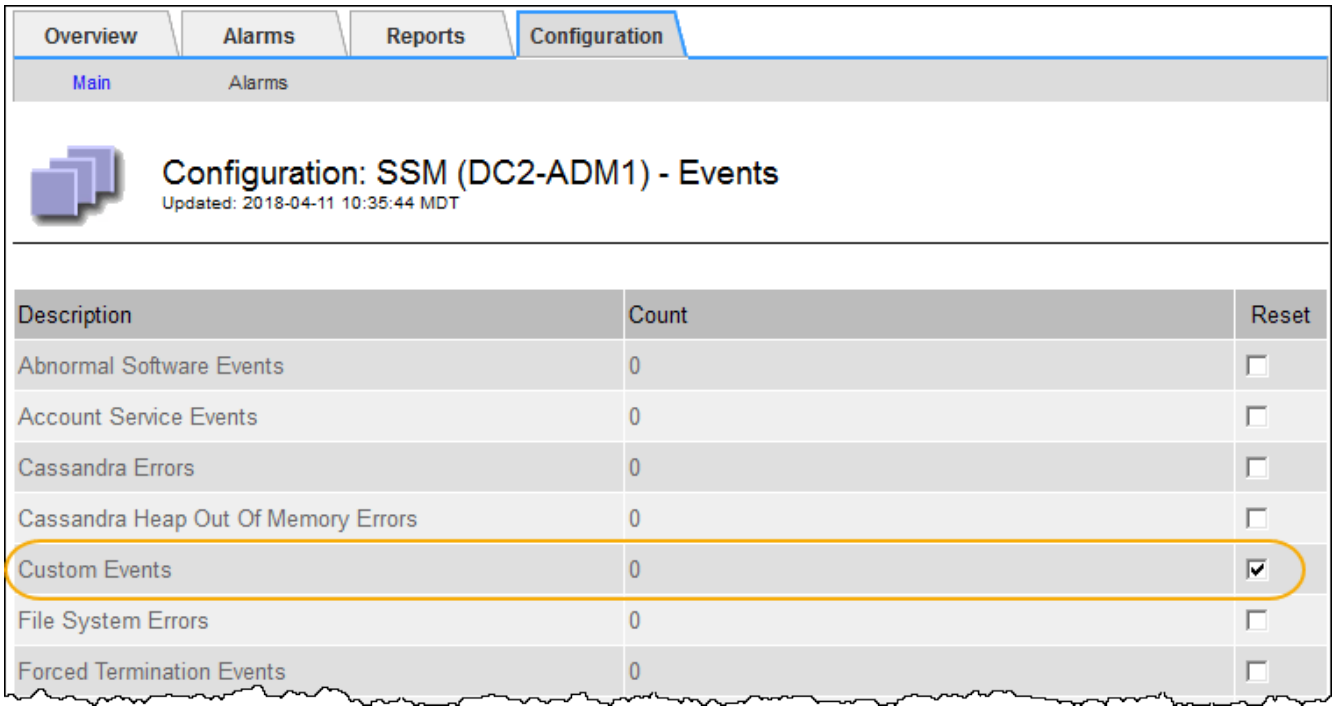

4. Select **Apply Changes**.

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. Each node in the grid also stores a copy of the audit information generated on the node.

For easy access to audit logs, you can configure audit client access for NFS. You can also access audit log files directly from the command line of the Admin Node.

Optionally, you can change the destination of audit logs and send audit information to an external syslog server. Local logs of audit records continue to be generated and stored when an external syslog server is configured. See Configure audit messages and log destinations.

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 Review audit logs.

Collect log files and system data

You can use the Grid Manager to retrieve log files and system data (including configuration data) for your StorageGRID system.

Before you begin

- You must be signed in to the Grid Manager on the primary Admin Node using a supported web browser.
- You have specific access permissions.
- You must have the provisioning passphrase.

About this task

You can use the Grid Manager to gather log files, system data, and configuration data from any grid node for the time period that you select. Data is collected and archived in a .tar.gz file that you can then download to your local computer.

Optionally, you can change the destination of audit logs and send audit information to an external syslog server. Local logs of audit records continue to be generated and stored when an external syslog server is configured. See Configure audit messages and log destinations.

Steps

1. Select **SUPPORT** > **Tools** > **Logs**.

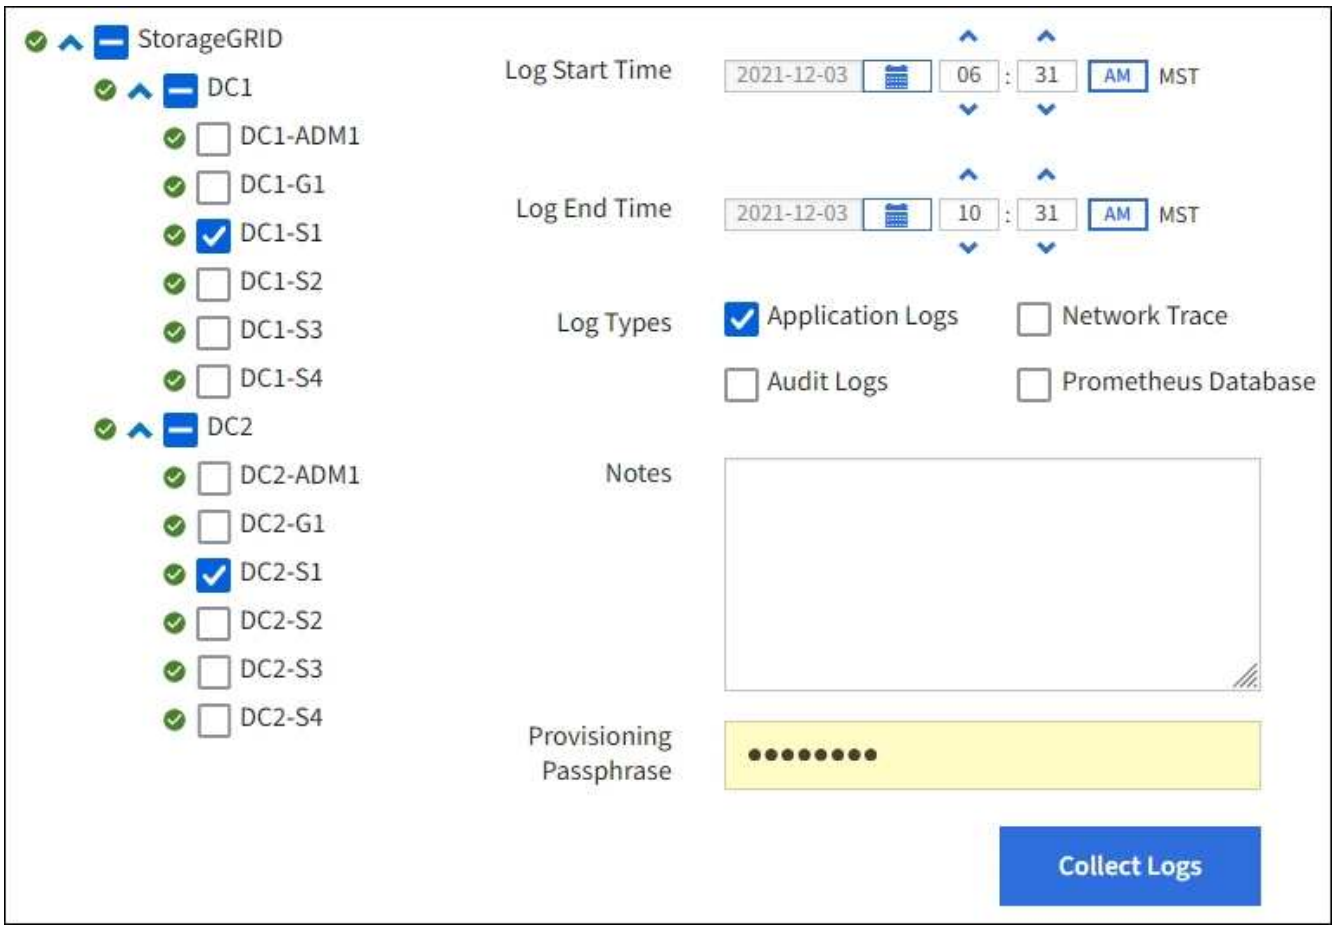

2. Select the grid nodes for which you want to collect log files.

   As required, you can collect log files for the entire grid or an entire data center site.

3. Select a **Start Time** and **End Time** to set the time range of the data to be included in the log files.

   If you select a very long time period or collect logs from all nodes in a large grid, the log archive could become too large to be stored on a node, or too large to be collected to the primary Admin Node for download. If this occurs, you must restart log collection with a smaller set of data.

4. Select the types of logs you want to collect.
   - **Application Logs**: Application-specific logs that technical support uses most frequently for troubleshooting. The logs collected are a subset of the available application logs.
   - **Audit Logs**: Logs containing the audit messages generated during normal system operation.
   - **Network Trace**: Logs used for network debugging.
   - **Prometheus Database**: Time series metrics from the services on all nodes.
5. Optionally, enter notes about the log files you are gathering in the **Notes** text box.

   You can use these notes to give technical support information about the problem that prompted you to collect the log files. Your notes are added to a file called `info.txt`, along with other information about the log file collection. The `info.txt` file is saved in the log file archive package.

6. Enter the provisioning passphrase for your StorageGRID system in the **Provisioning Passphrase** text box.
7. Select **Collect Logs**.

   When you submit a new request, the previous collection of log files is deleted.

   You can use the Logs page to monitor the progress of log file collection for each grid node.

   If you receive an error message about log size, try collecting logs for a shorter time period or for fewer nodes.

8. Select **Download** when log file collection is complete.

   The *.tar.gz* file contains all log files from all grid nodes where log collection was successful. Inside the combined *.tar.gz* file, there is one log file archive for each grid node.

After you finish

You can re-download the log file archive package later if you need to.

Optionally, you can select **Delete** to remove the log file archive package and free up disk space. The current log file archive package is automatically removed the next time you collect log files.

Manually trigger an AutoSupport package

To assist technical support in troubleshooting issues with your StorageGRID system, you can manually trigger an AutoSupport package to be sent.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You must have the Root access or Other grid configuration permission.

Steps

1. Select **SUPPORT** > **Tools** > **AutoSupport**.
2. On the **Actions** tab, select **Send User-Triggered AutoSupport**.

   StorageGRID attempts to send an AutoSupport package to the NetApp Support Site. If the attempt is successful, the **Most Recent Result** and **Last Successful Time** values on the **Results** tab are updated. If there is a problem, the **Most Recent Result** value updates to "Failed," and StorageGRID does not try to send the AutoSupport package again.

   After sending an User-triggered AutoSupport package, refresh the AutoSupport page in your browser after 1 minute to access the most recent results.

View the Grid Topology tree

The Grid Topology tree provides access to detailed information about StorageGRID system elements, including sites, grid nodes, services, and components. In most cases, you only need to access the Grid Topology tree when instructed in the documentation or when working with technical support.

To access the Grid Topology tree, select **SUPPORT** > **Tools** > **Grid topology**.

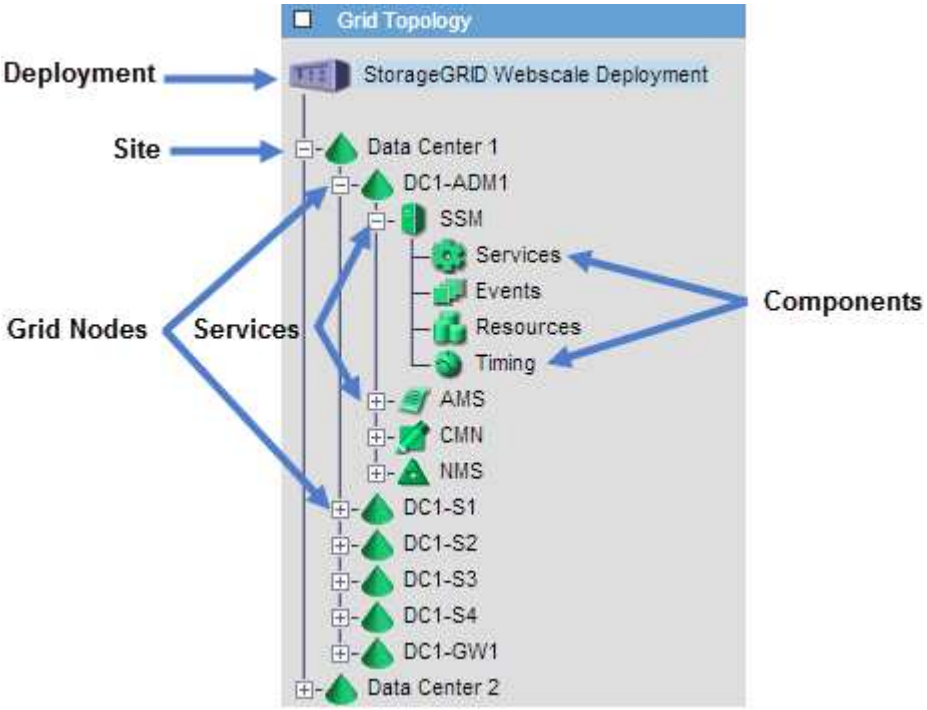

To expand or collapse the Grid Topology tree, click ⊞ or ⊟ at the site, node, or service level. To expand or collapse all items in the entire site or in each node, hold down the **<Ctrl>** key and click.

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 several 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.

Review support metrics

When troubleshooting an issue, you can work with technical support to review detailed metrics and charts for your StorageGRID system.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

The Metrics page allows you to access 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. See the list of commonly used Prometheus metrics.

Steps

1. As directed by technical support, select **SUPPORT** > **Tools** > **Metrics**.

   An example of the Metrics page is shown here:

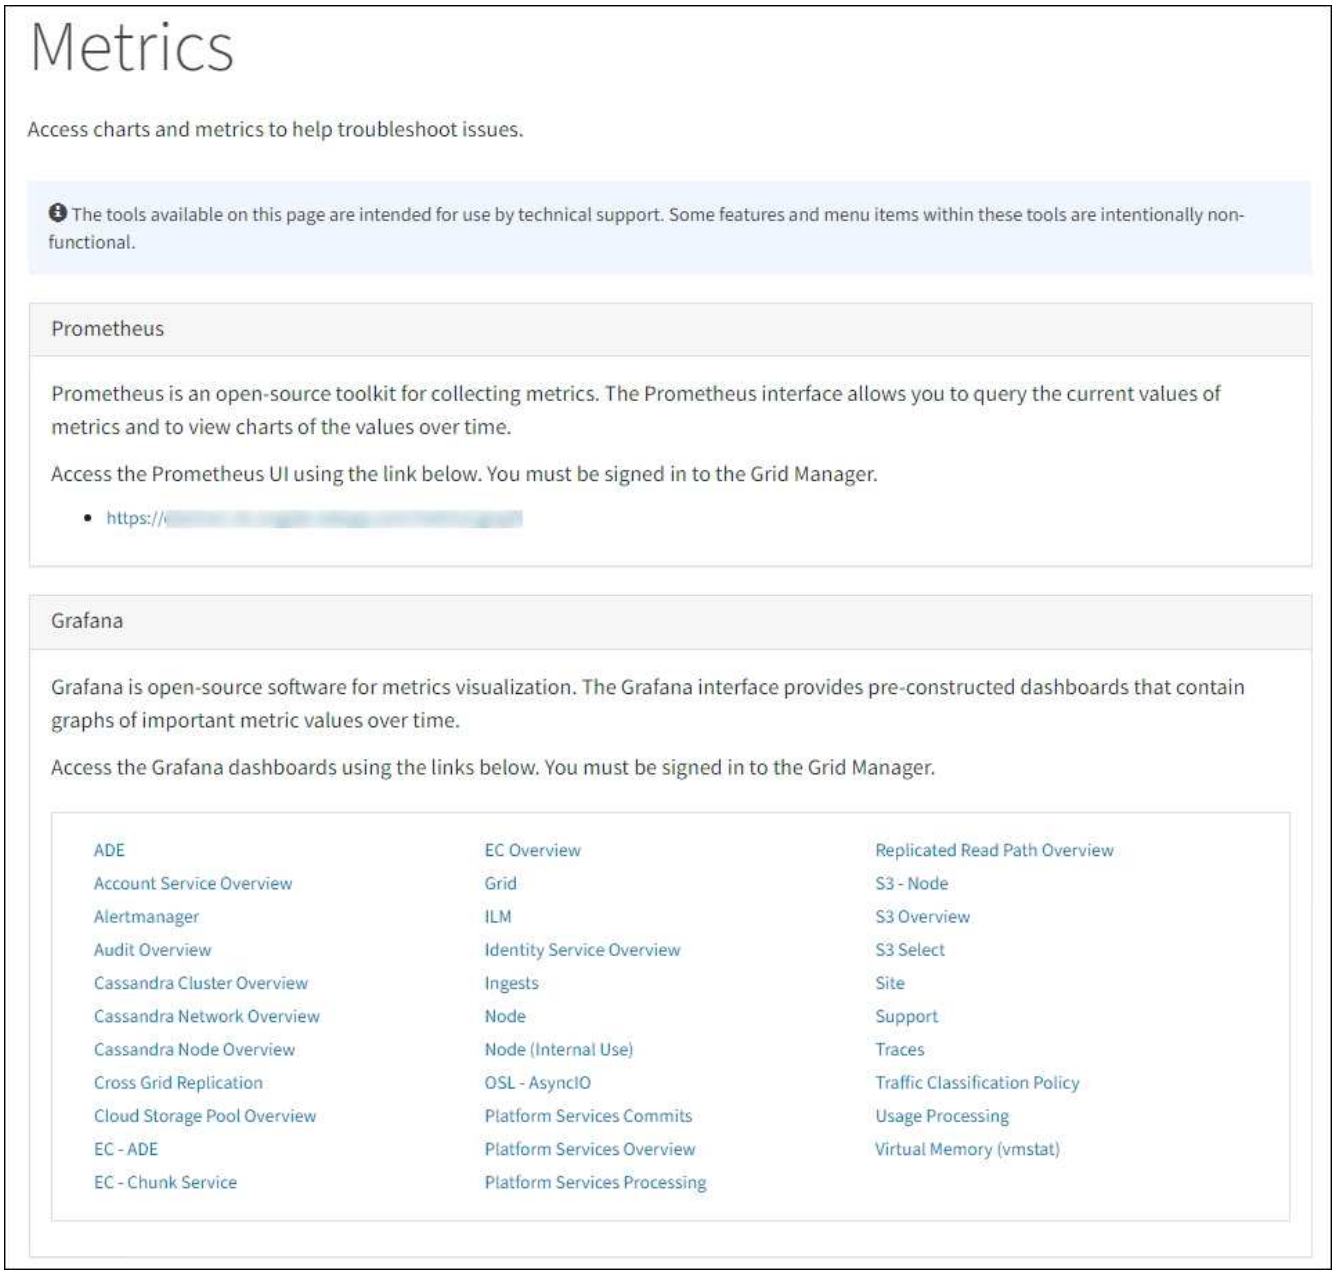

2. To query the current values of StorageGRID metrics and to view graphs of the values over time, click the link in the Prometheus section.

   The Prometheus interface appears. You can use this interface to execute queries on the available StorageGRID metrics and to graph StorageGRID metrics 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.

3. To access pre-constructed dashboards containing graphs of StorageGRID metrics over time, click the links in the Grafana section.

   The Grafana interface for the link you selected appears.

Run diagnostics

When troubleshooting an issue, you can work with technical support to run diagnostics on your StorageGRID system and review the results.

- Review support metrics
- Commonly used Prometheus metrics

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

The Diagnostics page performs a set of diagnostic checks on the current state of the grid. Each 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.

Steps

1. Select **SUPPORT** > **Tools** > **Diagnostics**.

   The Diagnostics page appears and lists the results for each diagnostic check. The results are sorted by severity (Caution, Attention, and then Normal). Within each severity, the results are sorted alphabetically.

   In this example, all diagnostics have a Normal status.

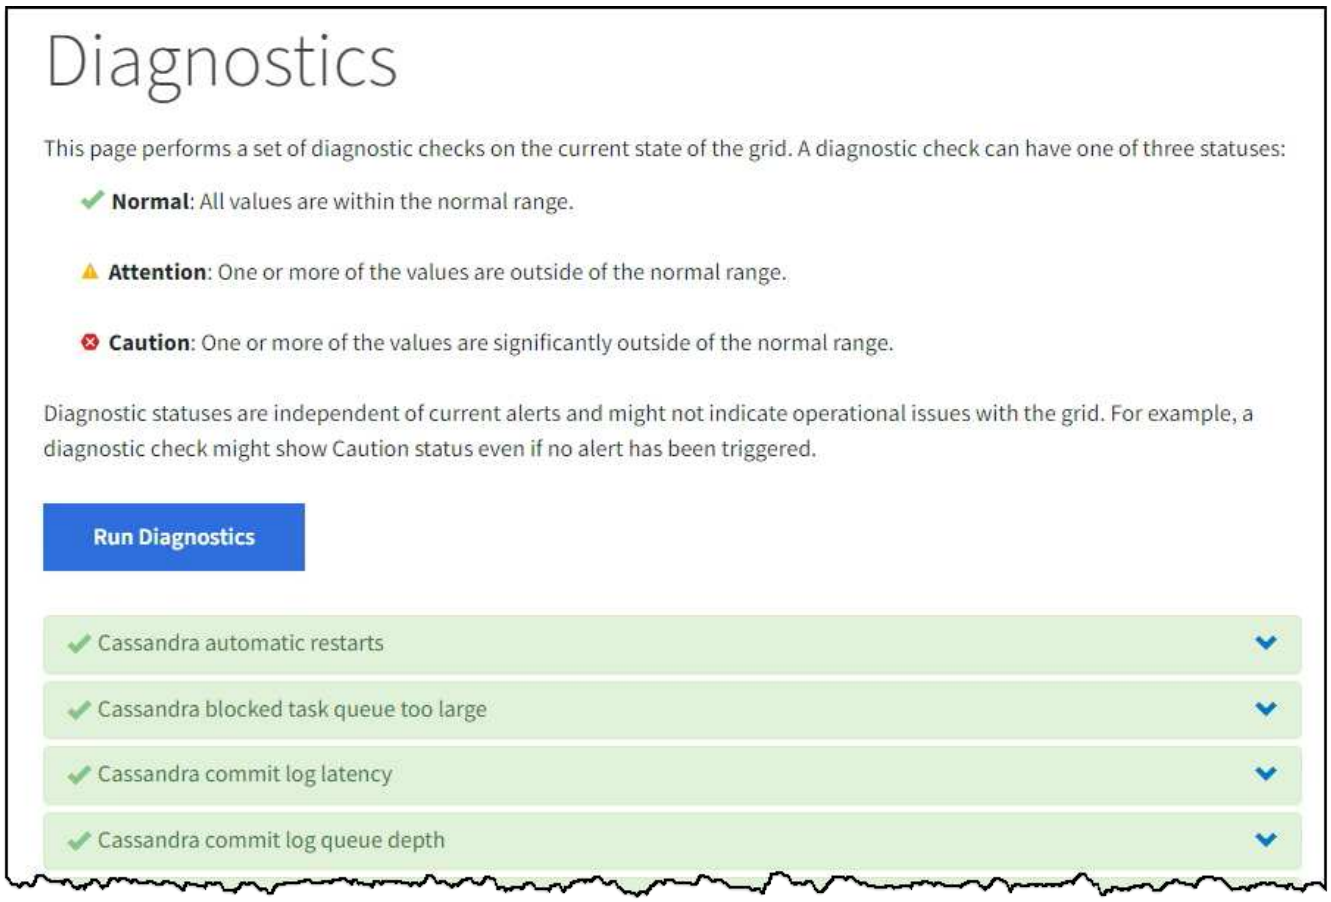

2. To learn more about a specific diagnostic, click anywhere in the row.

   Details about the diagnostic and its current results appear. The following details are listed:

   - **Status**: The current status of this diagnostic: Normal, Attention, or Caution.
   - **Prometheus query**: If used for the diagnostic, the Prometheus expression that was used to generate the status values. (A Prometheus expression is not used for all diagnostics.)
   - **Thresholds**: If available for the diagnostic, the system-defined thresholds for each abnormal diagnostic status. (Threshold values aren't used for all diagnostics.)

> You can't change these thresholds.

- **Status values**: A table showing the status and the value of the diagnostic throughout the StorageGRID system.
  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%

3. **Optional**: To see Grafana charts related to this diagnostic, click the **Grafana dashboard** link.

   This link is not displayed for all diagnostics.

   The related Grafana dashboard appears. In this example, the Node dashboard appears showing CPU Utilization over time for this node as well as other Grafana charts for the node.

> You can also access the pre-constructed Grafana dashboards from the Grafana section of the **SUPPORT** > **Tools** > **Metrics** page.

4. **Optional**: To see a chart of the Prometheus expression over time, click **View in Prometheus**.

   A Prometheus graph of the expression used in the diagnostic appears.

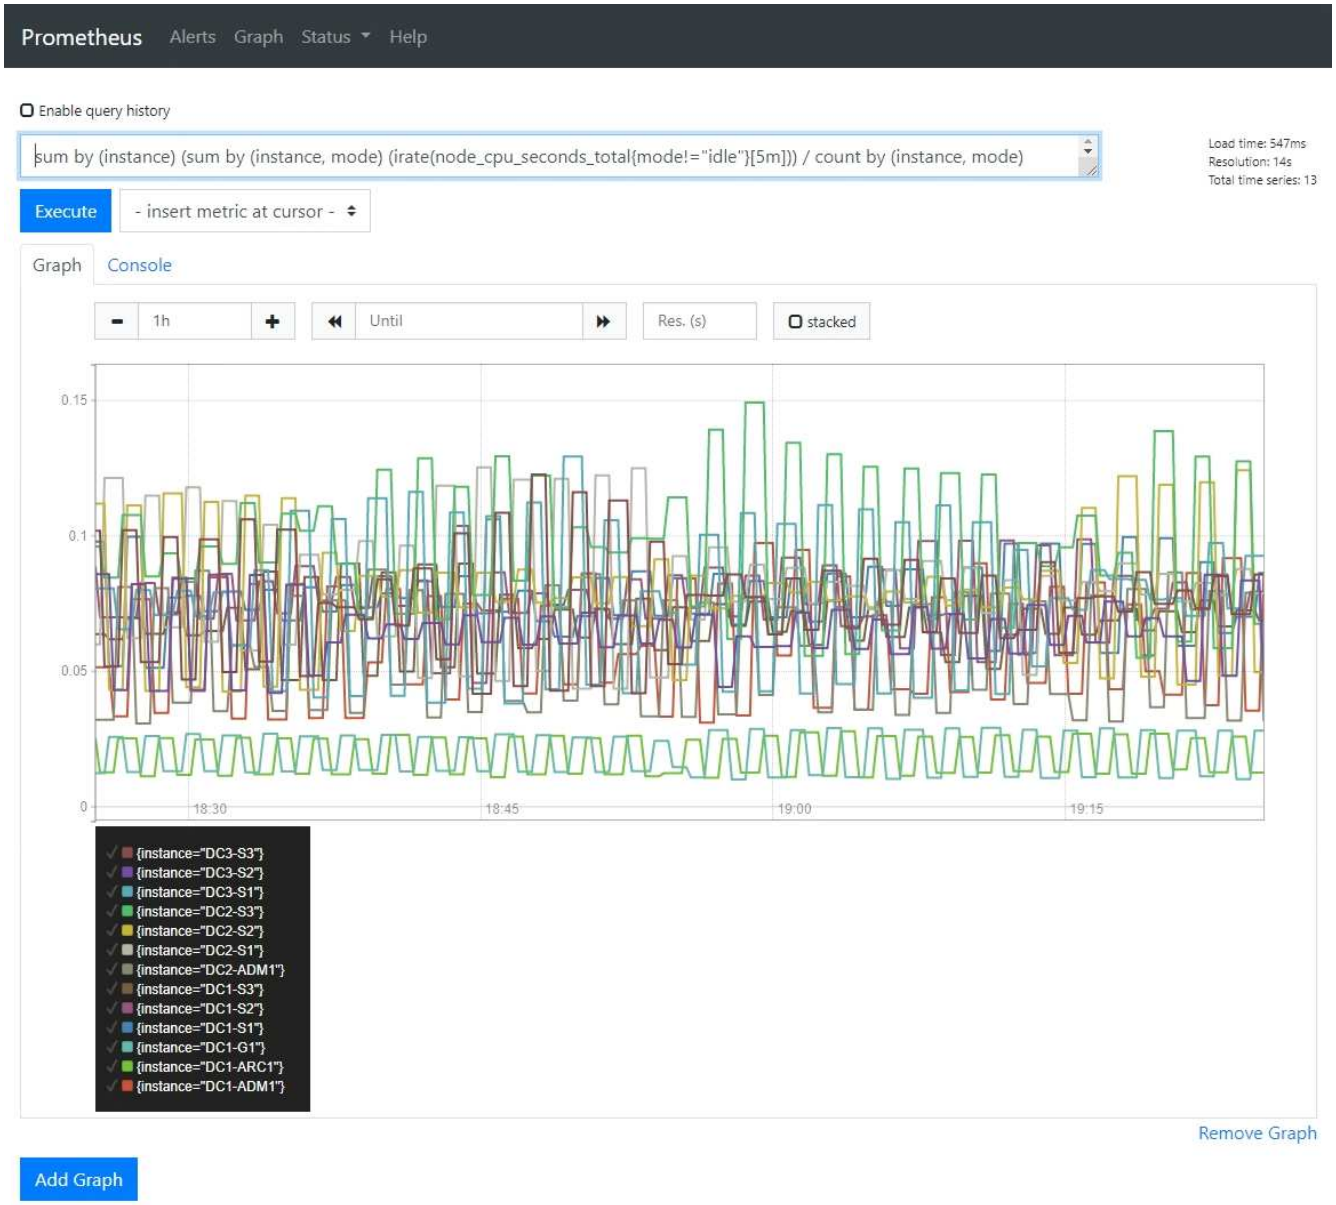

Create custom monitoring applications

You can build custom monitoring applications and dashboards using the StorageGRID metrics available from the Grid Management API.

If you want to monitor metrics that aren't displayed on an existing page of the Grid Manager, or if you want to create custom dashboards for StorageGRID, you can use the Grid Management API to query StorageGRID metrics.

You can also access Prometheus metrics directly with an external monitoring tool, such as Grafana. Using an external tool requires that you upload or generate an administrative client certificate to allow StorageGRID to authenticate the tool for security. See the instructions for administering StorageGRID.

To view the metrics API operations, including the complete list of the metrics that are available, go to the Grid Manager. From the top of the page, select the help icon and select **API documentation** > **metrics**.

metrics Operations on metrics

GET	/grid/metric-labels/{label}/values	Lists the values for a metric label
GET	/grid/metric-names	Lists all available metric names
GET	/grid/metric-query	Performs an instant metric query at a single point in time
GET	/grid/metric-query-range	Performs a metric query over a range of time

The details of how to implement a custom monitoring application are beyond the scope of this documentation.

Troubleshoot StorageGRID system

Troubleshoot a StorageGRID system: Overview

If you encounter a problem when using a StorageGRID system, refer to the tips and guidelines in this section for help in determining and resolving the issue.

Often, you can resolve problems on your own; however, you might need to escalate some issues to technical support.

Define the problem

The first step to solving a problem is to define the problem clearly.

This table provides examples of the types of information that you might collect to define a problem:

Question	Example response
What is the StorageGRID system doing or not doing? What are its symptoms?	Client applications are reporting that objects can't be ingested into StorageGRID.
When did the problem start?	Object ingest was first denied at about 14:50 on January 8, 2020.
How did you first notice the problem?	Notified by client application. Also received alert email notifications.
Does the problem happen consistently, or only sometimes?	Problem is ongoing.
If the problem happens regularly, what steps cause it to occur	Problem happens every time a client tries to ingest an object.

Question	Example response
If the problem happens intermittently, when does it occur? Record the times of each incident that you are aware of.	Problem is not intermittent.
Have you seen this problem before? How often have you had this problem in the past?	This is the first time I have seen this issue.

Assess the risk and impact on the system

After you have defined the problem, assess its risk and impact on the StorageGRID system. For example, the presence of critical alerts does not necessarily mean that the system is not delivering core services.

This table summarizes the impact the example problem is having on system operations:

Question	Example response
Can the StorageGRID system ingest content?	No.
Can client applications retrieve content?	Some objects can be retrieved and others can't.
Is data at risk?	No.
Is the ability to conduct business severely affected?	Yes, because client applications can't store objects to the StorageGRID system and data can't be retrieved consistently.

Collect data

After you have defined the problem and have assessed its risk and impact, collect data for analysis. The type of data that is most useful to collect depends upon the nature of the problem.

Type of data to collect	Why collect this data	Instructions
Create timeline of recent changes	Changes to your StorageGRID system, its configuration, or its environment can cause new behavior.	• Create a timeline of recent changes
Review alerts and alarms	Alerts and alarms can help you quickly determine the root cause of a problem by providing important clues as to the underlying issues that might be causing it. Review the list of current alerts and alarms to see if StorageGRID has identified the root cause of a problem for you. Review alerts and alarms triggered in the past for additional insights.	• View current and resolved alerts • Manage alarms (legacy system)

Type of data to collect	Why collect this data	Instructions
Monitor events	Events include any system error or fault events for a node, including errors such as network errors. Monitor events to learn more about issues or to help with troubleshooting.	• Monitor events
Identify trends using charts and text reports	Trends can provide valuable clues about when issues first appeared, and can help you understand how quickly things are changing.	• Use charts and graphs • Use text reports
Establish baselines	Collect information about the normal levels of various operational values. These baseline values, and deviations from these baselines, can provide valuable clues.	• Establish baselines
Perform ingest and retrieval tests	To troubleshoot performance issues with ingest and retrieval, use a workstation to store and retrieve objects. Compare results against those seen when using the client application.	• Monitor PUT and GET performance
Review audit messages	Review audit messages to follow StorageGRID operations in detail. The details in audit messages can be useful for troubleshooting many types of issues, including performance issues.	• Review audit messages
Check object locations and storage integrity	If you are having storage problems, verify that objects are being placed where you expect. Check the integrity of object data on a Storage Node.	• Monitor object verification operations • Confirm object data locations • Verify object integrity
Collect data for technical support	Technical support might ask you to collect data or review specific information to help troubleshoot issues.	• Collect log files and system data • Manually trigger an AutoSupport package • Review support metrics

Create a timeline of recent changes

When a problem occurs, you should consider what has changed recently and when those changes occurred.

- Changes to your StorageGRID system, its configuration, or its environment can cause new behavior.
- A timeline of changes can help you identify which changes might be responsible for an issue, and how each change might have affected its development.

Create a table of recent changes to your system that includes information about when each change occurred

and any relevant details about the change, such information about what else was happening while the change was in progress:

Time of change	Type of change	Details
For example: • When did you start the node recovery? • When did the software upgrade complete? • Did you interrupt the process?	What happened? What did you do?	Document any relevant details about the change. For example: • Details of the network changes. • Which hotfix was installed. • How client workloads changed. Make sure to note if more than one change was happening at the same time. For example, was this change made while an upgrade was in progress?

Examples of significant recent changes

Here are some examples of potentially significant changes:

- Was the StorageGRID system recently installed, expanded, or recovered?
- Has the system been upgraded recently? Was a hotfix applied?
- Has any hardware been repaired or changed recently?
- Has the ILM policy been updated?
- Has the client workload changed?
- Has the client application or its behavior changed?
- Have you changed load balancers, or added or removed a high availability group of Admin Nodes or Gateway Nodes?
- Have any tasks been started that might take a long time to complete? Examples include:
  - Recovery of a failed Storage Node
  - Storage Node decommissioning
- Have any changes been made to user authentication, such as adding a tenant or changing LDAP configuration?
- Is data migration taking place?
- Were platform services recently enabled or changed?
- Was compliance enabled recently?
- Have Cloud Storage Pools been added or removed?
- Have any changes been made to storage compression or encryption?
- Have there been any changes to the network infrastructure? For example, VLANs, routers, or DNS.
- Have any changes been made to NTP sources?
- Have any changes been made to the Grid, Admin, or Client Network interfaces?
- Have any configuration changes been made to the Archive Node?
- Have any other changes been made to the StorageGRID system or its environment?

Establish baselines

You can establish baselines for your system by recording the normal levels of various operational values. In the future, you can compare current values to these baselines to help detect and resolve abnormal values.

Property	Value	How to obtain
Average storage consumption	GB consumed/day Percent consumed/day	Go to the Grid Manager. On the Nodes page, select the entire grid or a site and go to the Storage tab. On the Storage Used - Object Data chart, find a period where the line is fairly stable. Position your cursor over the chart to estimate how much storage is consumed each day You can collect this information for the entire system or for a specific data center.
Average metadata consumption	GB consumed/day Percent consumed/day	Go to the Grid Manager. On the Nodes page, select the entire grid or a site and go to the Storage tab. On the Storage Used - Object Metadata chart, find a period where the line is fairly stable. Position your cursor over the chart to estimate how much metadata storage is consumed each day You can collect this information for the entire system or for a specific data center.
Rate of S3/Swift operations	Operations/second	On the Grid Manager dashboard, select **Performance** > **S3 operations** or **Performance** > **Swift operations**. To see ingest and retrieval rates and counts for a specific site or node, select **NODES** > ***site or Storage Node*** > **Objects**. Position your cursor over the Ingest and Retrieve chart for S3 or Swift.
Failed S3/Swift operations	Operations	Select **SUPPORT** > **Tools** > **Grid topology**. On the Overview tab in the API Operations section, view the value for S3 Operations - Failed or Swift Operations - Failed.
ILM evaluation rate	Objects/second	From the Nodes page, select ***grid*** > **ILM**. On the ILM Queue chart, find a period where the line is fairly stable. Position your cursor over the chart to estimate a baseline value for **Evaluation rate** for your system.

Property	Value	How to obtain
ILM scan rate	Objects/second	Select **NODES** > ***grid*** > **ILM**. On the ILM Queue chart, find a period where the line is fairly stable. Position your cursor over the chart to estimate a baseline value for **Scan rate** for your system.
Objects queued from client operations	Objects/second	Select **NODES** > ***grid*** > **ILM**. On the ILM Queue chart, find a period where the line is fairly stable. Position your cursor over the chart to estimate a baseline value for **Objects queued (from client operations)** for your system.
Average query latency	Milliseconds	Select **NODES** > ***Storage Node*** > **Objects**. In the Queries table, view the value for Average Latency.

Analyze data

Use the information that you collect to determine the cause of the problem and potential solutions.

The analysis is problem-dependent, but in general:

- Locate points of failure and bottlenecks using the alarms.
- Reconstruct the problem history using the alarm history and charts.
- Use charts to find anomalies and compare the problem situation with normal operation.

Escalation information checklist

If you can't resolve the problem on your own, contact technical support. Before contacting technical support, gather the information listed in the following table to facilitate problem resolution.

✔	Item	Notes
	Problem statement	What are the problem symptoms? When did the problem start? Does it happen consistently or intermittently? If intermittently, what times has it occurred? Define the problem
	Impact assessment	What is the severity of the problem? What is the impact to the client application? • Has the client connected successfully before? • Can the client ingest, retrieve, and delete data?

✓	Item	Notes
	StorageGRID System ID	Select **MAINTENANCE** > **System** > **License**. The StorageGRID System ID is shown as part of the current license.
	Software version	From the top of the Grid Manager, select the help icon and select **About** to see the StorageGRID version.
	Customization	Summarize how your StorageGRID system is configured. For example, list the following: • Does the grid use storage compression, storage encryption, or compliance? • Does ILM make replicated or erasure-coded objects? Does ILM ensure site redundancy? Do ILM rules use the Balanced, Strict, or Dual Commit ingest behaviors?
	Log files and system data	Collect log files and system data for your system. Select **SUPPORT** > **Tools** > **Logs**. You can collect logs for the entire grid, or for selected nodes. If you are collecting logs only for selected nodes, be sure to include at least one Storage Node that has the ADC service. (The first three Storage Nodes at a site include the ADC service.) Collect log files and system data
	Baseline information	Collect baseline information regarding ingest operations, retrieval operations, and storage consumption. Establish baselines
	Timeline of recent changes	Create a timeline that summarizes any recent changes to the system or its environment. Create a timeline of recent changes
	History of efforts to diagnose the issue	If you have taken steps to diagnose or troubleshoot the issue yourself, make sure to record the steps you took and the outcome.

Troubleshoot object and storage issues

Confirm object data locations

Depending on the problem, you might want to confirm where object data is being stored. For example, you might want to verify that the ILM policy is performing as expected and

object data is being stored where intended.

Before you begin

- You must have an object identifier, which can be one of:
  - **UUID**: The object's Universally Unique Identifier. Enter the UUID in all uppercase.
  - **CBID**: The object's unique identifier within StorageGRID . You can obtain an object's CBID from the audit log. Enter the CBID in all uppercase.
  - **S3 bucket and object key**: When an object is ingested through the S3 interface, the client application uses a bucket and object key combination to store and identify the object.
  - **Swift container and object name**: When an object is ingested through the Swift interface, the client application uses a container and object name combination to store and identify the object.

Steps

1. Select **ILM** > **Object metadata lookup**.
2. Type the object's identifier in the **Identifier** field.

   You can enter a UUID, CBID, S3 bucket/object-key, or Swift container/object-name.

3. If you want to look up a specific version of the object, enter the version ID (optional).

4. Select **Look Up**.

   The object metadata lookup results appear. This page lists the following types of information:

   - System metadata, including the object ID (UUID), the version ID (optional), the object name, the name of the container, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
   - Any custom user metadata key-value pairs associated with the object.
   - For S3 objects, any object tag key-value pairs associated with the object.
   - For replicated object copies, the current storage location of each copy.
   - For erasure-coded object copies, the current storage location of each fragment.
   - For object copies in a Cloud Storage Pool, the location of the object, including the name of the external bucket and the object's unique identifier.
   - For segmented objects and multipart objects, a list of object segments including segment identifiers and data sizes. For objects with more than 100 segments, only the first 100 segments are shown.
   - All object metadata in the unprocessed, internal storage format. This raw metadata includes internal

system metadata that is not guaranteed to persist from release to release.

The following example shows the object metadata lookup results for an S3 test object that is stored as two replicated copies.

System Metadata

Object ID	A12E96FF-B13F-4905-9E9E-45373F6E7DA8
Name	testobject
Container	source
Account	t-1582139188
Size	5.24 MB
Creation Time	2020-02-19 12:15:59 PST
Modified Time	2020-02-19 12:15:59 PST

Replicated Copies

Node	Disk Path
99-97	/var/local/rangedb/2/p/06/0B/00nM8H$|TFbnQQ}|CV2E
99-99	/var/local/rangedb/1/p/12/0A/00nM8H$|TFboW28|CXG%

Raw Metadata


```
{
    "TYPE": "CTNT",
    "CHND": "A12E96FF-B13F-4905-9E9E-45373F6E7DA8",
    "NAME": "testobject",
    "CBID": "0x8823DE7EC7C10416",
    "PHND": "FEA0AE51-534A-11EA-9FCD-31FF00C36D56",
    "PPTH": "source",
    "META": {
        "BASE": {
            "PAWS": "2",
```


Object store (storage volume) failures

The underlying storage on a Storage Node is divided into object stores. Object stores are also known as storage volumes.

You can view object store information for each Storage Node. Object stores are shown at the bottom of the **NODES** > ***Storage Node*** > **Storage** page.

Disk devices

Name	World Wide Name	I/O load	Read rate	Write rate
sdc(8:16,sdb)	N/A	0.05%	0 bytes/s	4 KB/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s
sdf(8:64,sde)	N/A	0.00%	0 bytes/s	82 bytes/s
sdg(8:80,sdf)	N/A	0.00%	0 bytes/s	82 bytes/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s
croot(8:1,sda1)	N/A	0.04%	0 bytes/s	4 KB/s
cvloc(8:2,sda2)	N/A	0.95%	0 bytes/s	52 KB/s

Volumes

Mount point	Device	Status	Size	Available	Write cache status
/	croot	Online	21.00 GB	14.73 GB	Unknown
/var/local	cvloc	Online	85.86 GB	80.94 GB	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.17 GB	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/2	sde	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/3	sdf	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/4	sdg	Online	107.32 GB	107.18 GB	Enabled

Object stores

ID	Size	Available	Replicated data	EC data	Object data (%)	Health
0000	107.32 GB	96.44 GB	1.55 MB	0 bytes	0.00%	No Errors
0001	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0002	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0003	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0004	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors

To see more details about each Storage Node, follow these steps:

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select ***site*** > ***Storage Node*** > **LDR** > **Storage** > **Overview** > **Main**.

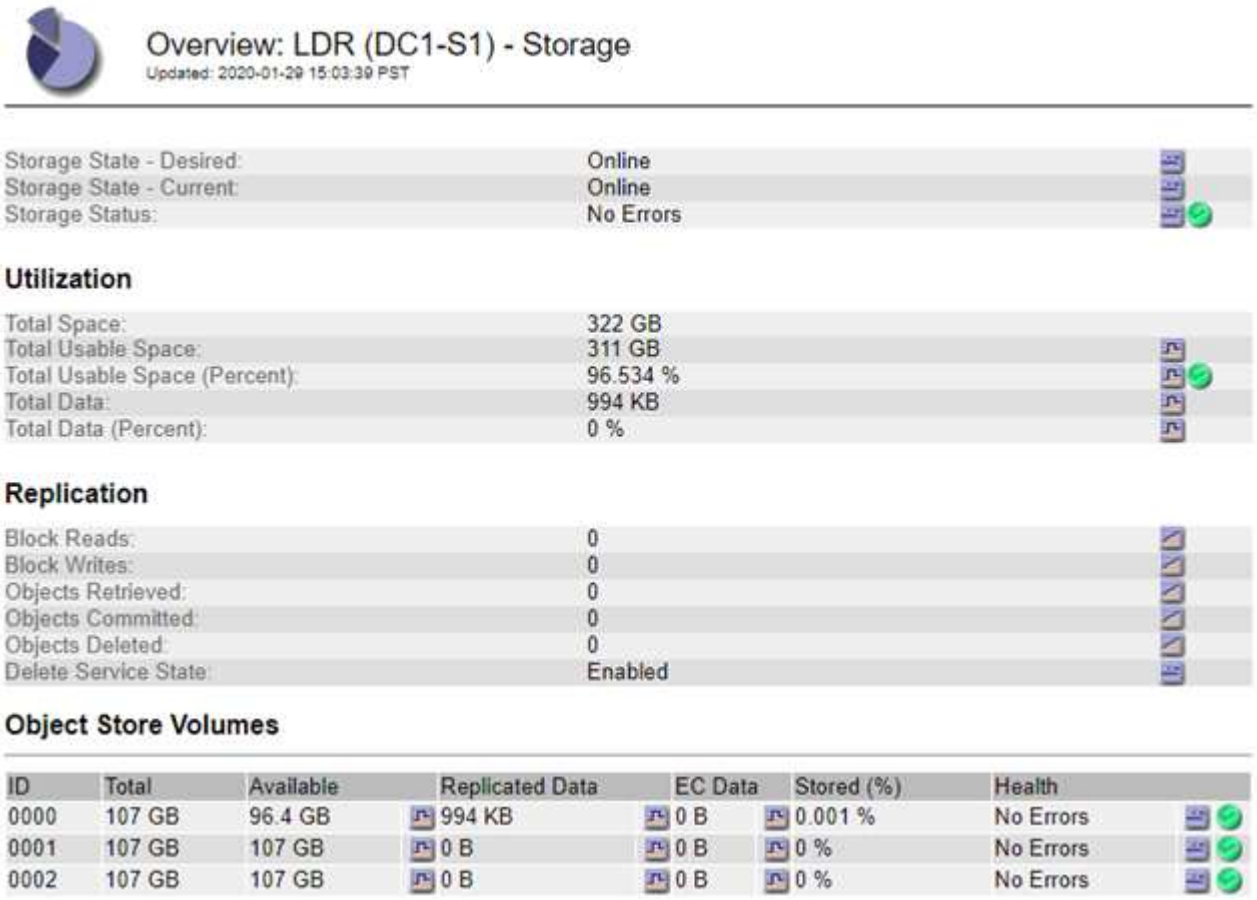

Depending on the nature of the failure, faults with a storage volume might be reflected in an alarm on the storage status or on the health of an object store. If a storage volume fails, you should repair the failed storage volume to restore the Storage Node to full functionality as soon as possible. If necessary, you can go to the **Configuration** tab and place the Storage Node in a read-only state so that the StorageGRID system can use it for data retrieval while you prepare for a full recovery of the server.

Verify object integrity

The StorageGRID system verifies the integrity of object data on Storage Nodes, checking for both corrupt and missing objects.

There are two verification processes: background verification and object existence check (formerly called foreground verification). They work together to ensure data integrity. Background verification runs automatically, and continuously checks the correctness of object data. Object existence check can be triggered by a user to more quickly verify the existence (although not the correctness) of objects.

What is background verification?

The background verification process automatically and continuously checks Storage Nodes for corrupt copies of object data, and automatically attempts to repair any issues that it finds.

Background verification checks the integrity of replicated objects and erasure-coded objects, as follows:

- **Replicated objects**: If the background verification process finds a replicated object that is corrupt, the corrupt copy is removed from its location and quarantined elsewhere on the Storage Node. Then, a new uncorrupted copy is generated and placed to satisfy the active ILM policies. The new copy might not be placed on the Storage Node that was used for the original copy.

  > Corrupt object data is quarantined rather than deleted from the system, so that it can still be accessed. For more information about accessing quarantined object data, contact technical support.

- **Erasure-coded objects**: If the background verification process detects that a fragment of an erasure-coded object is corrupt, StorageGRID automatically attempts to rebuild the missing fragment in place on the same Storage Node, using the remaining data and parity fragments. If the corrupted fragment can't be rebuilt, an attempt is made to retrieve another copy of the object. If retrieval is successful, an ILM evaluation is performed to create a replacement copy of the erasure-coded object.

  The background verification process checks objects on Storage Nodes only. It does not check objects on Archive Nodes or in a Cloud Storage Pool. Objects must be older than four days to qualify for background verification.

Background verification runs at a continuous rate that is designed not to interfere with ordinary system activities. Background verification can't be stopped. However you can increase the background verification rate to more quickly verify the contents of a Storage Node if you suspect a problem.

Alerts and alarms (legacy) related to background verification

If the system detects a corrupt object that it can't correct automatically (because the corruption prevents the object from being identified), the **Unidentified corrupt object detected** alert is triggered.

If background verification can't replace a corrupted object because it can't locate another copy, the **Objects lost** alert is triggered.

Change the background verification rate

You can change the rate at which background verification checks replicated object data on a Storage Node if you have concerns about data integrity.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

You can change the Verification Rate for background verification on a Storage Node:

- Adaptive: Default setting. The task is designed to verify at a maximum of 4 MB/s or 10 objects/s (whichever is exceeded first).
- High: Storage verification proceeds quickly, at a rate that can slow ordinary system activities.

Use the High verification rate only when you suspect that a hardware or software fault might have corrupted object data. After the High priority background verification completes, the Verification Rate automatically resets to Adaptive.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select ***Storage Node*** > **LDR** > **Verification**.
3. Select **Configuration** > **Main**.
4. Go to **LDR** > **Verification** > **Configuration** > **Main**.
5. Under Background Verification, select **Verification Rate** > **High** or **Verification Rate** > **Adaptive**.

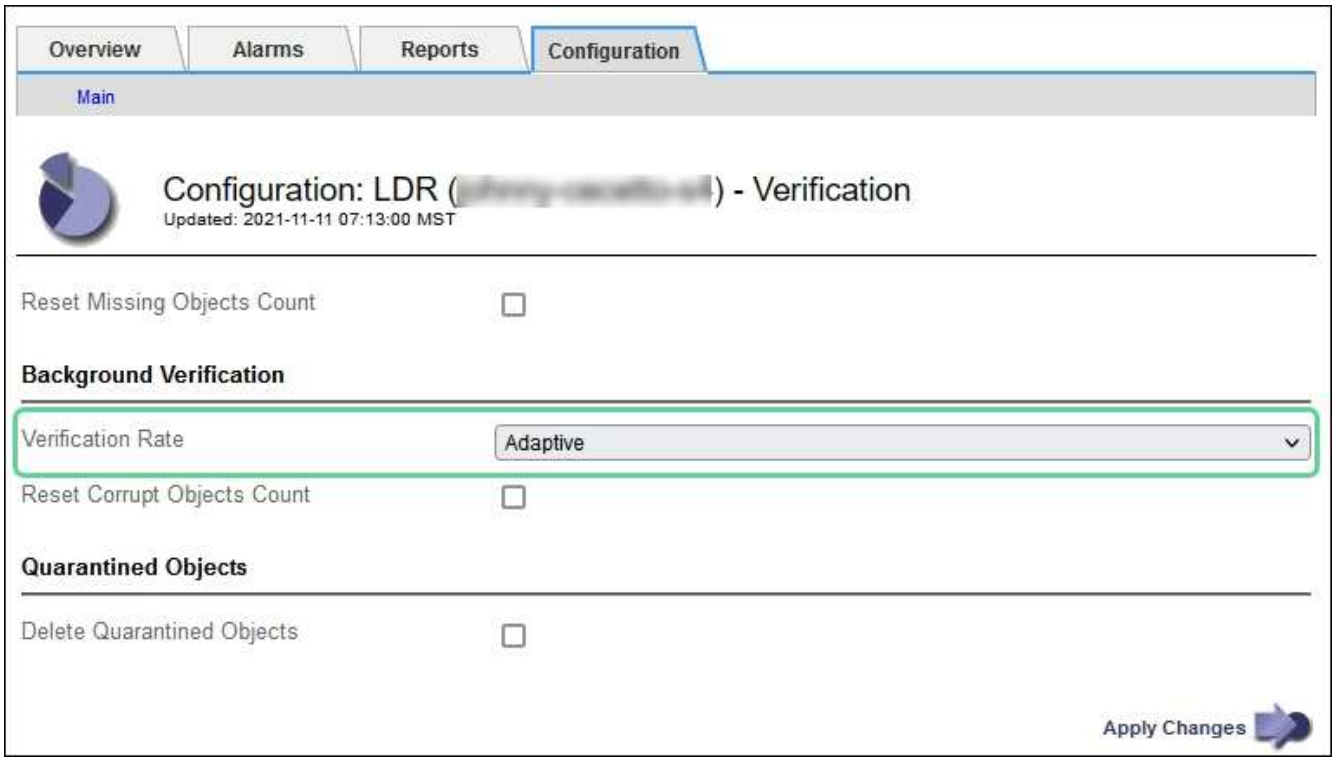

> Setting the Verification Rate to High triggers the VPRI (Verification Rate) legacy alarm at the Notice level.

6. Click **Apply Changes**.
7. Monitor the results of background verification for replicated objects.
   a. Go to **NODES** > ***Storage Node*** > **Objects**.
   b. In the Verification section, monitor the values for **Corrupt Objects** and **Corrupt Objects Unidentified**.

      If background verification finds corrupt replicated object data, the **Corrupt Objects** metric is incremented, and StorageGRID attempts to extract the object identifier from the data, as follows:

      - If the object identifier can be extracted, StorageGRID automatically creates a new copy of the object data. The new copy can be made anywhere in the StorageGRID system that satisfies the active ILM policies.
      - If the object identifier can't be extracted (because it has been corrupted), the **Corrupt Objects Unidentified** metric is incremented, and the **Unidentified corrupt object detected** alert is triggered.

   c. If corrupt replicated object data is found, contact technical support to determine the root cause of the corruption.
8. Monitor the results of background verification for erasure-coded objects.

If background verification finds corrupt fragments of erasure-coded object data, the Corrupt Fragments Detected attribute is incremented. StorageGRID recovers by rebuilding the corrupt fragment in place on the same Storage Node.

   a. Select **SUPPORT** > **Tools** > **Grid topology**.
   b. Select ***Storage Node*** > **LDR** > **Erasure Coding**.
   c. In the Verification Results table, monitor the Corrupt Fragments Detected (ECCD) attribute.

9. After corrupt objects have been automatically restored by the StorageGRID system, reset the count of corrupt objects.
   a. Select **SUPPORT** > **Tools** > **Grid topology**.
   b. Select ***Storage Node*** > **LDR** > **Verification** > **Configuration**.
   c. Select **Reset Corrupt Object Count**.
   d. Click **Apply Changes**.
10. If you are confident that quarantined objects aren't required, you can delete them.

> If the **Objects lost** alert or the LOST (Lost Objects) legacy alarm was triggered, technical support might want to access quarantined objects to help debug the underlying issue or to attempt data recovery.

   a. Select **SUPPORT** > **Tools** > **Grid topology**.
   b. Select ***Storage Node*** > **LDR** > **Verification** > **Configuration**.
   c. Select **Delete Quarantined Objects**.
   d. Select **Apply Changes**.

What is object existence check?

Object existence check verifies whether all expected replicated copies of objects and erasure-coded fragments exist on a Storage Node. Object existence check does not verify the object data itself (background verification does that); instead, it provides a way to verify the integrity of storage devices, especially if a recent hardware issue could have affected data integrity.

Unlike background verification, which occurs automatically, you must manually start an object existence check job.

Object existence check reads the metadata for every object stored in StorageGRID and verifies the existence of both replicated object copies and erasure-coded object fragments. Any missing data is handled as follows:

- **Replicated copies**: If a copy of replicated object data is missing, StorageGRID automatically attempts to replace the copy from a copy stored elsewhere in the system. The Storage Node runs an existing copy through an ILM evaluation, which will determine that the current ILM policy is no longer being met for this object because another copy is missing. A new copy is generated and placed to satisfy the system's active ILM policies. This new copy might not be placed in the same location where the missing copy was stored.
- **Erasure-coded fragments**: If a fragment of an erasure-coded object is missing, StorageGRID automatically attempts to rebuild the missing fragment in place on the same Storage Node using the remaining fragments. If the missing fragment can't be rebuilt (because too many fragments have been lost), ILM attempts to find another copy of the object, which it can use to generate a new erasure-coded fragment.

Run object existence check

You create and run one object existence check job at a time. When you create a job, you select the Storage Nodes and volumes you want to verify. You also select the consistency for the job.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have the Maintenance or Root access permission.
- You have ensured that the Storage Nodes you want to check are online. Select **NODES** to view the table of nodes. Ensure that no alert icons appear next to the node name for the nodes you want to check.
- You have ensured that the following procedures are **not** running on the nodes you want to check:
  - Grid expansion to add a Storage Node
  - Storage Node decommission
  - Recovery of a failed storage volume
  - Recovery of a Storage Node with a failed system drive
  - EC rebalance
  - Appliance node clone

Object existence check does not provide useful information while these procedures are in progress.

About this task

An object existence check job can take days or weeks to complete, depending on the number of objects in the grid, the selected storage nodes and volumes, and the selected consistency. You can run only one job at a time, but you can select multiple Storage Nodes and volumes at the same time.

Steps

1. Select **MAINTENANCE** > **Tasks** > **Object existence check**.
2. Select **Create job**. The Create an object existence check job wizard appears.
3. Select the nodes containing the volumes you want to verify. To select all online nodes, select the **Node name** checkbox in the column header.

   You can search by node name or site.

   You can’t select nodes that aren’t connected to the grid.

4. Select **Continue**.
5. Select one or more volumes for each node in the list. You can search for volumes using the storage volume number or node name.

   To select all volumes for each node you selected, select the **Storage volume** checkbox in the column header.

6. Select **Continue**.
7. Select the consistency for the job.

   The consistency determines how many copies of object metadata are used for the object existence check.

   - **Strong-site**: Two copies of metadata at a single site.

- **Strong-global**: Two copies of metadata at each site.
- **All** (default): All three copies of metadata at each site.

  For more information about consistency, see the descriptions in the wizard.

8. Select **Continue**.
9. Review and verify your selections. You can select **Previous** to go to a previous step in the wizard to update your selections.

   An Object existence check job is generated and runs until one of the following occurs:

   - The job completes.
   - You pause or cancel the job. You can resume a job that you have paused, but you can't resume a job that you have canceled.
   - The job stalls. The **Object existence check has stalled** alert is triggered. Follow the corrective actions specified for the alert.
   - The job fails. The **Object existence check has failed** alert is triggered. Follow the corrective actions specified for the alert.
   - A "Service unavailable" or an "Internal server error" message appears. After one minute, refresh the page to continue monitoring the job.

   > As needed, you can navigate away from the Object existence check page and return to continue monitoring the job.

10. As the job runs, view the **Active job** tab and note the value of Missing object copies detected.

    This value represents the total number of missing copies of replicated objects and erasure-coded objects with one or more missing fragments.

    If the number of Missing object copies detected is greater than 100, there might be an issue with the Storage Node's storage.

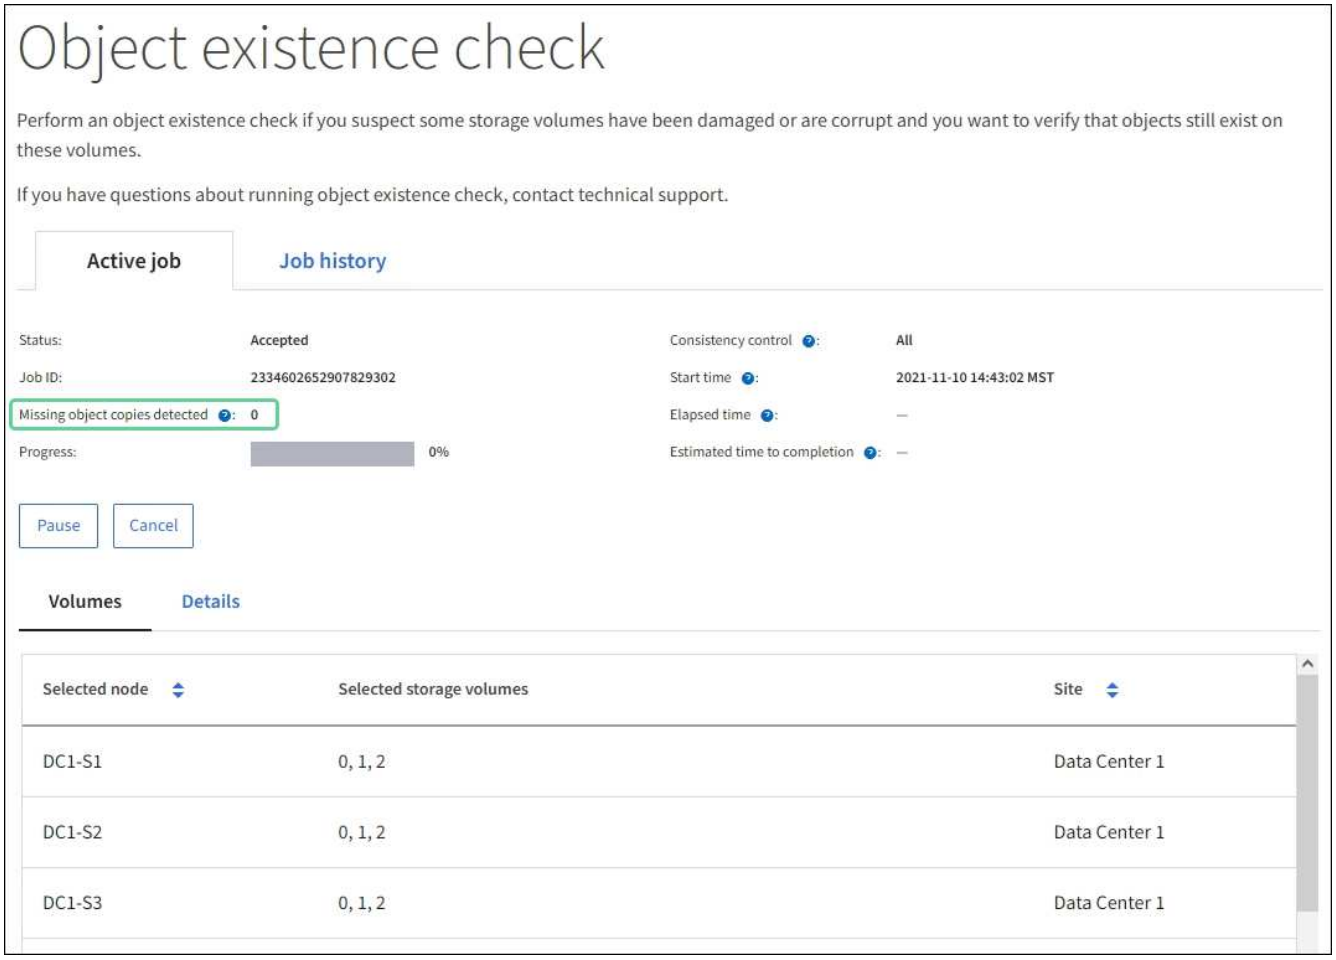

11. After the job has completed, take any additional required actions:
    - If Missing object copies detected is zero, then no issues were found. No action is required.
    - If Missing object copies detected is greater than zero and the **Objects lost** alert has not been triggered, then all missing copies were repaired by the system. Verify that any hardware issues have been corrected to prevent future damage to object copies.
    - If Missing object copies detected is greater than zero and the **Objects lost** alert has been triggered, then data integrity could be affected. Contact technical support.
    - You can investigate lost object copies by using grep to extract the LLST audit messages: `grep LLST audit_file_name`.

      This procedure is similar to the one for investigating lost objects, although for object copies you search for `LLST` instead of `OLST`.

12. If you selected the strong-site or strong-global consistency for the job, wait approximately three weeks for metadata consistency and then rerun the job on the same volumes again.

    When StorageGRID has had time to achieve metadata consistency for the nodes and volumes included in the job, rerunning the job could clear erroneously reported missing object copies or cause additional object copies to be checked if they were missed.

    a. Select **MAINTENANCE** > **Object existence check** > **Job history**.

    b. Determine which jobs are ready to be rerun:

       i. Look at the **End time** column to determine which jobs were run more than three weeks ago.

ii. For those jobs, scan the Consistency control column for strong-site or strong-global.

c. Select the checkbox for each job you want to rerun, then select **Rerun**.

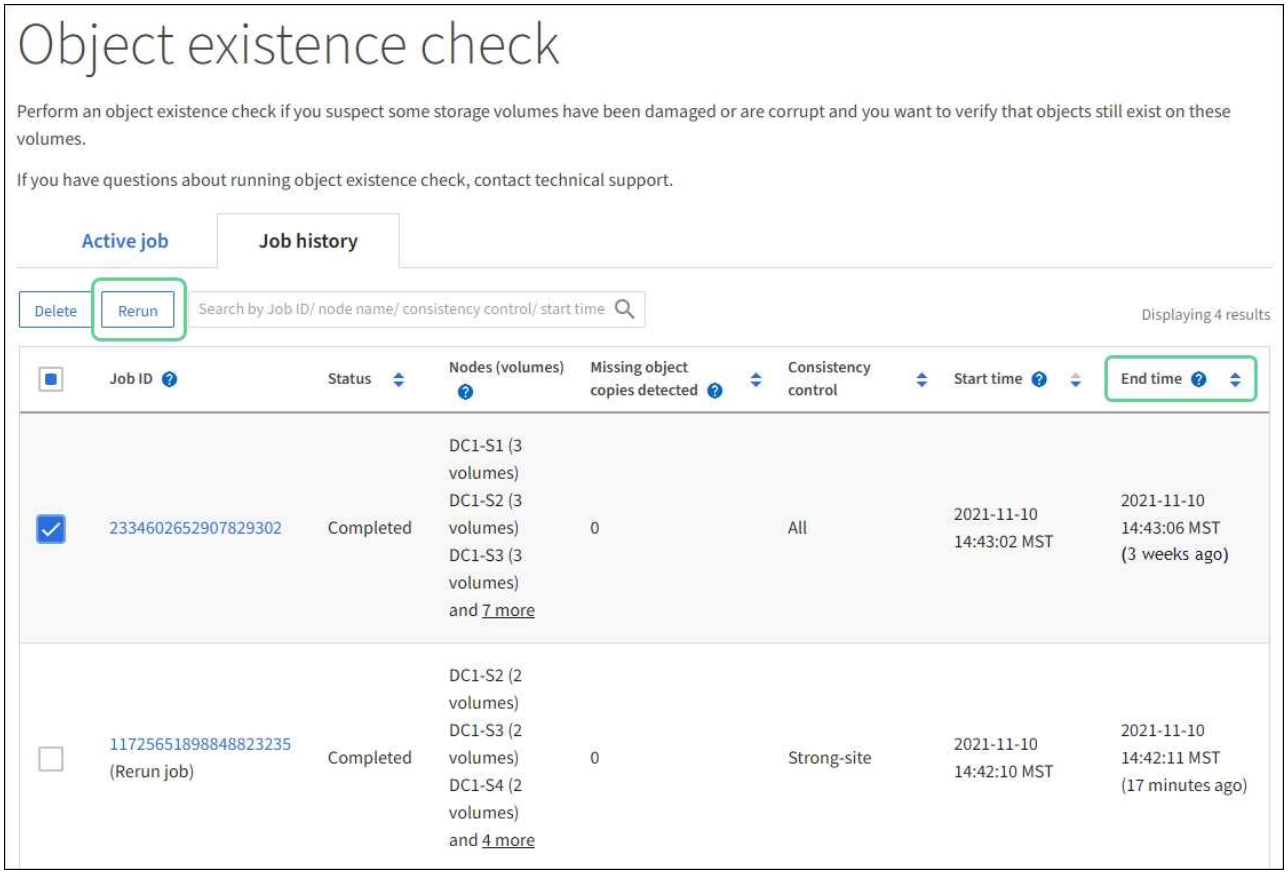

d. In the Rerun jobs wizard, review the selected nodes and volumes and the consistency.

e. When you are ready to rerun the jobs, select **Rerun**.

The Active job tab appears. All the jobs you selected are rerun as one job at a consistency of strong-site. A **Related jobs** field in the Details section lists the job IDs for the original jobs.

After you finish

If you still have concerns about data integrity, go to **SUPPORT** > **Tools** > **Grid topology** > ***site*** > ***Storage Node*** > **LDR** > **Verification** > **Configuration** > **Main** and increase the Background Verification Rate. Background verification checks the correctness of all stored object data and repairs any issues that it finds. Finding and repairing potential issues as quickly as possible reduces the risk of data loss.

Troubleshoot S3 PUT Object size too large alert

The S3 PUT Object size too large alert is triggered if a tenant attempts a non-multipart PutObject operation that exceeds the S3 size limit of 5 GiB.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

Determine which tenants use objects that are larger than 5 GiB, so you can notify them.

Steps

1. Go to **CONFIGURATION** > **Monitoring** > **Audit and syslog server**.
2. If Client Writes are Normal, access the audit log:
   a. Enter `ssh admin@primary_Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

   e. Enter `cd /var/local/log`
   f. Identify which tenants are using objects larger than 5 GiB.
      i. Enter `zgrep SPUT * | egrep "CSIZ\(UI64\):[0-9]*[5-9][0-9]{9}"`
      ii. For each audit message in the results, look at `S3AI` field to determine the tenant account ID. Use the other fields in the message to determine which IP address was used by the client, the bucket, and the object:

Code	Description
`SAIP`	Source IP
`S3AI`	Tenant ID
`S3BK`	Bucket
`S3KY`	Object
`CSIZ`	Size (bytes)

Example audit log results

```
audit.log:2023-01-05T18:47:05.525999
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1672943621106262][TIME(UI64):80
4317333][SAIP(IPAD):"10.96.99.127"][S3AI(CSTR):"933908492661540043
43"][SACC(CSTR):"bhavna"][S3AK(CSTR):"06OX85M40Q90Y280B7YT"][SUSR(
CSTR):"urn:sgws:identity::93390849266154004343:root"][SBAI(CSTR):"
93390849266154004343"][SBAC(CSTR):"bhavna"][S3BK(CSTR):"test"][S3K
Y(CSTR):"large-
object"][CBID(UI64):0x077EA25F3B36C69A][UUID(CSTR):"A80219A2-CD1E-
466F-9094-
B9C0FDE2FFA3"][CSIZ(UI64):6040000000][MTME(UI64):1672943621338958]
[AVER(UI32):10][ATIM(UI64):1672944425525999][ATYP(FC32):SPUT][ANID
(UI32):12220829][AMID(FC32):S3RQ][ATID(UI64):4333283179807659119]]
```

3. If Client Writes aren't Normal, use the tenant ID from the alert to identify the tenant:
   a. Go to **SUPPORT** > **Tools** > **Logs**. Collect application logs for the Storage Node in the alert. Specify 15 minutes before and after the alert.
   b. Extract the file and go to `bycast.log`:

      ```
      /GID<grid_id>_<time_stamp>/<site_node>/<time_stamp>/grid/bycast.log
      ```

   c. Search the log for `method=PUT` and identify the client in the `clientIP` field.

      Example bycast.log

      ```
      Jan  5 18:33:41 BHAVNAJ-DC1-S1-2-65 ADE: |12220829 1870864574 S3RQ
      %CEA 2023-01-05T18:33:41.208790| NOTICE   1404 af23cb66b7e3efa5 S3RQ:
      EVENT_PROCESS_CREATE - connection=1672943621106262 method=PUT
      name=</test/4MiB-0> auth=<V4> clientIP=<10.96.99.127>
      ```

4. Inform tenants that the maximum PutObject size is 5 GiB and to use multipart uploads for objects greater than 5 GiB.
5. Ignore the alert for one week if the application has been changed.

Troubleshoot lost and missing object data

Troubleshoot lost and missing object data: Overview

Objects can be retrieved for several reasons, including read requests from a client application, background verifications of replicated object data, ILM re-evaluations, and the restoration of object data during the recovery of a Storage Node.

The StorageGRID system uses location information in an object's metadata to determine from which location to retrieve the object. If a copy of the object is not found in the expected location, the system attempts to retrieve another copy of the object from elsewhere in the system, assuming that the ILM policy contains a rule to make two or more copies of the object.

If this retrieval is successful, the StorageGRID system replaces the missing copy of the object. Otherwise, the **Objects lost** alert is triggered, as follows:

- For replicated copies, if another copy can't be retrieved, the object is considered lost, and the alert is triggered.
- For erasure-coded copies, if a copy can't be retrieved from the expected location, the Corrupt Copies Detected (ECOR) attribute is incremented by one before an attempt is made to retrieve a copy from another location. If no other copy is found, the alert is triggered.

You should investigate all **Objects lost** alerts immediately to determine the root cause of the loss and to determine if the object might still exist in an offline, or otherwise currently unavailable, Storage Node or Archive Node. See Investigate lost objects.

In the case where object data without copies is lost, there is no recovery solution. However, you must reset the Lost objects counter to prevent known lost objects from masking any new lost objects. See Reset lost and missing object counts.

Investigate lost objects

When the **Objects lost** alert is triggered, you must investigate immediately. Collect information about the affected objects and contact technical support.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- You must have the `Passwords.txt` file.

About this task

The **Objects lost** alert indicates that StorageGRID believes that there are no copies of an object in the grid. Data might have been permanently lost.

Investigate lost object alerts immediately. You might need to take action to prevent further data loss. In some cases, you might be able to restore a lost object if you take prompt action.

Steps

1. Select **NODES**.
2. Select ***Storage Node*** > **Objects**.
3. Review the number of Lost objects shown in the Object counts table.

   This number indicates the total number of objects this grid node detects as missing from the entire StorageGRID system. The value is the sum of the Lost objects counters of the Data store component within the LDR and DDS services.

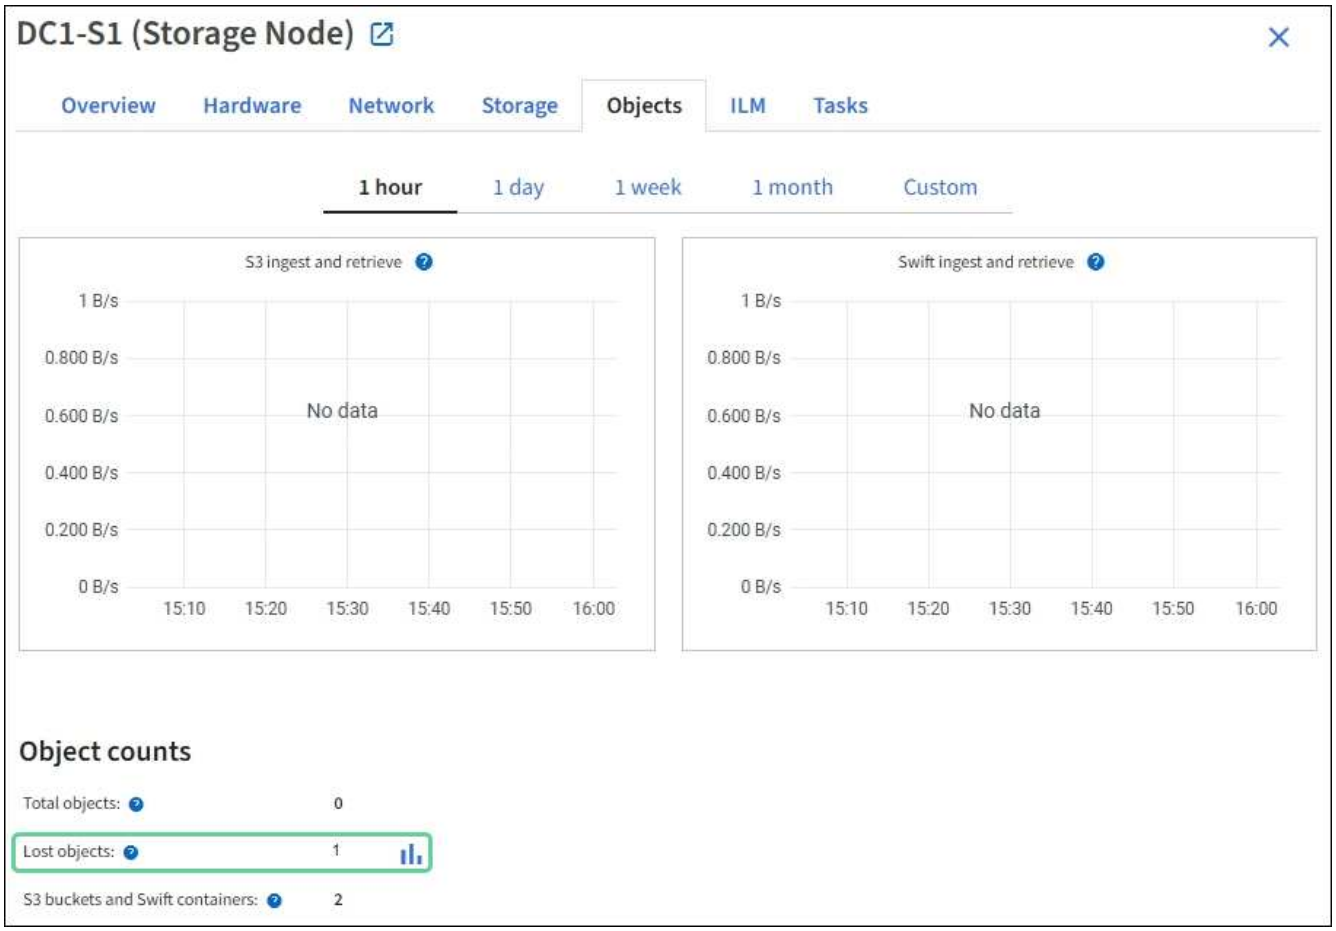

4. From an Admin Node, access the audit log to determine the unique identifier (UUID) of the object that triggered the **Objects lost** alert:
   a. Log in to the grid node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
      When you are logged in as root, the prompt changes from `$` to `#`.
   b. Change to the directory where the audit logs are located. Enter: `cd /var/local/log/`
   c. Use grep to extract the Object Lost (OLST) audit messages. Enter: `grep OLST audit_file_name`
   d. Note the UUID value included in the message.

```
>Admin: # grep OLST audit.log
2020-02-12T19:18:54.780426
[AUDT:[CBID(UI64):0x38186FE53E3C49A5][UUID(CSTR):926026C4-00A4-449B-
AC72-BCCA72DD1311]
[PATH(CSTR):"source/cats"][NOID(UI32):12288733][VOLI(UI64):3222345986
][RSLT(FC32):NONE][AVER(UI32):10]
[ATIM(UI64):1581535134780426][ATYP(FC32):OLST][ANID(UI32):12448208][A
MID(FC32):ILMX][ATID(UI64):7729403978647354233]]
```

5. Use the `ObjectByUUID` command to find the object by its identifier (UUID), and then determine if data is at risk.
   a. Telnet to localhost 1402 to access the LDR console.
   b. Enter: `/proc/OBRP/ObjectByUUID UUID_value`

      In this first example, the object with `UUID 926026C4-00A4-449B-AC72-BCCA72DD1311` has two locations listed.

```
ade 12448208: /proc/OBRP > ObjectByUUID 926026C4-00A4-449B-AC72-
BCCA72DD1311

{
    "TYPE(Object Type)": "Data object",
    "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
    "NAME": "cats",
    "CBID": "0x38186FE53E3C49A5",
    "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-
ACBB00BB82DD",
    "PPTH(Parent path)": "source",
    "META": {
        "BASE(Protocol metadata)": {
            "PAWS(S3 protocol version)": "2",
            "ACCT(S3 account ID)": "44084621669730638018",
            "*ctp(HTTP content MIME type)": "binary/octet-stream"
        },
        "BYCB(System metadata)": {
            "CSIZ(Plaintext object size)": "5242880",
            "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
            "BSIZ(Content block size)": "5252084",
            "CVER(Content block version)": "196612",
            "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
            "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
```

```
            "ITME": "1581534970983000"
        },
        "CMSM": {
            "LATM(Object last access time)": "2020-02-
12T19:16:10.983000"
        },
        "AWS3": {
            "LOCC": "us-east-1"
        }
    },
    "CLCO\(Locations\)": \[
        \{
            "Location Type": "CLDI\(Location online\)",
            "NOID\(Node ID\)": "12448208",
            "VOLI\(Volume ID\)": "3222345473",
            "Object File Path":
"/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila\#3udu",
            "LTIM\(Location timestamp\)": "2020-02-
12T19:36:17.880569"
        \},
        \{
            "Location Type": "CLDI\(Location online\)",
            "NOID\(Node ID\)": "12288733",
            "VOLI\(Volume ID\)": "3222345984",
            "Object File Path":
"/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb\#3s;L",
            "LTIM\(Location timestamp\)": "2020-02-
12T19:36:17.934425"
        }
    ]
}
```


In the second example, the object with `UUID 926026C4-00A4-449B-AC72-BCCA72DD1311` has no locations listed.

```
ade 12448208: / > /proc/OBRP/ObjectByUUID 926026C4-00A4-449B-AC72-
BCCA72DD1311

{
    "TYPE(Object Type)": "Data object",
    "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
    "NAME": "cats",
    "CBID": "0x38186FE53E3C49A5",
    "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-
ACBB00BB82DD",
    "PPTH(Parent path)": "source",
    "META": {
        "BASE(Protocol metadata)": {
            "PAWS(S3 protocol version)": "2",
            "ACCT(S3 account ID)": "44084621669730638018",
            "*ctp(HTTP content MIME type)": "binary/octet-stream"
        },
        "BYCB(System metadata)": {
            "CSIZ(Plaintext object size)": "5242880",
            "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
            "BSIZ(Content block size)": "5252084",
            "CVER(Content block version)": "196612",
            "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
            "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
            "ITME": "1581534970983000"
        },
        "CMSM": {
            "LATM(Object last access time)": "2020-02-
12T19:16:10.983000"
        },
        "AWS3": {
            "LOCC": "us-east-1"
        }
    }
}
```

c. Review the output of /proc/OBRP/ObjectByUUID, and take the appropriate action:

Metadata	Conclusion
No object found ("ERROR":"")	If the object is not found, the message "ERROR":"" is returned. If the object is not found, you can reset the count of **Objects lost** to clear the alert. The lack of an object indicates that the object was intentionally deleted.
Locations > 0	If there are locations listed in the output, the **Objects lost** alert might be a false positive. Confirm that the objects exist. Use the Node ID and filepath listed in the output to confirm that the object file is in the listed location. (The procedure for searching for potentially lost objects explains how to use the Node ID to find the correct Storage Node.) If the objects exist, you can reset the count of **Objects lost** to clear the alert.
Locations = 0	If there are no locations listed in the output, the object is potentially missing. You can try to search for and restore the object yourself, or you can contact technical support. Technical support might ask you to determine if there is a storage recovery procedure in progress. See the information about restoring object data using Grid Manager and restoring object data to a storage volume.

Search for and restore potentially lost objects

It might be possible to find and restore objects that have triggered a Lost Objects (LOST) alarm and a **Object lost** alert and that you have identified as potentially lost.

Before you begin

- You have the UUID of any lost object, as identified in Investigate lost objects.
- You have the `Passwords.txt` file.

About this task

You can follow this procedure to look for replicated copies of the lost object elsewhere in the grid. In most cases, the lost object will not be found. However, in some cases, you might be able to find and restore a lost replicated object if you take prompt action.

> Contact technical support for assistance with this procedure.

Steps

1. From an Admin Node, search the audit logs for possible object locations:
   a. Log in to the grid node:
      i. Enter the following command: `ssh admin@grid_node_IP`

ii. Enter the password listed in the `Passwords.txt` file.

iii. Enter the following command to switch to root: `su -`

iv. Enter the password listed in the `Passwords.txt` file.
When you are logged in as root, the prompt changes from `$` to `#`.

b. Change to the directory where the audit logs are located: `cd /var/local/log/`

c. Use grep to extract the audit messages associated with the potentially lost object and send them to an output file. Enter: `grep uuid-valueaudit_file_name > output_file_name`

For example:

```
Admin: # grep 926026C4-00A4-449B-AC72-BCCA72DD1311 audit.log >
messages_about_lost_object.txt
```

d. Use grep to extract the Location Lost (LLST) audit messages from this output file. Enter: `grep LLST output_file_name`

For example:

```
Admin: # grep LLST messages_about_lost_objects.txt
```

An LLST audit message looks like this example message.

```
[AUDT:\[NOID\(UI32\):12448208\][CBIL(UI64):0x38186FE53E3C49A5]
[UUID(CSTR):"926026C4-00A4-449B-AC72-BCCA72DD1311"][LTYP(FC32):CLDI]
[PCLD\(CSTR\):"/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%\#3tN6"\]
[TSRC(FC32):SYST][RSLT(FC32):NONE][AVER(UI32):10][ATIM(UI64):
1581535134379225][ATYP(FC32):LLST][ANID(UI32):12448208][AMID(FC32):CL
SM]
[ATID(UI64):7086871083190743409]]
```

e. Find the PCLD field and the NOID field in the LLST message.

If present, the value of PCLD is the complete path on disk to the missing replicated object copy. The value of NOID is the node id of the LDR where a copy of the object might be found.

If you find an object location, you might be able to restore the object.

f. Find the Storage Node associated with this LDR node ID. In the Grid Manager, select **SUPPORT** > **Tools** > **Grid topology**. Then select ***Data Center*** > ***Storage Node*** > **LDR**.

The Node ID for the LDR service is in the Node Information table. Review the information for each Storage Node until you find the one that hosts this LDR.

2. Determine if the object exists on the Storage Node indicated in the audit message:

a. Log in to the grid node:

i. Enter the following command: `ssh admin@grid_node_IP`

ii. Enter the password listed in the `Passwords.txt` file.

iii. Enter the following command to switch to root: `su -`

iv. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

b. Determine if the file path for the object exists.

For the file path of the object, use the value of PCLD from the LLST audit message.

For example, enter:

```
ls '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6'
```

> Always enclose the object file path in single quotes in commands to escape any special characters.

- If the object path is not found, the object is lost and can't be restored using this procedure. Contact technical support.
- If the object path is found, continue with the next step. You can attempt to restore the found object back to StorageGRID.

3. If the object path was found, attempt to restore the object to StorageGRID:

a. From the same Storage Node, change the ownership of the object file so that it can be managed by StorageGRID. Enter: `chown ldr-user:bycast 'file_path_of_object'`

b. Telnet to localhost 1402 to access the LDR console. Enter: `telnet 0 1402`

c. Enter: `cd /proc/STOR`

d. Enter: `Object_Found 'file_path_of_object'`

For example, enter:

```
Object_Found '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6'
```

Issuing the `Object\_Found` command notifies the grid of the object's location. It also triggers the active ILM policies, which make additional copies as specified in each policy.

> If the Storage Node where you found the object is offline, you can copy the object to any Storage Node that is online. Place the object in any /var/local/rangedb directory of the online Storage Node. Then, issue the `Object\_Found` command using that file path to the object.

- If the object can't be restored, the `Object\_Found` command fails. Contact technical support.
- If the object was successfully restored to StorageGRID, a success message appears. For example:

```
ade 12448208: /proc/STOR > Object_Found
'/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6'

ade 12448208: /proc/STOR > Object found succeeded.
First packet of file was valid. Extracted key: 38186FE53E3C49A5
Renamed '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6' to
'/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila#3udu'
```

Continue with the next step.

4. If the object was successfully restored to StorageGRID, verify that new locations were created.
   a. Enter: `cd /proc/OBRP`
   b. Enter: `ObjectByUUID UUID_value`

   The following example shows that there are two locations for the object with UUID 926026C4-00A4-449B-AC72-BCCA72DD1311.

```
ade 12448208: /proc/OBRP > ObjectByUUID 926026C4-00A4-449B-AC72-
BCCA72DD1311

{
    "TYPE(Object Type)": "Data object",
    "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
    "NAME": "cats",
    "CBID": "0x38186FE53E3C49A5",
    "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-
ACBB00BB82DD",
    "PPTH(Parent path)": "source",
    "META": {
        "BASE(Protocol metadata)": {
            "PAWS(S3 protocol version)": "2",
            "ACCT(S3 account ID)": "44084621669730638018",
            "*ctp(HTTP content MIME type)": "binary/octet-stream"
        },
        "BYCB(System metadata)": {
            "CSIZ(Plaintext object size)": "5242880",
            "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
            "BSIZ(Content block size)": "5252084",
            "CVER(Content block version)": "196612",
            "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
            "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
            "ITME": "1581534970983000"
```

```
        },
        "CMSM": {
            "LATM(Object last access time)": "2020-02-
12T19:16:10.983000"
        },
        "AWS3": {
            "LOCC": "us-east-1"
        }
    },
    "CLCO\(Locations\)": \[
        \{
            "Location Type": "CLDI\(Location online\)",
            "NOID\(Node ID\)": "12448208",
            "VOLI\(Volume ID\)": "3222345473",
            "Object File Path":
"/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila\#3udu",
            "LTIM\(Location timestamp\)": "2020-02-
12T19:36:17.880569"
        \},
        \{
            "Location Type": "CLDI\(Location online\)",
            "NOID\(Node ID\)": "12288733",
            "VOLI\(Volume ID\)": "3222345984",
            "Object File Path":
"/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb\#3s;L",
            "LTIM\(Location timestamp\)": "2020-02-
12T19:36:17.934425"
        }
    ]
}
```

c. Sign out of the LDR console. Enter: `exit`

5. From an Admin Node, search the audit logs for the ORLM audit message for this object to confirm that information lifecycle management (ILM) has placed copies as required.

   a. Log in to the grid node:

      i. Enter the following command: `ssh admin@grid_node_IP`

      ii. Enter the password listed in the `Passwords.txt` file.

      iii. Enter the following command to switch to root: `su -`

      iv. Enter the password listed in the `Passwords.txt` file.
      When you are logged in as root, the prompt changes from `$` to `#`.

   b. Change to the directory where the audit logs are located: `cd /var/local/log/`

   c. Use grep to extract the audit messages associated with the object to an output file. Enter: `grep uuid-valueaudit_file_name > output_file_name`

For example:

```
Admin: # grep 926026C4-00A4-449B-AC72-BCCA72DD1311 audit.log >
messages_about_restored_object.txt
```

d. Use grep to extract the Object Rules Met (ORLM) audit messages from this output file. Enter: `grep ORLM output_file_name`

For example:

```
Admin: # grep ORLM messages_about_restored_object.txt
```

An ORLM audit message looks like this example message.

```
[AUDT:[CBID(UI64):0x38186FE53E3C49A5][RULE(CSTR):"Make 2 Copies"]
[STAT(FC32):DONE][CSIZ(UI64):0][UUID(CSTR):"926026C4-00A4-449B-AC72-
BCCA72DD1311"]
[LOCS(CSTR):"**CLDI 12828634 2148730112**, CLDI 12745543 2147552014"]
[RSLT(FC32):SUCS][AVER(UI32):10][ATYP(FC32):ORLM][ATIM(UI64):15633982
30669]
[ATID(UI64):15494889725796157557][ANID(UI32):13100453][AMID(FC32):BCM
S]]
```

e. Find the LOCS field in the audit message.

If present, the value of CLDI in LOCS is the node ID and the volume ID where an object copy has been created. This message shows that the ILM has been applied and that two object copies have been created in two locations in the grid.

6. Reset the lost and missing object counts in the Grid Manager.

Reset lost and missing object counts

After investigating the StorageGRID system and verifying that all recorded lost objects are permanently lost or that it is a false alarm, you can reset the value of the Lost Objects attribute to zero.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

You can reset the Lost Objects counter from either of the following pages:

- **SUPPORT** > **Tools** > **Grid topology** > ***Site*** > ***Storage Node*** > **LDR** > **Data Store** > **Overview** > **Main**
- **SUPPORT** > **Tools** > **Grid topology** > ***Site*** > ***Storage Node*** > **DDS** > **Data Store** > **Overview** > **Main**

These instructions show resetting the counter from the **LDR** > **Data Store** page.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select ***Site*** > ***Storage Node*** > **LDR** > **Data Store** > **Configuration** for the Storage Node that has the **Objects lost** alert or the LOST alarm.
3. Select **Reset Lost Objects Count**.

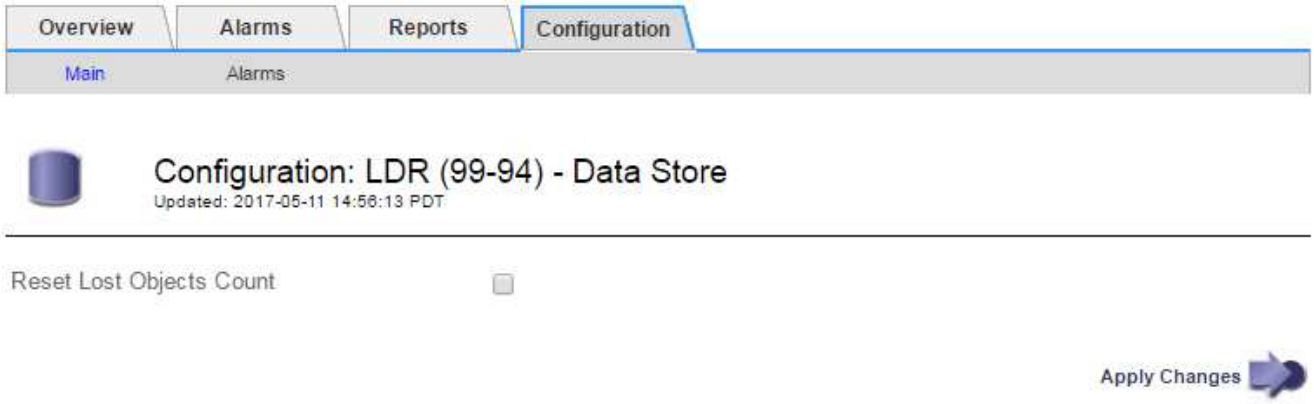

4. Click **Apply Changes**.

   The Lost Objects attribute is reset to 0 and the **Objects lost** alert and the LOST alarm clear, which can take a few minutes.

5. Optionally, reset other related attribute values that might have been incremented in the process of identifying the lost object.

   a. Select ***Site*** > ***Storage Node*** > **LDR** > **Erasure Coding** > **Configuration**.
   b. Select **Reset Reads Failure Count** and **Reset Corrupt Copies Detected Count**.
   c. Click **Apply Changes**.
   d. Select ***Site*** > ***Storage Node*** > **LDR** > **Verification** > **Configuration**.
   e. Select **Reset Missing Objects Count** and **Reset Corrupt Objects Count**.
   f. If you are confident that quarantined objects aren't required, you can select **Delete Quarantined Objects**.

      Quarantined objects are created when background verification identifies a corrupt replicated object copy. In most cases StorageGRID automatically replaces the corrupt object, and it is safe to delete the quarantined objects. However, if the **Objects lost** alert or the LOST alarm is triggered, technical support might want to access the quarantined objects.

   g. Click **Apply Changes**.

   It can take a few moments for the attributes to reset after you click **Apply Changes**.

Troubleshoot the Low object data storage alert

The **Low object data storage** alert monitors how much space is available for storing object data on each Storage Node.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

The **Low object data storage** alert is triggered when the total amount of replicated and erasure-coded object data on a Storage Node meets one of the conditions configured in the alert rule.

By default, a major alert is triggered when this condition evaluates as true:

```
(storagegrid_storage_utilization_data_bytes/
(storagegrid_storage_utilization_data_bytes +
storagegrid_storage_utilization_usable_space_bytes)) >=0.90
```

In this condition:

- `storagegrid_storage_utilization_data_bytes` is an estimate of the total size of replicated and erasure-coded object data for a Storage Node.
- `storagegrid_storage_utilization_usable_space_bytes` is the total amount of object storage space remaining for a Storage Node.

If a major or minor **Low object data storage** alert is triggered, you should perform an expansion procedure as soon as possible.

Steps

1. Select **ALERTS** > **Current**.

   The Alerts page appears.

2. From the table of alerts, expand the **Low object data storage** alert group, if required, and select the alert you want to view.

   Select the alert, not the heading for a group of alerts.

3. Review the details in the dialog box, and note the following:
   - Time triggered
   - The name of the site and node
   - The current values of the metrics for this alert
4. Select **NODES** > ***Storage Node or Site*** > **Storage**.
5. Position your cursor over the Storage Used - Object Data graph.

   The following values are shown:

   - **Used (%)**: The percentage of the Total usable space that has been used for object data.
   - **Used**: The amount of the Total usable space that has been used for object data.
   - **Replicated data**: An estimate of the amount of replicated object data on this node, site, or grid.
   - **Erasure-coded data**: An estimate of the amount of erasure-coded object data on this node, site, or

grid.

- **Total**: The total amount of usable space on this node, site, or grid.
  The Used value is the `storagegrid_storage_utilization_data_bytes` metric.

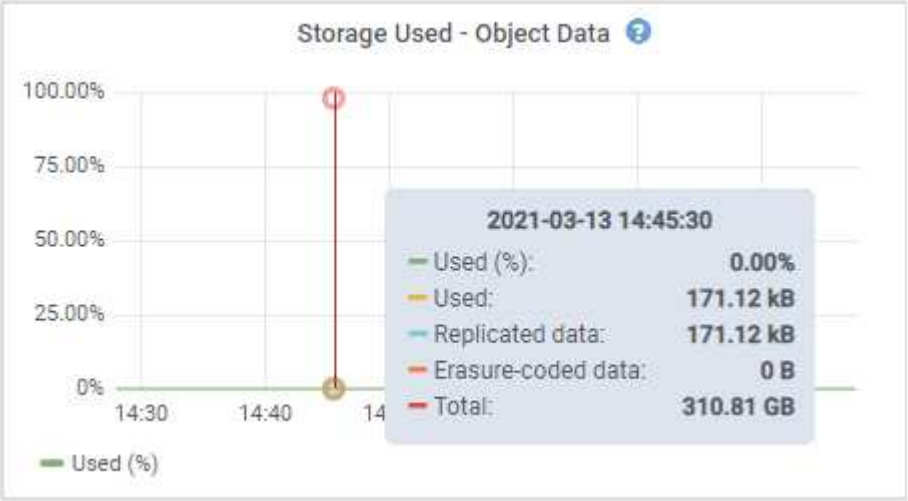

6. Select the time controls above the graph to view storage use over different time periods.

   Looking at storage use over time can help you understand how much storage was used before and after the alert was triggered and can help you estimate how long it might take for the node's remaining space to become full.

7. As soon as possible, add storage capacity to your grid.

   You can add storage volumes (LUNs) to existing Storage Nodes, or you can add new Storage Nodes.

   For more information, see Manage full Storage Nodes.

Related information

Troubleshoot the Storage Status (SSTS) alarm (legacy)

Troubleshoot Low read-only watermark override alerts

If you use custom values for storage volume watermarks, you might need to resolve the **Low read-only watermark override** alert. If possible, you should update your system to start using the optimized values.

In previous releases, the three storage volume watermarks were global settings — the same values applied to every storage volume on every Storage Node. Starting in StorageGRID 11.6, the software can optimize these watermarks for each storage volume, based on the size of the Storage Node and the relative capacity of the volume.

When you upgrade to StorageGRID 11.6 or higher, optimized read-only and read-write watermarks are automatically applied to all storage volumes, unless either of the following is true:

- Your system is close to capacity and would not be able to accept new data if optimized watermarks were applied. StorageGRID will not change watermark settings in this case.
- You previously set any of the storage volume watermarks to a custom value. StorageGRID will not override

custom watermark settings with optimized values. However, StorageGRID might trigger the **Low read-only watermark override** alert if your custom value for the Storage Volume Soft Read-Only Watermark is too small.

Understand the alert

If you use custom values for storage volume watermarks, the **Low read-only watermark override** alert might be triggered for one or more Storage Nodes.

Each instance of the alert indicates that the custom value of the **Storage Volume Soft Read-Only Watermark** is smaller than the minimum optimized value for that Storage Node. If you continue to use the custom setting, the Storage Node might run critically low on space before it can safely transition to the read-only state. Some storage volumes might become inaccessible (automatically unmounted) when the node reaches capacity.

For example, suppose you previously set the **Storage Volume Soft Read-Only Watermark** to 5 GB. Now suppose that StorageGRID has calculated the following optimized values for the four storage volumes in Storage Node A:

Volume 0	12 GB
Volume 1	12 GB
Volume 2	11 GB
Volume 3	15 GB

The **Low read-only watermark override** alert is triggered for Storage Node A because your custom watermark (5 GB) is smaller than the minimum optimized value for all volumes in that node (11 GB). If you continue using the custom setting, the node might run critically low on space before it can safely transition to the read-only state.

Resolve the alert

Follow these steps if one or more **Low read-only watermark override** alerts have been triggered. You can also use these instructions if you currently use custom watermark settings and want to start using optimized settings even if no alerts have been triggered.

Before you begin

- You have completed the upgrade to StorageGRID 11.6 or higher.
- You are signed in to the Grid Manager using a supported web browser.
- You have the Root access permission.

About this task

You can resolve the **Low read-only watermark override** alert by updating custom watermark settings to the new watermark overrides. However, if one or more Storage Nodes are close to full or you have special ILM requirements, you should first view the optimized storage watermarks and determine if it is safe to use them.

Assess object data usage for entire grid

Steps

1. Select **NODES**.

2. For each site in the grid, expand the list of nodes.
3. Review the percentage values shown in the **Object data used** column for each Storage Node at every site.

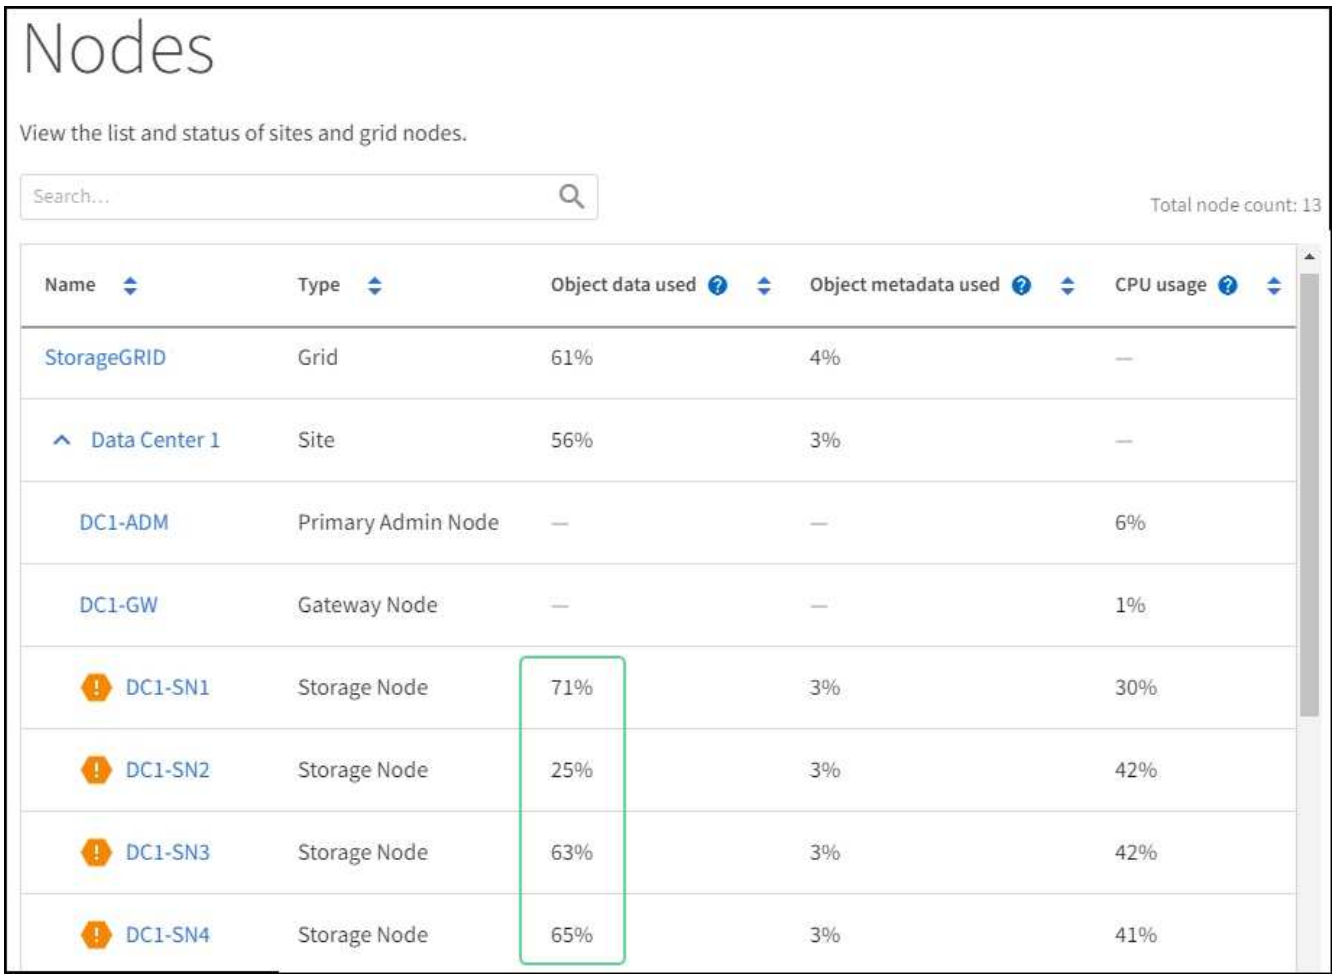

Name	Type	Object data used	Object metadata used	CPU usage
StorageGRID	Grid	61%	4%	—
Data Center 1	Site	56%	3%	—
DC1-ADM	Primary Admin Node	—	—	6%
DC1-GW	Gateway Node	—	—	1%
DC1-SN1	Storage Node	71%	3%	30%
DC1-SN2	Storage Node	25%	3%	42%
DC1-SN3	Storage Node	63%	3%	42%
DC1-SN4	Storage Node	65%	3%	41%

4. Follow the appropriate step:
   a. If none of the Storage Nodes are close to full (for example, all **Object data used** values are less than 80%), you can start using the override settings. Go to Use optimized watermarks.
   b. If ILM rules use Strict ingest behavior or if specific storage pools are close to full, perform the steps in View optimized storage watermarks and Determine if you can use optimized watermarks.

View optimized storage watermarks

StorageGRID uses two Prometheus metrics to show the optimized values it has calculated for the **Storage Volume Soft Read-Only Watermark**. You can view the minimum and maximum optimized values for each Storage Node in your grid.

Steps

1. Select **SUPPORT** > **Tools** > **Metrics**.
2. In the Prometheus section, select the link to access the Prometheus user interface.
3. To see the recommended minimum soft read-only watermark, enter the following Prometheus metric, and select **Execute**:

```
storagegrid_storage_volume_minimum_optimized_soft_readonly_watermark
```

The last column shows the minimum optimized value of the Soft Read-Only Watermark for all storage volumes on each Storage Node. If this value is greater than the custom setting for the **Storage Volume Soft Read-Only Watermark**, the **Low read-only watermark override** alert is triggered for the Storage Node.

4. To see the recommended maximum soft read-only watermark, enter the following Prometheus metric, and select **Execute**:

   ```
   storagegrid_storage_volume_maximum_optimized_soft_readonly_watermark
   ```

   The last column shows the maximum optimized value of the Soft Read-Only Watermark for all storage volumes on each Storage Node.

5. Note the maximum optimized value for each Storage Node.

Determine if you can use optimized watermarks

Steps

1. Select **NODES**.
2. Repeat these steps for each online Storage Node:
   a. Select ***Storage Node*** > **Storage**.
   b. Scroll down to the Object Stores table.
   c. Compare the **Available** value for each object store (volume) to the maximum optimized watermark you noted for that Storage Node.
3. If at least one volume on every online Storage Node has more space available than maximum optimized watermark for that node, go to Use optimized watermarks to start using the optimized watermarks.

   Otherwise, expand the grid as soon as possible. Either add storage volumes to an existing node or add new Storage Nodes. Then, go to Use optimized watermarks to update watermark settings.

4. If you need to continue using custom values for the storage volume watermarks, silence or disable the **Low read-only watermark override** alert.

   > The same custom watermark values are applied to every storage volume on every Storage Node. Using smaller-than-recommended values for storage volume watermarks might cause some storage volumes to become inaccessible (automatically unmounted) when the node reaches capacity.

Use optimized watermarks

Steps

1. Go to **SUPPORT** > **Other** > **Storage watermarks**.
2. Select the **Use optimized values** checkbox.
3. Select **Save**.

Optimized storage volume watermark settings are now in effect for each storage volume, based on the size of the Storage Node and the relative capacity of the volume.

Troubleshoot the Storage Status (SSTS) alarm

The Storage Status (SSTS) alarm is triggered if a Storage Node has insufficient free space remaining for object storage.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.

About this task

The SSTS (Storage Status) alarm is triggered at the Notice level when the amount of free space on every volume in a Storage Node falls below the value of the Storage Volume Soft Read Only Watermark (**CONFIGURATION** > **System** > **Storage options**).

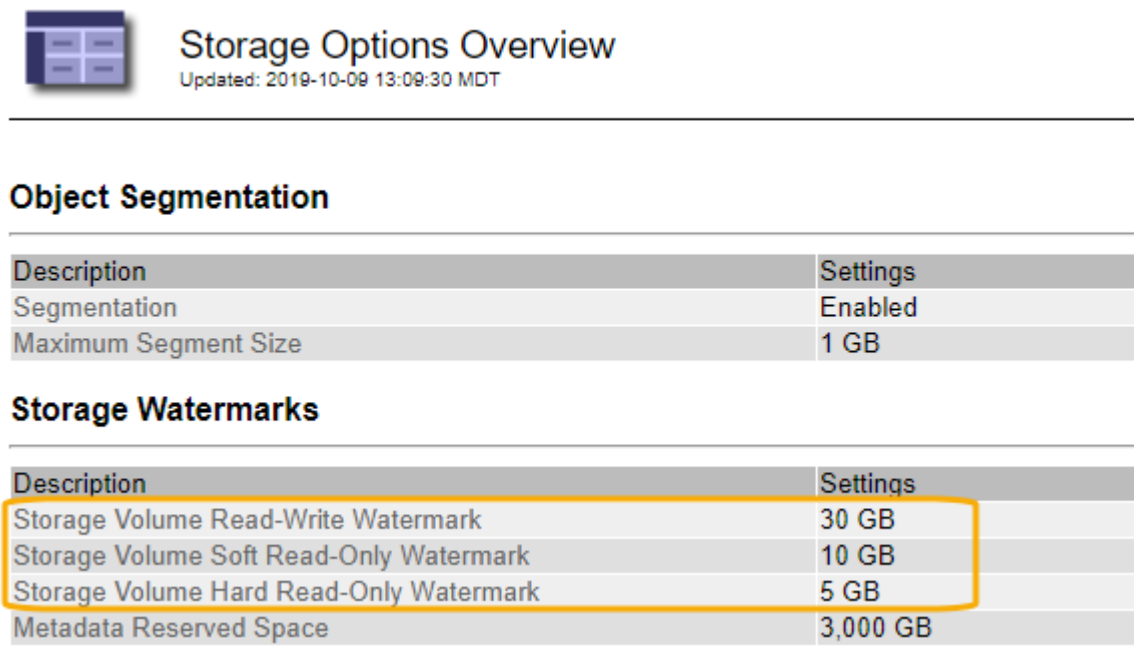

For example, suppose the Storage Volume Soft Read-Only Watermark is set to 10 GB, which is its default value. The SSTS alarm is triggered if less than 10 GB of usable space remains on each storage volume in the Storage Node. If any of the volumes has 10 GB or more of available space, the alarm is not triggered.

If an SSTS alarm has been triggered, you can follow these steps to better understand the issue.

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Current alarms**.
2. From the Service column, select the data center, node, and service that are associated with the SSTS alarm.

   The Grid Topology page appears. The Alarms tab shows the active alarms for the node and service you selected.

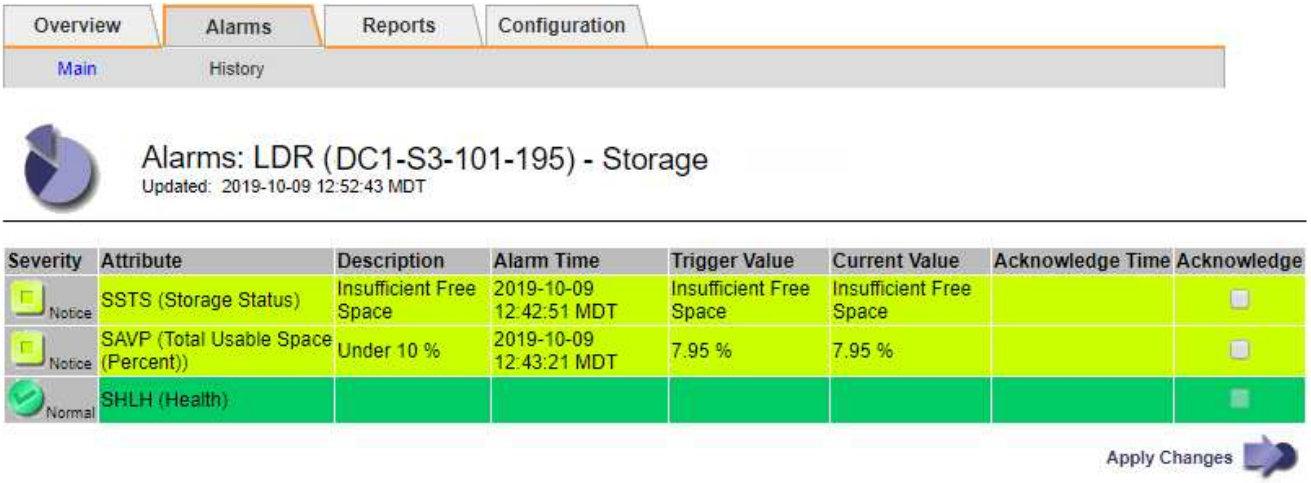

In this example, both the SSTS (Storage Status) and SAVP (Total Usable Space (Percent)) alarms have been triggered at the Notice level.

> Typically, both the SSTS alarm and the SAVP alarm are triggered at about the same time; however, whether both alarms are triggered depends on the the watermark setting in GB and the SAVP alarm setting in percent.

3. To determine how much usable space is actually available, select **LDR** > **Storage** > **Overview**, and find the Total Usable Space (STAS) attribute.

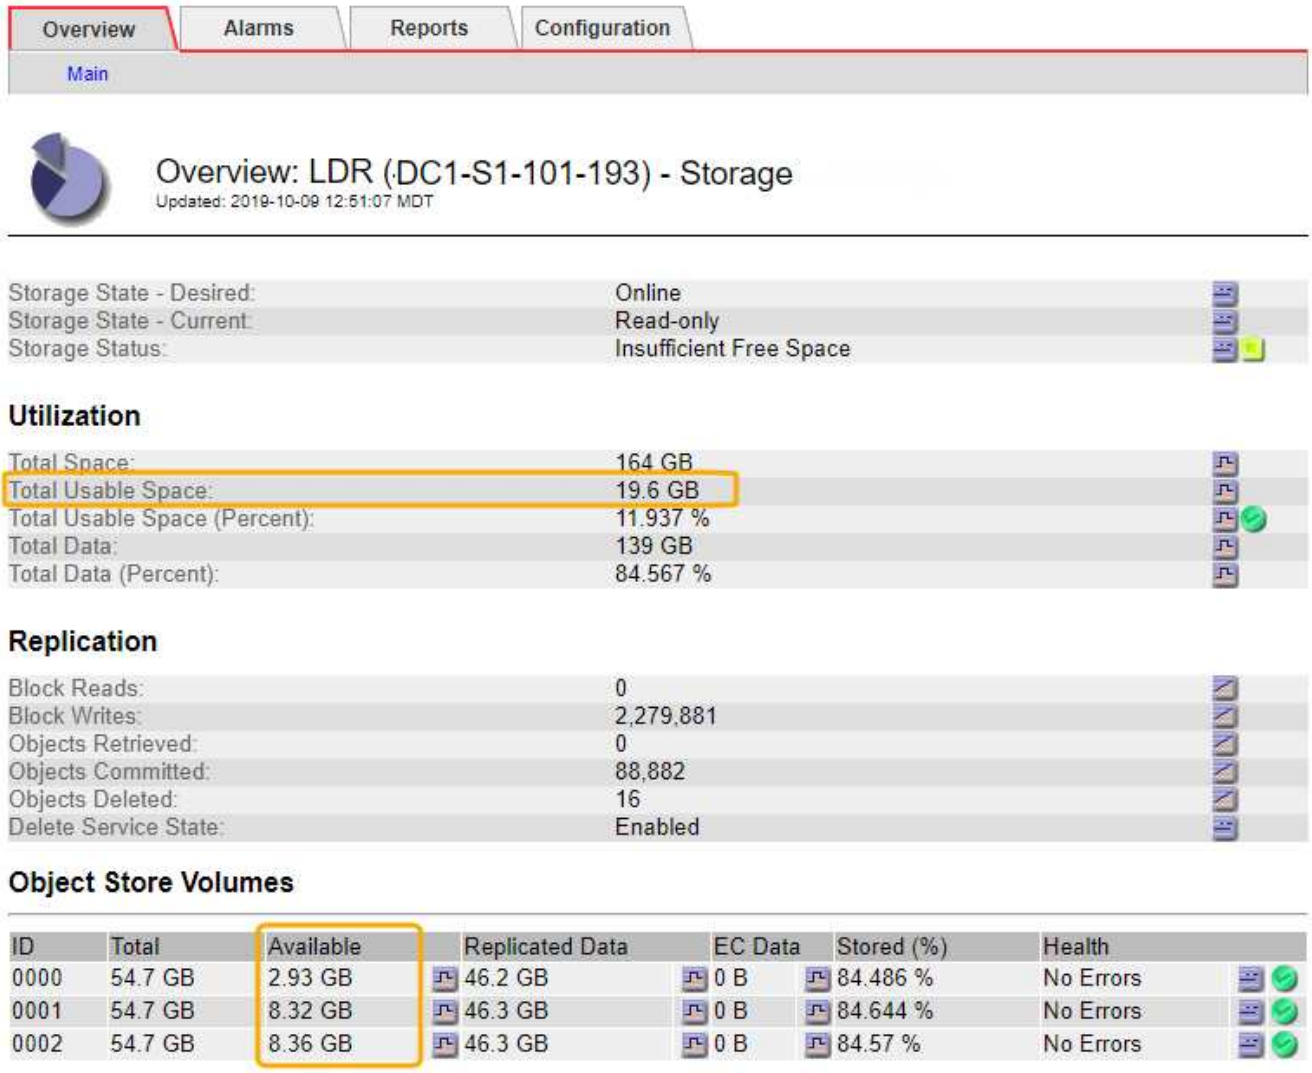

In this example, only 19.6 GB of the 164 GB of space on this Storage Node remains available. Note that the total value is the sum of the **Available** values for the three object store volumes. The SSTS alarm was triggered because each of the three storage volumes had less than 10 GB of available space.

4. To understand how storage has been used over time, select the **Reports** tab, and plot Total Usable Space over the last few hours.

   In this example, Total Usable Space dropped from roughly 155 GB at 12:00 to 20 GB at 12:35, which corresponds to the time at which the SSTS alarm was triggered.

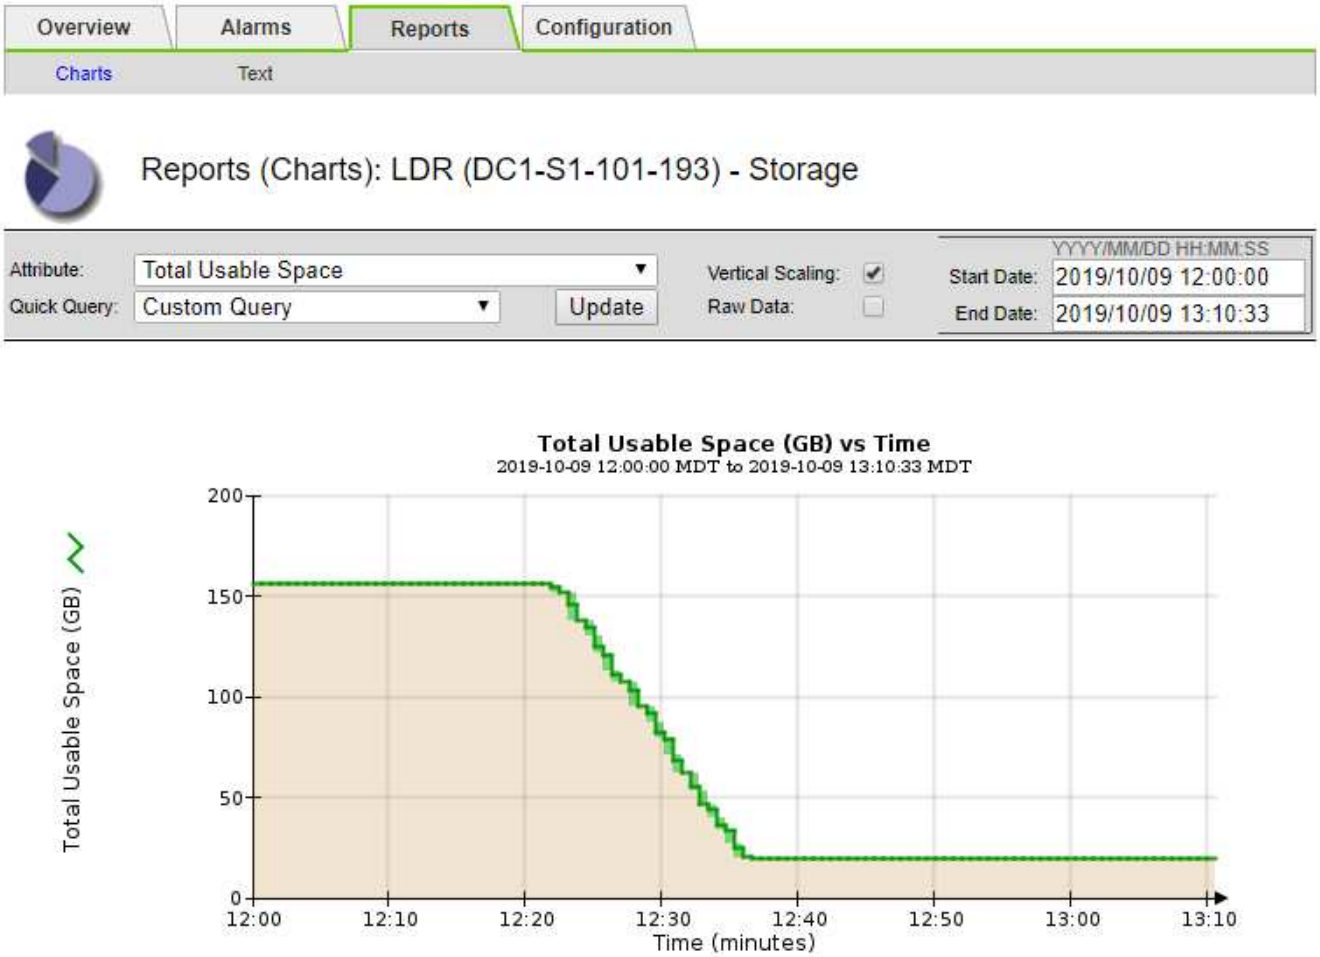

5. To understand how storage is being used as a percent of the total, plot Total Usable Space (Percent) over the last few hours.

   In this example, the total usable space dropped from 95% to just over 10% at approximately the same time.

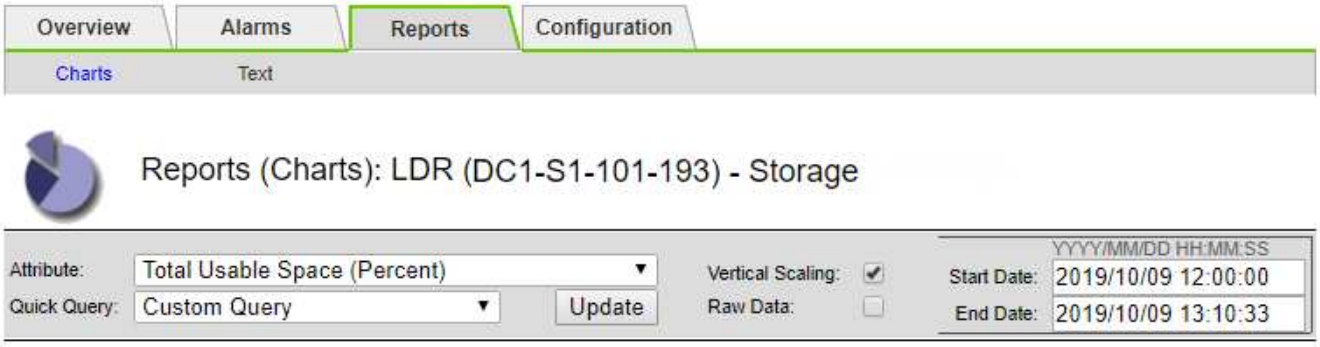

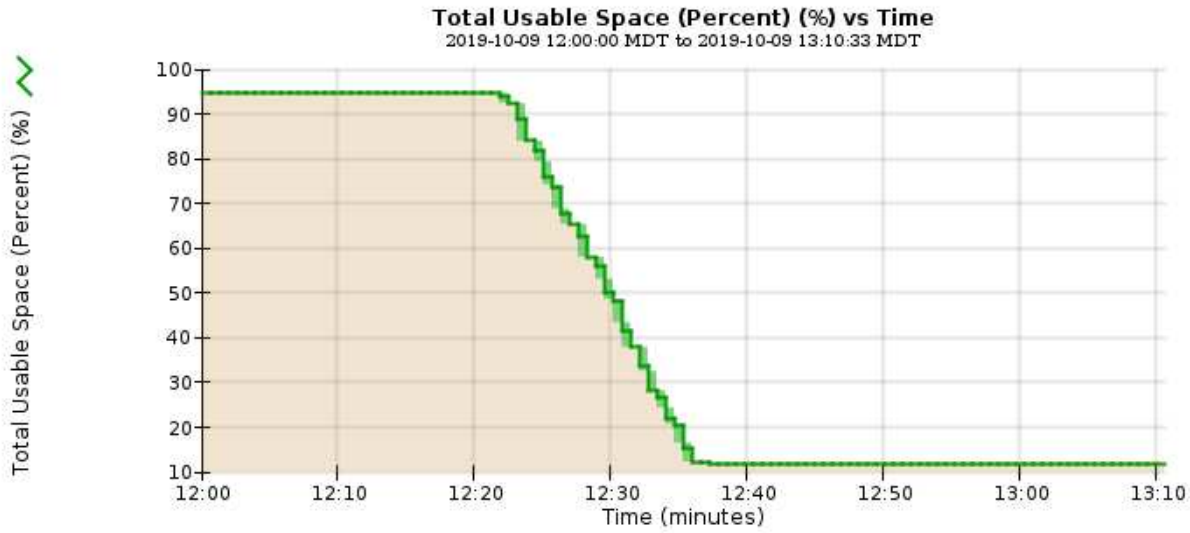

6. As required, add storage capacity.

   Also see Manage full Storage Nodes.

Troubleshoot delivery of platform services messages (SMTT alarm)

The Total Events (SMTT) alarm is triggered in the Grid Manager if a platform service message is delivered to an destination that can't accept the data.

About this task

For example, an S3 multipart upload can succeed even though the associated replication or notification message can't be delivered to the configured endpoint. Or, a message for CloudMirror replication can fail to be delivered if the metadata is too long.

The SMTT alarm contains a Last Event message that says, `Failed to publish notifications for bucket-name object key` for the last object whose notification failed.

Event messages are also listed in the `/var/local/log/bycast-err.log` log file. See the Log files reference.

For additional information, see the Troubleshoot platform services. You might need to access the tenant from the Tenant Manager to debug a platform service error.

Steps

1. To view the alarm, select **NODES** > ***site*** > ***grid node*** > **Events**.
2. View Last Event at the top of the table.

   Event messages are also listed in `/var/local/log/bycast-err.log`.

3. Follow the guidance provided in the SMTT alarm contents to correct the issue.
4. Select **Reset event counts**.
5. Notify the tenant of the objects whose platform services messages have not been delivered.
6. Instruct the tenant to trigger the failed replication or notification by updating the object's metadata or tags.

Troubleshoot metadata issues

You can perform several tasks to help determine the source of metadata problems.

Low metadata storage alert

If the **Low metadata storage** alert is triggered, you must add new Storage Nodes.

Before you begin

- You are signed in to the Grid Manager using a supported web browser.

About this task

StorageGRID reserves a certain amount of space on volume 0 of each Storage Node for object metadata. This space is known as the actual reserved space, and it is subdivided into the space allowed for object metadata (the allowed metadata space) and the space required for essential database operations, such as compaction and repair. The allowed metadata space governs overall object capacity.

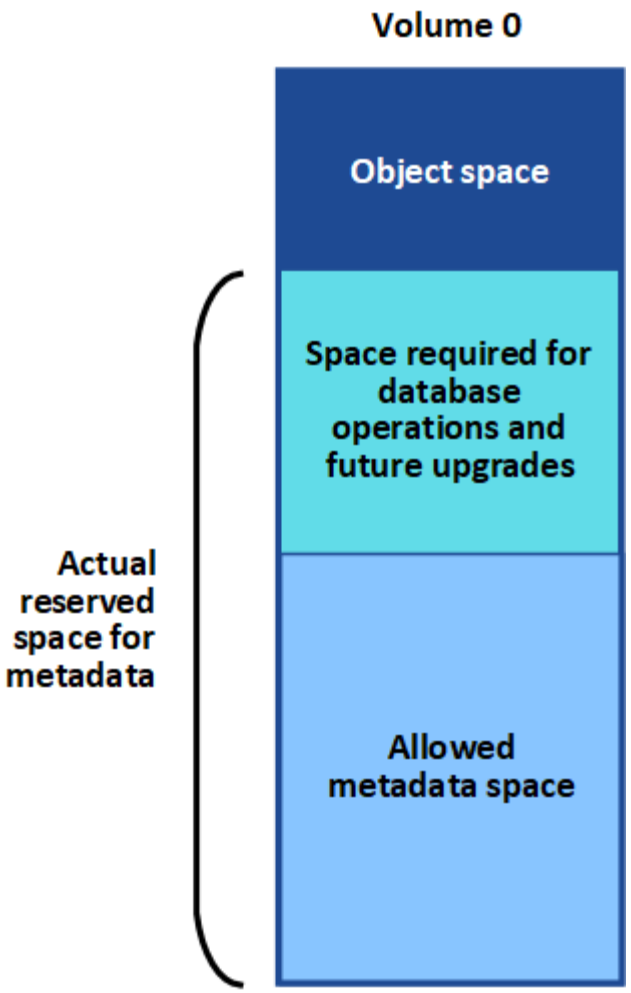

If object metadata consumes more than 100% of the space allowed for metadata, database operations can't run efficiently and errors will occur.

You can monitor object metadata capacity for each Storage Node to help you anticipate errors and correct them before they occur.

StorageGRID uses the following Prometheus metric to measure how full the allowed metadata space is:

```
storagegrid_storage_utilization_metadata_bytes/storagegrid_storage_utiliza
tion_metadata_allowed_bytes
```

When this Prometheus expression reaches certain thresholds, the **Low metadata storage** alert is triggered.

- **Minor**: Object metadata is using 70% or more of the allowed metadata space. You should add new Storage Nodes as soon as possible.
- **Major**: Object metadata is using 90% or more of the allowed metadata space. You must add new Storage Nodes immediately.

> When object metadata is using 90% or more of the allowed metadata space, a warning appears on the dashboard. If this warning appears, you must add new Storage Nodes immediately. You must never allow object metadata to use more than 100% of the allowed space.

- **Critical**: Object metadata is using 100% or more of the allowed metadata space and is starting to consume the space required for essential database operations. You must stop the ingest of new objects, and you must add new Storage Nodes immediately.

In the following example, object metadata is using more than 100% of the allowed metadata space. This is a critical situation, which will result in inefficient database operation and errors.

The following Storage Nodes are using more than 90% of the space allowed for object metadata:

Node	% Used	Used	Allowed
DC1-S2-227	104.51%	6.73 GB	6.44 GB
DC1-S3-228	104.36%	6.72 GB	6.44 GB
DC2-S2-233	104.20%	6.71 GB	6.44 GB
DC1-S1-226	104.20%	6.71 GB	6.44 GB
DC2-S3-234	103.43%	6.66 GB	6.44 GB

Undesirable results can occur if object metadata uses more than 100% of the allowed space. You must add new Storage Nodes immediately or contact support.

> If the size of volume 0 is smaller than the Metadata Reserved Space storage option (for example, in a non-production environment), the calculation for the **Low metadata storage** alert might be inaccurate.

Steps

1. Select **ALERTS** > **Current**.
2. From the table of alerts, expand the **Low metadata storage** alert group, if required, and select the specific alert you want to view.
3. Review the details in the alert dialog box.
4. If a major or critical **Low metadata storage** alert has been triggered, perform an expansion to add Storage Nodes immediately.

   > Because StorageGRID keeps complete copies of all object metadata at each site, the metadata capacity of the entire grid is limited by the metadata capacity of the smallest site. If you need to add metadata capacity to one site, you should also expand any other sites by the same number of Storage Nodes.

   After you perform the expansion, StorageGRID redistributes the existing object metadata to the new nodes, which increases the overall metadata capacity of the grid. No user action is required. The **Low metadata storage** alert is cleared.

Services: Status - Cassandra (SVST) alarm

The Services: Status - Cassandra (SVST) alarm indicates that you might need to rebuild the Cassandra database for a Storage Node. Cassandra is used as the metadata store for StorageGRID.

Before you begin

- You must be signed in to the Grid Manager using a supported web browser.
- You have specific access permissions.
- You must have the `Passwords.txt` file.

About this task

If Cassandra is stopped for more than 15 days (for example, the Storage Node is powered off), Cassandra will not start when the node is brought back online. You must rebuild the Cassandra database for the affected DDS

service.

You can run diagnostics to obtain additional information about the current state of your grid.

> If two or more of the Cassandra database services are down for more than 15 days, contact technical support, and don't proceed with the steps below.

Steps

1. Select **SUPPORT** > **Tools** > **Grid topology**.
2. Select ***Site*** > ***Storage Node*** > **SSM** > **Services** > **Alarms** > **Main** to display alarms.

   This example shows that the SVST alarm was triggered.

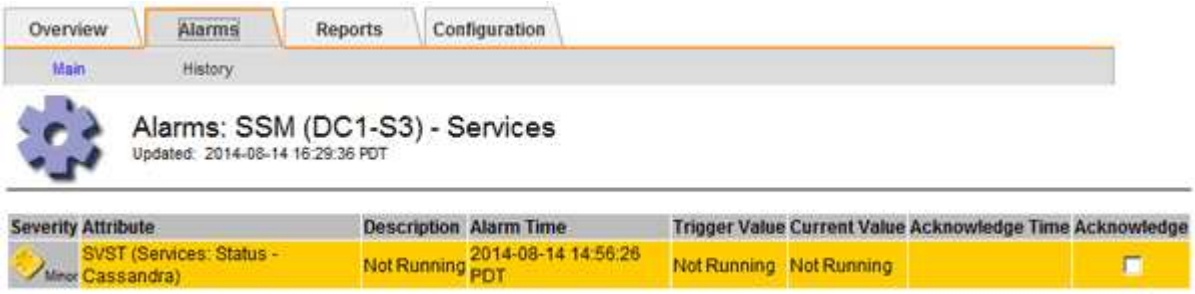

   The SSM Services Main page also indicates that Cassandra is not running.

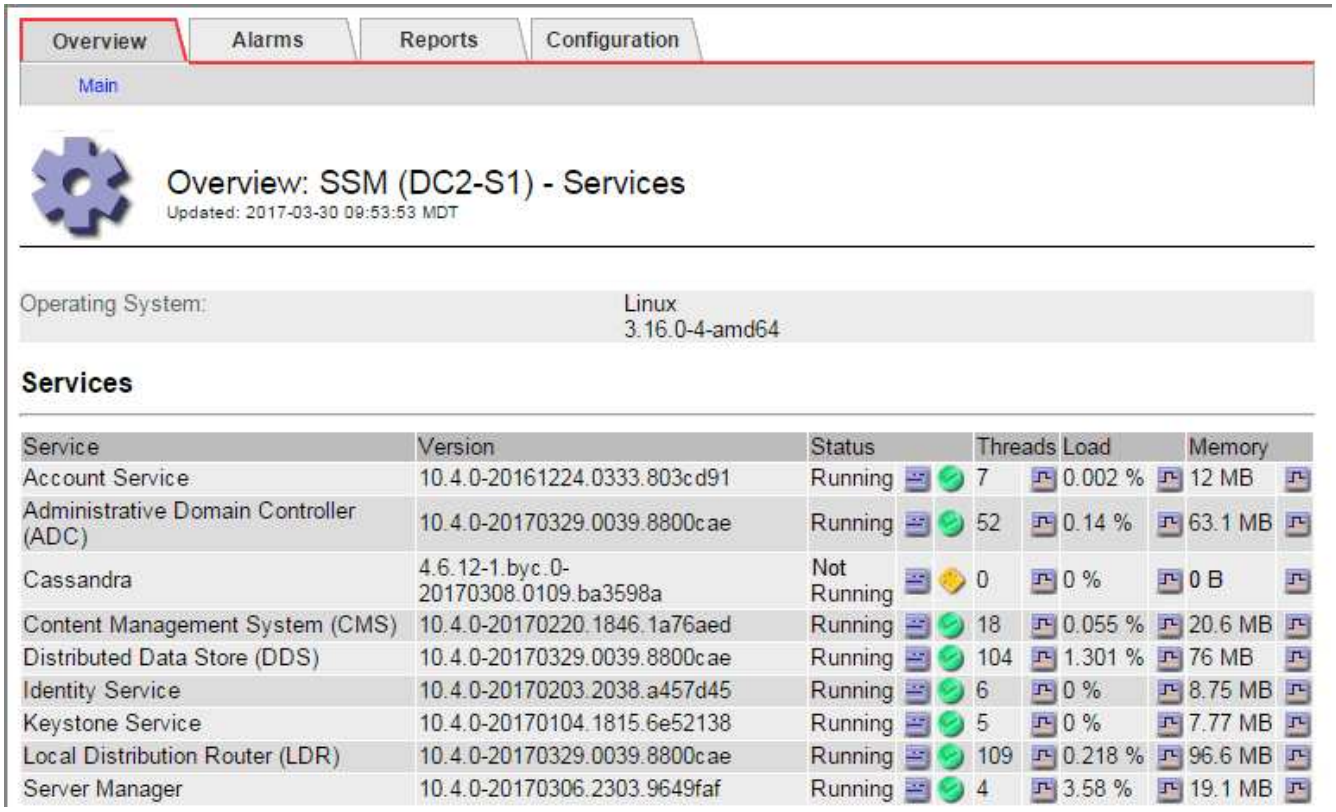

3. Try restarting Cassandra from the Storage Node:
   a. Log in to the grid node:
      i. Enter the following command: `ssh admin@grid_node_IP`

ii. Enter the password listed in the `Passwords.txt` file.

iii. Enter the following command to switch to root: `su -`

iv. Enter the password listed in the `Passwords.txt` file.
When you are logged in as root, the prompt changes from `$` to `#`.

b. Enter: `/etc/init.d/cassandra status`

c. If Cassandra is not running, restart it: `/etc/init.d/cassandra restart`

4. If Cassandra does not restart, determine how long Cassandra has been down. If Cassandra has been down for longer than 15 days, you must rebuild the Cassandra database.

> If two or more of the Cassandra database services are down, contact technical support, and don't proceed with the steps below.

You can determine how long Cassandra has been down by charting it or by reviewing the servermanager.log file.

5. To chart Cassandra:

   a. Select **SUPPORT** > **Tools** > **Grid topology**. Then select ***Site*** > ***Storage Node*** > **SSM** > **Services** > **Reports** > **Charts**.

   b. Select **Attribute** > **Service: Status - Cassandra**.

   c. For **Start Date**, enter a date that is at least 16 days before the current date. For **End Date**, enter the current date.

   d. Click **Update**.

   e. If the chart shows Cassandra as being down for more than 15 days, rebuild the Cassandra database.

   The following chart example shows that Cassandra has been down for at least 17 days.

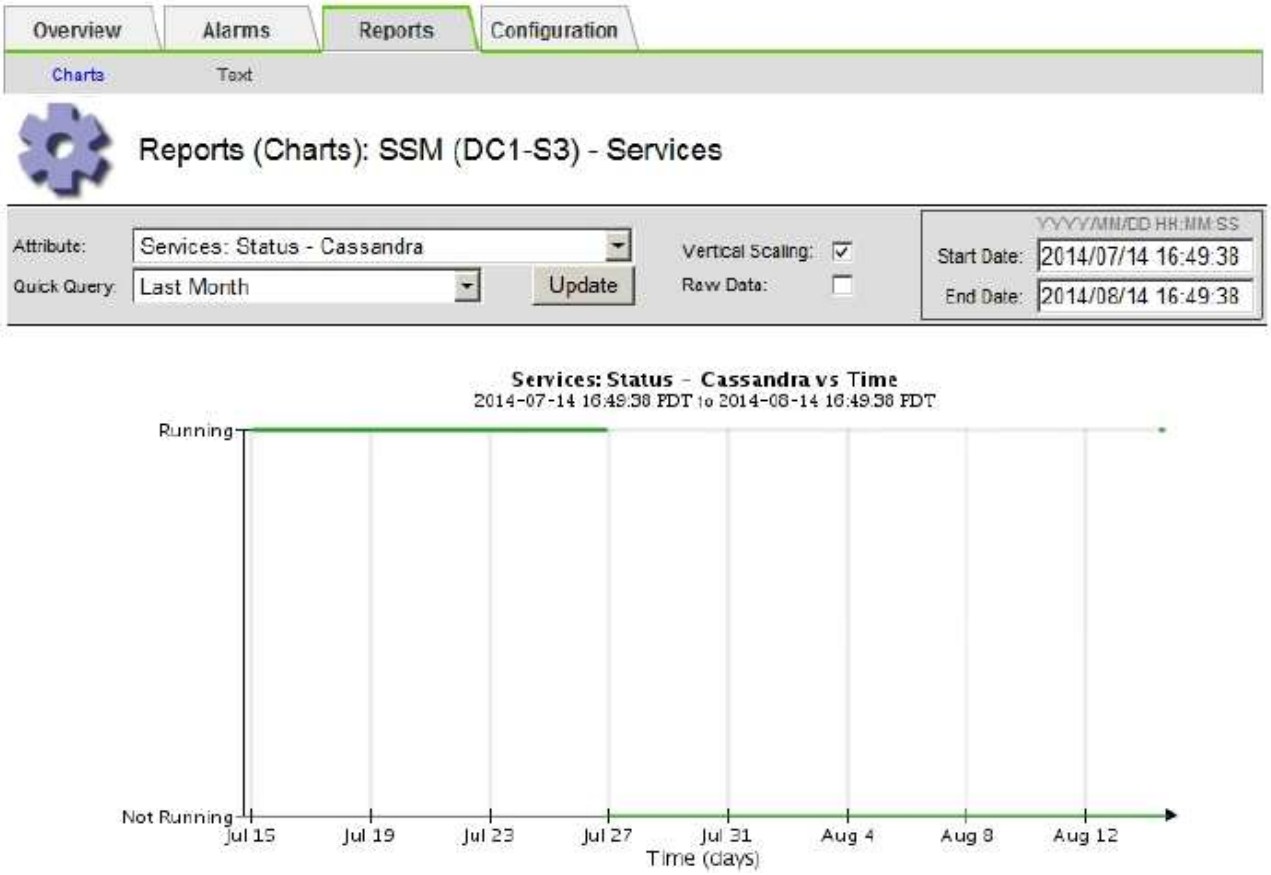

6. To review the servermanager.log file on the Storage Node:
   a. Log in to the grid node:
      i. Enter the following command: `ssh admin@grid_node_IP`
      ii. Enter the password listed in the `Passwords.txt` file.
      iii. Enter the following command to switch to root: `su -`
      iv. Enter the password listed in the `Passwords.txt` file.
      When you are logged in as root, the prompt changes from `$` to `#`.
   b. Enter: `cat /var/local/log/servermanager.log`

      The contents of the servermanager.log file are displayed.

      If Cassandra has been down for longer than 15 days, the following message is displayed in the servermanager.log file:

```
"2014-08-14 21:01:35 +0000 | cassandra | cassandra not
started because it has been offline for longer than
its 15 day grace period - rebuild cassandra
```

   c. Make sure the timestamp of this message is the time when you attempted restarting Cassandra as instructed in step Restart Cassandra from the Storage Node.

      There can be more than one entry for Cassandra; you must locate the most recent entry.

d. If Cassandra has been down for longer than 15 days, you must rebuild the Cassandra database.

   For instructions, see Recover Storage Node down more than 15 days.

e. Contact technical support if alarms don't clear after Cassandra is rebuilt.

Cassandra Out of Memory errors (SMTT alarm)

A Total Events (SMTT) alarm is triggered when the Cassandra database has an out-of-memory error. If this error occurs, contact technical support to work through the issue.

About this task

If an out-of-memory error occurs for the Cassandra database, a heap dump is created, a Total Events (SMTT) alarm is triggered, and the Cassandra Heap Out Of Memory Errors count is incremented by one.

Steps

1. To view the event, select **SUPPORT** > **Tools** > **Grid topology** > **Configuration**.
2. Verify that the Cassandra Heap Out Of Memory Errors count is 1 or greater.

   You can run diagnostics to obtain additional information about the current state of your grid.

3. Go to `/var/local/core/`, compress the `Cassandra.hprof` file, and send it to technical support.
4. Make a backup of the `Cassandra.hprof` file, and delete it from the `/var/local/core/ directory`.

   This file can be as large as 24 GB, so you should remove it to free up space.

5. After the issue is resolved, select the **Reset** checkbox for the Cassandra Heap Out Of Memory Errors count. Then select **Apply Changes**.

   To reset event counts, you must have the Grid topology page configuration permission.

Troubleshoot certificate errors

If you see a security or certificate issue when you try to connect to StorageGRID using a web browser, an S3 or Swift client, or an external monitoring tool, you should check the certificate.

About this task

Certificate errors can cause problems when you try to connect to StorageGRID using the Grid Manager, Grid Management API, Tenant Manager, or the Tenant Management API. Certificate errors can also occur when you try to connect with an S3 or Swift client or external monitoring tool.

If you are accessing the Grid Manager or Tenant Manager using a domain name instead of an IP address, the browser shows a certificate error without an option to bypass if either of the following occurs:

- Your custom management interface certificate expires.
- You revert from a custom management interface certificate to the default server certificate.

The following example shows a certificate error when the custom management interface certificate expired:

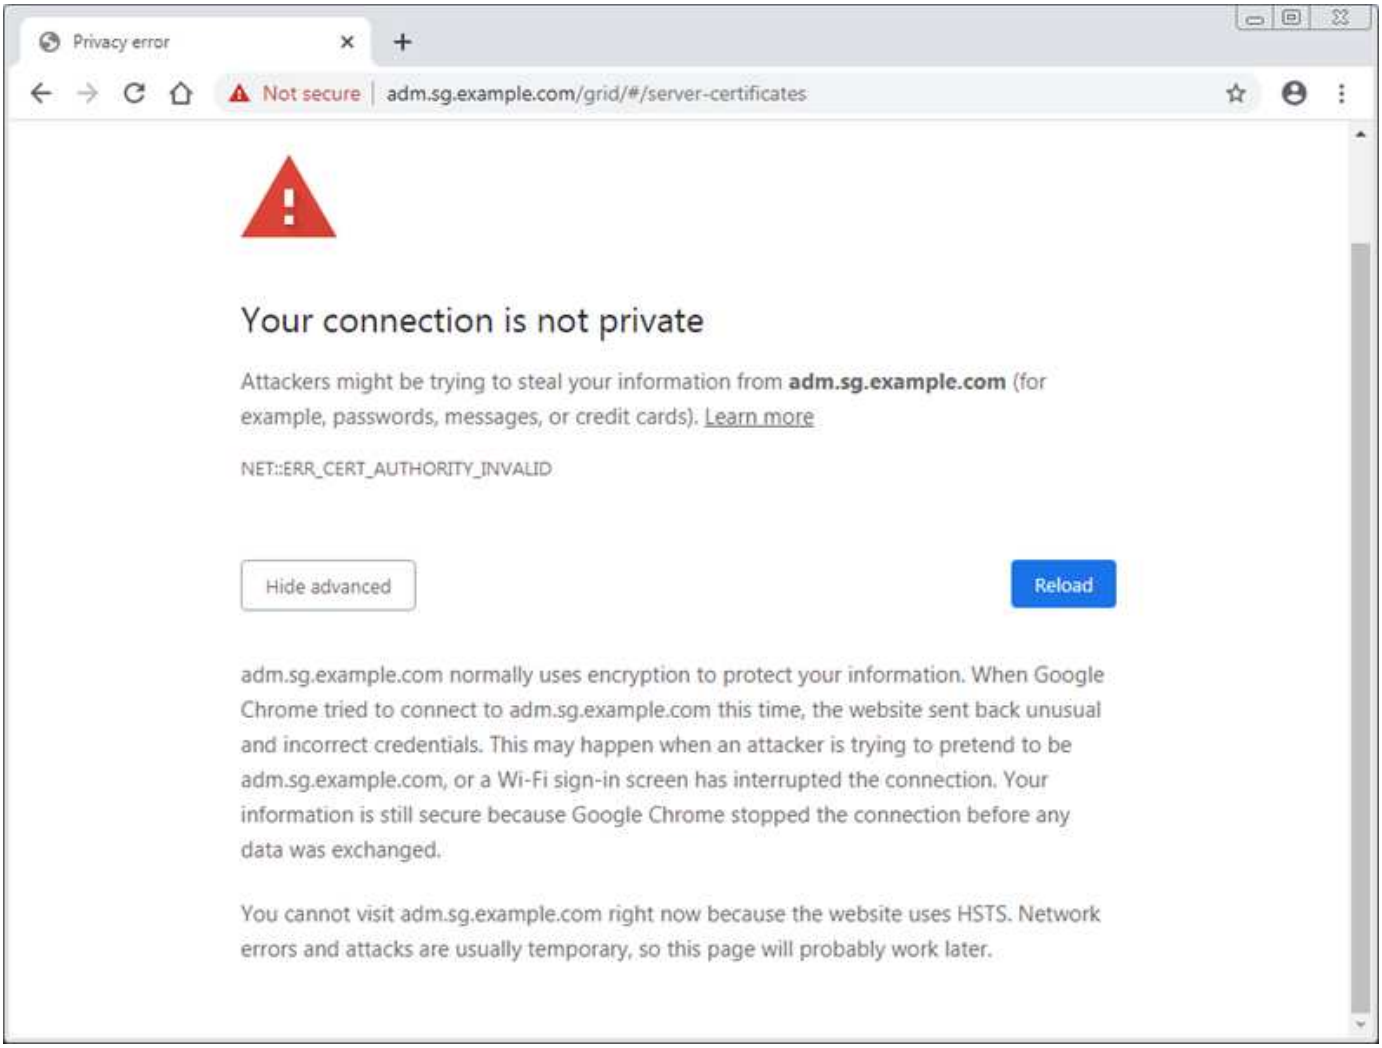

To ensure that operations aren't disrupted by a failed server certificate, the **Expiration of server certificate for Management Interface** alert is triggered when the server certificate is about to expire.

When you are using client certificates for external Prometheus integration, certificate errors can be caused by the StorageGRID management interface certificate or by client certificates. The **Expiration of client certificates configured on the Certificates page** alert is triggered when a client certificate is about to expire.

Steps

If you received an alert notification about an expired certificate, access the certificate details:
. Select **CONFIGURATION** > **Security** > **Certificates** and then select the appropriate certificate tab.

1. Check the validity period of the certificate.
   Some web browsers and S3 or Swift clients don't accept certificates with a validity period greater than 398 days.
2. If the certificate has expired or will expire soon, upload or generate a new certificate.
   - For a server certificate, see the steps for configuring a custom server certificate for the Grid Manager and the Tenant Manager.
   - For a client certificate, see the steps for configuring a client certificate.
3. For server certificate errors, try either or both of the following options:
   - Ensure that the Subject Alternative Name (SAN) of the certificate is populated, and that the SAN matches the IP address or host name of the node that you are connecting to.

- If you are attempting to connect to StorageGRID using a domain name:
  i. Enter the IP address of the Admin Node instead of the domain name to bypass the connection error and access the Grid Manager.
  ii. From the Grid Manager, select **CONFIGURATION** > **Security** > **Certificates** and then select the appropriate certificate tab to install a new custom certificate or continue with the default certificate.
  iii. In the instructions for administering StorageGRID, see the steps for configuring a custom server certificate for the Grid Manager and the Tenant Manager.

Troubleshoot Admin Node and user interface issues

There are several tasks you can perform to help determine the source of issues related to Admin Nodes and the StorageGRID user interface.

Sign-on errors

If you experience an error when you are signing in to a StorageGRID Admin Node, your system might have an issue with the identity federation configuration, a networking or hardware problem, an issue with Admin Node services, or an issue with the Cassandra database on connected Storage Nodes.

Before you begin

- You have the `Passwords.txt` file.
- You have specific access permissions.

About this task

Use these troubleshooting guidelines if you see any of the following error messages when attempting to sign in to an Admin Node:

- `Your credentials for this account were invalid. Please try again.`
- `Waiting for services to start…`
- `Internal server error. The server encountered an error and could not complete your request. Please try again. If the problem persists, contact Technical Support.`
- `Unable to communicate with server. Reloading page…`

Steps

1. Wait 10 minutes, and try signing in again.

   If the error is not resolved automatically, go to the next step.

2. If your StorageGRID system has more than one Admin Node, try signing in to the Grid Manager from another Admin Node.
   - If you are able to sign in, you can use the **Dashboard**, **NODES**, **Alerts**, and **SUPPORT** options to help determine the cause of the error.
   - If you have only one Admin Node or you still can't sign in, go to the next step.
3. Determine if the node's hardware is offline.
4. If single sign-on (SSO) is enabled for your StorageGRID system, refer to the steps for configuring single sign-on.

You might need to temporarily disable and re-enable SSO for a single Admin Node to resolve any issues.

> If SSO is enabled, you can't sign on using a restricted port. You must use port 443.

5. Determine if the account you are using belongs to a federated user.

   If the federated user account is not working, try signing in to the Grid Manager as a local user, such as root.

   - If the local user can sign in:
     i. Review any displayed alarms.
     ii. Select **CONFIGURATION** > **Access Control** > **Identity federation**.
     iii. Click **Test Connection** to validate your connection settings for the LDAP server.
     iv. If the test fails, resolve any configuration errors.
   - If the local user can't sign in and you are confident that the credentials are correct, go to the next step.

6. Use Secure Shell (ssh) to log in to the Admin Node:
   a. Enter the following command: `ssh admin@Admin_Node_IP`
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

7. View the status of all services running on the grid node: `storagegrid-status`

   Make sure the nms, mi, nginx, and mgmt api services are all running.

   The output is updated immediately if the status of a service changes.

```
$ storagegrid-status
Host Name                       99-211
IP Address                      10.96.99.211
Operating System Kernel         4.19.0         Verified
Operating System Environment    Debian 10.1    Verified
StorageGRID Webscale Release    11.4.0         Verified
Networking                                     Verified
Storage Subsystem                              Verified
Database Engine                 5.5.9999+default Running
Network Monitoring              11.4.0         Running
Time Synchronization            1:4.2.8p10+dfsg Running
ams                             11.4.0         Running
cmn                             11.4.0         Running
nms                             11.4.0         Running
ssm                             11.4.0         Running
mi                              11.4.0         Running
dynip                           11.4.0         Running
nginx                           1.10.3         Running
tomcat                          9.0.27         Running
grafana                         6.4.3          Running
mgmt api                        11.4.0         Running
prometheus                      11.4.0         Running
persistence                     11.4.0         Running
ade exporter                    11.4.0         Running
alertmanager                    11.4.0         Running
attrDownPurge                   11.4.0         Running
attrDownSamp1                   11.4.0         Running
attrDownSamp2                   11.4.0         Running
node exporter                   0.17.0+ds      Running
sg snmp agent                   11.4.0         Running
```

8. Confirm that the nginx-gw service is running `# service nginx-gw status`
9. Use Lumberjack to collect logs: `# /usr/local/sbin/lumberjack.rb`

   If the failed authentication happened in the past, you can use the --start and --end Lumberjack script options to specify the appropriate time range. Use lumberjack -h for details on these options.

   The output to the terminal indicates where the log archive has been copied.

10. Review the following logs:
    - `/var/local/log/bycast.log`
    - `/var/local/log/bycast-err.log`
    - `/var/local/log/nms.log`
    - `**/*commands.txt`

11. If you could not identify any issues with the Admin Node, issue either of the following commands to determine the IP addresses of the three Storage Nodes that run the ADC service at your site. Typically, these are the first three Storage Nodes that were installed at the site.

```
# cat /etc/hosts
```

```
# vi /var/local/gpt-data/specs/grid.xml
```

Admin Nodes use the ADC service during the authentication process.

12. From the Admin Node, log in to each of the ADC Storage Nodes, using the IP addresses you identified.
    a. Enter the following command: `ssh admin@grid_node_IP`
    b. Enter the password listed in the `Passwords.txt` file.
    c. Enter the following command to switch to root: `su -`
    d. Enter the password listed in the `Passwords.txt` file.

    When you are logged in as root, the prompt changes from `$` to `#`.

13. View the status of all services running on the grid node: `storagegrid-status`

    Make sure the idnt, acct, nginx, and cassandra services are all running.

14. Repeat steps Use Lumberjack to collect logs and Review logs to review the logs on the Storage Nodes.
15. If you are unable to resolve the issue, contact technical support.

    Provide the logs you collected to technical support. See also Log files reference.

User interface issues

The user interface for the Grid Manager or the Tenant Manager might not respond as expected after StorageGRID software is upgraded.

Steps

1. Make sure you are using a supported web browser.

   > Browser support can change with each StorageGRID release. Confirm you are using browser that is supported by your StorageGRID version.

2. Clear your web browser cache.

   Clearing the cache removes outdated resources used by the previous version of StorageGRID software, and permits the user interface to operate correctly again. For instructions, see the documentation for your web browser.

Unavailable Admin Node

If the StorageGRID system includes multiple Admin Nodes, you can use another Admin Node to check the status of an unavailable Admin Node.

Before you begin

You have specific access permissions.

Steps

1. From an available Admin Node, sign in to the Grid Manager using a supported web browser.
2. Select **SUPPORT** > **Tools** > **Grid topology**.
3. Select ***Site*** **> unavailable Admin Node** > **SSM** > **Services** > **Overview** > **Main**.
4. Look for services that have a status of Not Running and that might also be displayed in blue.

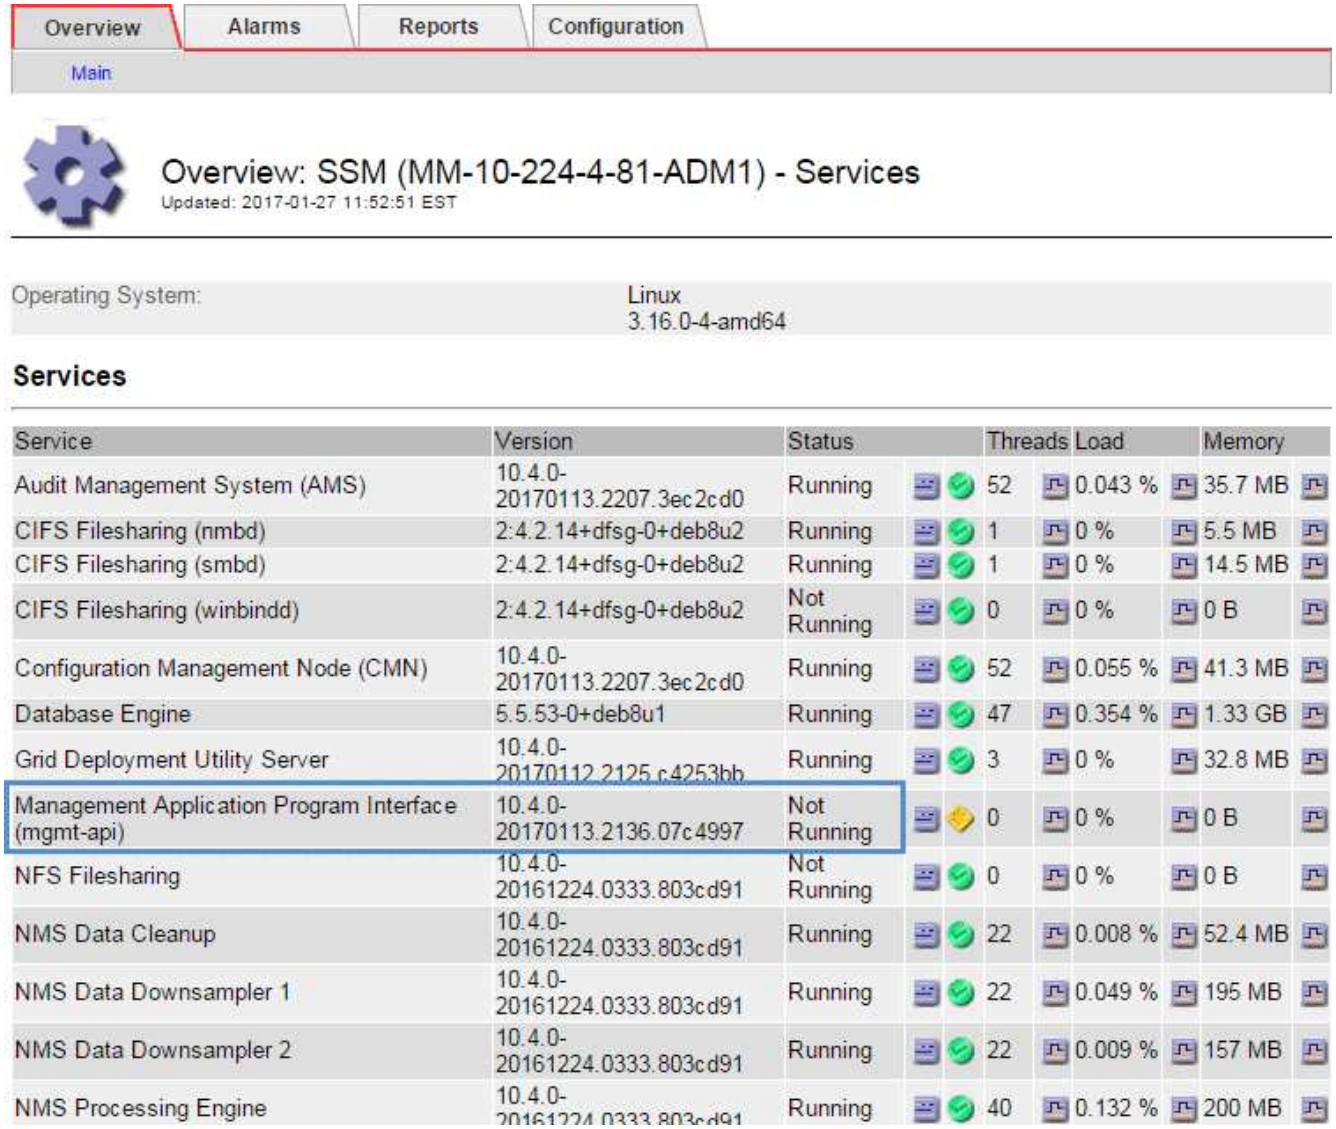

Service	Version	Status	Threads	Load	Memory
Audit Management System (AMS)	10.4.0-20170113.2207.3ec2cd0	Running	52	0.043 %	35.7 MB
CIFS Filesharing (nmbd)	2:4.2.14+dfsg-0+deb8u2	Running	1	0 %	5.5 MB
CIFS Filesharing (smbd)	2:4.2.14+dfsg-0+deb8u2	Running	1	0 %	14.5 MB
CIFS Filesharing (winbindd)	2:4.2.14+dfsg-0+deb8u2	Not Running	0	0 %	0 B
Configuration Management Node (CMN)	10.4.0-20170113.2207.3ec2cd0	Running	52	0.055 %	41.3 MB
Database Engine	5.5.53-0+deb8u1	Running	47	0.354 %	1.33 GB
Grid Deployment Utility Server	10.4.0-20170112.2125.c4253bb	Running	3	0 %	32.8 MB
Management Application Program Interface (mgmt-api)	10.4.0-20170113.2136.07c4997	Not Running	0	0 %	0 B
NFS Filesharing	10.4.0-20161224.0333.803cd91	Not Running	0	0 %	0 B
NMS Data Cleanup	10.4.0-20161224.0333.803cd91	Running	22	0.008 %	52.4 MB
NMS Data Downsampler 1	10.4.0-20161224.0333.803cd91	Running	22	0.049 %	195 MB
NMS Data Downsampler 2	10.4.0-20161224.0333.803cd91	Running	22	0.009 %	157 MB
NMS Processing Engine	10.4.0-20161224.0333.803cd91	Running	40	0.132 %	200 MB

5. Determine if alarms have been triggered.
6. Take the appropriate actions to resolve the issue.

Troubleshoot network, hardware, and platform issues

There are several tasks you can perform to help determine the source of issues related to StorageGRID network, hardware, and platform issues.

"422: Unprocessable Entity" errors

The error 422: Unprocessable Entity can occur for different reasons. Check the error message to determine what caused your issue.

If you see one of the listed error messages, take the recommended action.

Error message	Root cause and corrective action

`422: Unprocessable Entity`
 	
`Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration. Unable to authenticate, please verify your username and password: LDAP Result Code 8 "Strong Auth Required": 00002028: LdapErr: DSID-0C090256, comment: The server requires binds to turn on integrity checking if SSL\TLS are not already active on the connection, data 0, v3839`
	This message might occur if you select the **Do not use TLS** option for Transport Layer Security (TLS) when configuring identity federation using Windows Active Directory (AD).
	
	Using the **Do not use TLS** option is not supported for use with AD servers that enforce LDAP signing. You must select either the **Use STARTTLS** option or the **Use LDAPS** option for TLS.

Error message	Root cause and corrective action

`422: Unprocessable Entity`
 	
`Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration.Unable to begin TLS, verify your certificate and TLS configuration: LDAP Result Code 200 "Network Error": TLS handshake failed (EOF)`
	This message appears if you try to use an unsupported cipher to make a Transport Layer Security (TLS) connection from StorageGRID to an external system used for identify federation or Cloud Storage Pools.
	
	Check the ciphers that are offered by the external system. The system must use one of the ciphers supported by StorageGRID for outgoing TLS connections, as shown in the instructions for administering StorageGRID.

Grid Network MTU mismatch alert

The **Grid Network MTU mismatch** alert is triggered when the maximum transmission unit (MTU) setting for the Grid Network interface (eth0) differs significantly across nodes in the grid.

About this task

The differences in MTU settings could indicate that some, but not all, eth0 networks are configured for jumbo frames. An MTU size mismatch of greater than 1000 might cause network performance problems.

Steps

1. List the MTU settings for eth0 on all nodes.
   - Use the query provided in the Grid Manager.
   - Navigate to *`primary Admin Node IP address`*`/metrics/graph` and enter the following query: `node_network_mtu_bytes{device="eth0"}`
2. Modify the MTU settings as necessary to ensure they are the same for the Grid Network interface (eth0) on all nodes.
   - For Linux- and VMware-based nodes, use the following command: `/usr/sbin/change-ip.py [-h] [-n node] mtu network [network...]`

     Example: `change-ip.py -n node 1500 grid admin`

     Note: On Linux-based nodes, if the desired MTU value for the network in the container exceeds the value already configured on the host interface, you must first configure the host interface to have the desired MTU value, and then use the `change-ip.py` script to change the MTU value of the network in the container.

     Use the following arguments for modifying the MTU on Linux- or VMware-based nodes.

Positional arguments	Description
`mtu`	The MTU to set. Must be in the range 1280 to 9216.
`network`	The networks to apply the MTU to. Include one or more of the following network types: • grid • admin • client

Optional arguments	Description
`-h, - help`	Show the help message and exit.
`-n node, --node node`	The node. The default is the local node.

Network Receive Error (NRER) alarm

Network Receive Error (NRER) alarms can be caused by connectivity issues between StorageGRID and your network hardware. In some cases, NRER errors can clear without manual intervention. If the errors don't clear, take the recommended actions.

About this task

NRER alarms can be caused by the following issues with networking hardware that connects to StorageGRID:

- Forward error correction (FEC) is required and not in use
- Switch port and NIC MTU mismatch
- High link error rates
- NIC ring buffer overrun

Steps

1. Follow the troubleshooting steps for all potential causes of the NRER alarm given your network configuration.
2. Perform the following steps depending on the cause of the error:

FEC mismatch

> These steps are applicable only for NRER errors caused by FEC mismatch on StorageGRID appliances.

a. Check the FEC status of the port in the switch attached to your StorageGRID appliance.
b. Check the physical integrity of the cables from the appliance to the switch.
c. If you want to change FEC settings to try to resolve the NRER alarm, first ensure that the appliance is configured for **Auto** mode on the Link Configuration page of the StorageGRID Appliance Installer (see the instructions for your appliance:
   - SGF6112
   - SG6000
   - SG5700
   - SG110 and SG1100
   - SG100 and SG1000
d. Change the FEC settings on the switch ports. The StorageGRID appliance ports will adjust their FEC settings to match, if possible.

   You can't configure FEC settings on StorageGRID appliances. Instead, the appliances attempt to discover and mirror the FEC settings on the switch ports they are connected to. If the links are forced to 25-GbE or 100-GbE network speeds, the switch and NIC might fail to negotiate a common FEC setting. Without a common FEC setting, the network will fall back to "no-FEC" mode. When FEC is not enabled, the connections are more susceptible to errors caused by electrical noise.

   > StorageGRID appliances support Firecode (FC) and Reed Solomon (RS) FEC, as well as no FEC.

Switch port and NIC MTU mismatch

If the error is caused by a switch port and NIC MTU mismatch, check that the MTU size configured on the node is the same as the MTU setting for the switch port.

The MTU size configured on the node might be smaller than the setting on the switch port the node is connected to. If a StorageGRID node receives an Ethernet frame larger than its MTU, which is possible with this configuration, the NRER alarm might be reported. If you believe this is what is happening, either change the MTU of the switch port to match the StorageGRID network interface MTU, or change the MTU of the StorageGRID network interface to match the switch port, depending on your end-to-end MTU goals or requirements.

> For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The **Grid Network MTU mismatch** alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values don't have to be the same for all network types. See Troubleshoot the Grid Network MTU mismatch alert for more information.

> Also see Change MTU setting.

High link error rates

a. Enable FEC, if not already enabled.
b. Verify that your network cabling is of good quality and is not damaged or improperly connected.
c. If the cables don't appear to be the problem, contact technical support.

> You might notice high error rates in an environment with high electrical noise.

NIC ring buffer overrun

If the error is a NIC ring buffer overrun, contact technical support.

The ring buffer can be overrun when the StorageGRID system is overloaded and unable to process network events in a timely manner.

3. After you resolve the underlying problem, reset the error counter.
   a. Select **SUPPORT** > **Tools** > **Grid topology**.
   b. Select ***site*** > ***grid node*** > **SSM** > **Resources** > **Configuration** > **Main**.
   c. Select **Reset Receive Error Count** and click **Apply Changes**.

Related information

Alarms reference (legacy system)

Time synchronization errors

You might see issues with time synchronization in your grid.

If you encounter time synchronization problems, verify that you have specified at least four external NTP sources, each providing a Stratum 3 or better reference, and that all external NTP sources are operating normally and are accessible by your StorageGRID nodes.

> When specifying the external NTP source for a production-level StorageGRID installation, don't use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

Linux: Network connectivity issues

You might see issues with network connectivity for StorageGRID nodes hosted on Linux hosts.

MAC address cloning

In some cases, network issues can be resolved by using MAC address cloning. If you are using virtual hosts, set the value of the MAC address cloning key for each of your networks to "true" in your node configuration file. This setting causes the MAC address of the StorageGRID container to use the MAC address of the host. To create node configuration files, see the instructions for Red Hat Enterprise Linux or Ubuntu or Debian.

> Create separate virtual network interfaces for use by the Linux host OS. Using the same network interfaces for the Linux host OS and the StorageGRID container might cause the host OS to become unreachable if promiscuous mode has not been enabled on the hypervisor.

For more information about enabling MAC cloning, see the instructions for Red Hat Enterprise Linux or Ubuntu

or Debian.

Promiscuous mode

If you don't want to use MAC address cloning and would rather allow all interfaces to receive and transmit data for MAC addresses other than the ones assigned by the hypervisor, ensure that the security properties at the virtual switch and port group levels are set to **Accept** for Promiscuous Mode, MAC Address Changes, and Forged Transmits. The values set on the virtual switch can be overridden by the values at the port group level, so ensure that settings are the same in both places.

For more information about using Promiscuous Mode, see the instructions for Red Hat Enterprise Linux or Ubuntu or Debian.

Linux: Node status is "orphaned"

A Linux node in an orphaned state usually indicates that either the storagegrid service or the StorageGRID node daemon controlling the node's container died unexpectedly.

About this task

If a Linux node reports that it is in an orphaned state, you should:

- Check logs for errors and messages.
- Attempt to start the node again.
- If necessary, use container engine commands to stop the existing node container.
- Restart the node.

Steps

1. Check logs for both the service daemon and the orphaned node for obvious errors or messages about exiting unexpectedly.
2. Log in to the host as root or using an account with sudo permission.
3. Attempt to start the node again by running the following command: `$ sudo storagegrid node start node-name`

   ```
   $ sudo storagegrid node start DC1-S1-172-16-1-172
   ```

   If the node is orphaned, the response is

   ```
   Not starting ORPHANED node DC1-S1-172-16-1-172
   ```

4. From Linux, stop the container engine and any controlling storagegrid-node processes. For example:`sudo docker stop --time secondscontainer-name`

   For `seconds`, enter the number of seconds you want to wait for the container to stop (typically 15 minutes or less). For example:

   ```
   sudo docker stop --time 900 storagegrid-DC1-S1-172-16-1-172
   ```

5. Restart the node: `storagegrid node start node-name`

```
storagegrid node start DC1-S1-172-16-1-172
```

Linux: Troubleshoot IPv6 support

You might need to enable IPv6 support in the kernel if you have installed StorageGRID nodes on Linux hosts and you notice that IPv6 addresses have not been assigned to the node containers as expected.

About this task

You can see the IPv6 address that has been assigned to a grid node in the following locations in the Grid Manager:

- Select **NODES**, and select the node. Then, select **Show more** next to **IP Addresses** on the Overview tab.

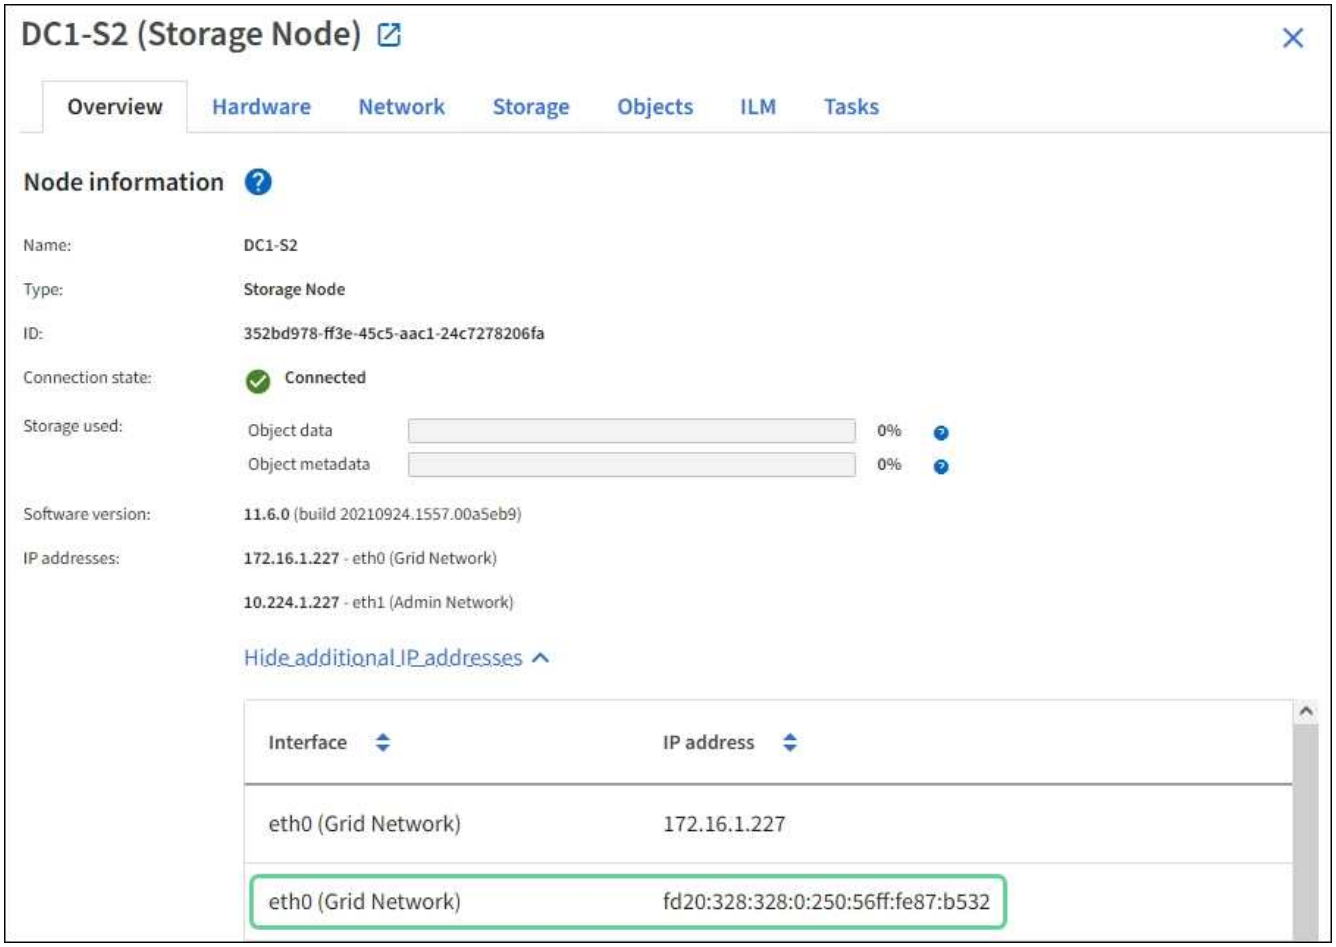

- Select **SUPPORT** > **Tools** > **Grid topology**. Then, select ***node*** > **SSM** > **Resources**. If an IPv6 address has been assigned, it is listed below the IPv4 address in the **Network Addresses** section.

If the IPv6 address is not shown and the node is installed on a Linux host, follow these steps to enable IPv6 support in the kernel.

Steps

1. Log in to the host as root or using an account with sudo permission.

2. Run the following command: `sysctl net.ipv6.conf.all.disable_ipv6`

```
root@SG:~ # sysctl net.ipv6.conf.all.disable_ipv6
```

The result should be 0.

```
net.ipv6.conf.all.disable_ipv6 = 0
```

> If the result is not 0, see the documentation for your operating system for changing `sysctl` settings. Then, change the value to 0 before continuing.

3. Enter the StorageGRID node container: `storagegrid node enter node-name`
4. Run the following command: `sysctl net.ipv6.conf.all.disable_ipv6`

```
root@DC1-S1:~ # sysctl net.ipv6.conf.all.disable_ipv6
```

The result should be 1.

```
net.ipv6.conf.all.disable_ipv6 = 1
```

> If the result is not 1, this procedure does not apply. Contact technical support.

5. Exit the container: `exit`

```
root@DC1-S1:~ # exit
```

6. As root, edit the following file: `/var/lib/storagegrid/settings/sysctl.d/net.conf`.

```
sudo vi /var/lib/storagegrid/settings/sysctl.d/net.conf
```

7. Locate the following two lines and remove the comment tags. Then, save and close the file.

```
net.ipv6.conf.all.disable_ipv6 = 0
```

```
net.ipv6.conf.default.disable_ipv6 = 0
```

8. Run these commands to restart the StorageGRID container:

```
storagegrid node stop node-name
```

```
storagegrid node start node-name
```

Troubleshoot an external syslog server

The following table describes the error messages that might be related using to an external syslog server and lists corrective actions.

For more information about sending audit information to an external syslog server, see:

- Considerations for using an external syslog server
- Configure audit messages and external syslog server

Error message	Description and recommended actions
Cannot resolve hostname	The FQDN you entered for the syslog server could not be resolved to an IP address. 1. Check the hostname you entered. If you entered an IP address, make sure it is a valid IP address in W.X.Y.Z ("dotted decimal") notation. 2. Check that the DNS servers are configured correctly. 3. Confirm that each node can access the IP addresses for the DNS server.
Connection refused	A TCP or TLS connection to the syslog server was refused. There might be no service listening on the TCP or TLS port for the host, or a firewall might be blocking access. 1. Check that you entered the correct FQDN or IP address, port, and protocol for the syslog server. 2. Confirm that the host for the syslog service is running a syslog daemon that is listening on the specified port. 3. Confirm that a firewall is not blocking access to TCP/TLS connections from the nodes to the IP and port of the syslog server.
Network unreachable	The syslog server is not on a directly attached subnet. A router returned an ICMP failure message to indicate it could not forward the test messages from the listed nodes to the syslog server. 1. Check that you entered the correct FQDN or IP address for the syslog server. 2. For each node listed, check the Grid Network Subnet List, the Admin Networks Subnet Lists, and the Client Network gateways. Confirm these are configured to route traffic to the syslog server over the expected network interface and gateway (Grid, Admin, or Client).

Error message	Description and recommended actions
Host unreachable	The syslog server is on a directly attached subnet (subnet used by the listed nodes for their Grid, Admin, or Client IP addresses). The nodes attempted to send test messages, but did not receive responses to ARP requests for the syslog server's MAC address. 1. Check that you entered the correct FQDN or IP address for the syslog server. 2. Check that the host running the syslog service is up.
Connection timed out	A TCP/TLS connection attempt was made, but no response was received from the syslog server for a long time. There might be a routing misconfiguration or a firewall might be dropping traffic without sending any response (a common configuration). 1. Check that you entered the correct FQDN or IP address for the syslog server. 2. For each node listed, check the Grid Network Subnet List, the Admin Networks Subnet Lists, and the Client Network gateways. Confirm these are configured to route traffic to the syslog server using the network interface and gateway (Grid, Admin, or Client) over which you expect the syslog server to be reached. 3. Confirm that a firewall is not blocking access to TCP/TLS connections from the nodes listed to the IP and port of the syslog server.
Connection closed by partner	A TCP connection to the syslog server was successfully established but was later closed. Reasons for this might include: • The syslog server might have been restarted or rebooted. • The node and the syslog server might have different TCP/TLS settings. • An intermediate firewall might be closing idle TCP connections. • A non-syslog server listening on the syslog server port might have closed the connection. To resolve this issue: 1. Check that you entered the correct FQDN or IP address, port, and protocol for the syslog server. 2. If you are using TLS, confirm the syslog server is also using TLS. If you are using TCP, confirm the syslog server is also using TCP. 3. Check that an intermediate firewall is not configured to close idle TCP connections.

Error message	Description and recommended actions
TLS certificate error	The server certificate received from the syslog server was not compatible with the CA certificate bundle and client certificate you provided. 1. Confirm that the CA certificate bundle and client certificate (if any) are compatible with the server certificate on the syslog server. 2. Confirm that the identities in the server certificate from the syslog server include the expected IP or FQDN values.
Forwarding suspended	Syslog records are no longer being forwarded to the syslog server and StorageGRID is unable to detect the reason. Review the debugging logs provided with this error to attempt to determine the root cause.
TLS session terminated	The syslog server terminated the TLS session and StorageGRID is unable to detect the reason. 1. Review the debugging logs provided with this error to attempt to determine the root cause. 2. Check that you entered the correct FQDN or IP address, port, and protocol for the syslog server. 3. If you are using TLS, confirm the syslog server is also using TLS. If you are using TCP, confirm the syslog server is also using TCP. 4. Confirm that the CA certificate bundle and client certificate (if any) are compatible with the server certificate from the syslog server. 5. Confirm that the identities in the server certificate from the syslog server include the expected IP or FQDN values.
Results query failed	The Admin Node used for syslog server configuration and testing is unable to request test results from the nodes listed. One or more nodes might be down. 1. Follow standard troubleshooting steps to ensure that the nodes are online and all expected services are running. 2. Restart the miscd service on the nodes listed.

Review audit logs

Review audit logs: Overview

These instructions contain information about the structure and content of StorageGRID audit messages and audit logs. You can use this information to read and analyze the audit trail of system activity.

These instructions are for administrators responsible for producing reports of system activity and usage that require analysis of the StorageGRID system's audit messages.

To use the text log file, you must have access to the configured audit share on the Admin Node.

For information about configuring audit message levels and using an external syslog server, see Configure audit messages and log destinations.

Audit message flow and retention

All StorageGRID services generate audit messages during normal system operation. You should understand how these audit messages move through the StorageGRID system to the `audit.log` file.

Audit message flow

Audit messages are processed by Admin Nodes and by those Storage Nodes that have an Administrative Domain Controller (ADC) service.

As shown in the audit message flow diagram, each StorageGRID node sends its audit messages to one of the ADC services at the data center site. The ADC service is automatically enabled for the first three Storage Nodes installed at each site.

In turn, each ADC service acts as a relay and sends its collection of audit messages to every Admin Node in the StorageGRID system, which gives each Admin Node a complete record of system activity.

Each Admin Node stores audit messages in text log files; the active log file is named `audit.log`.

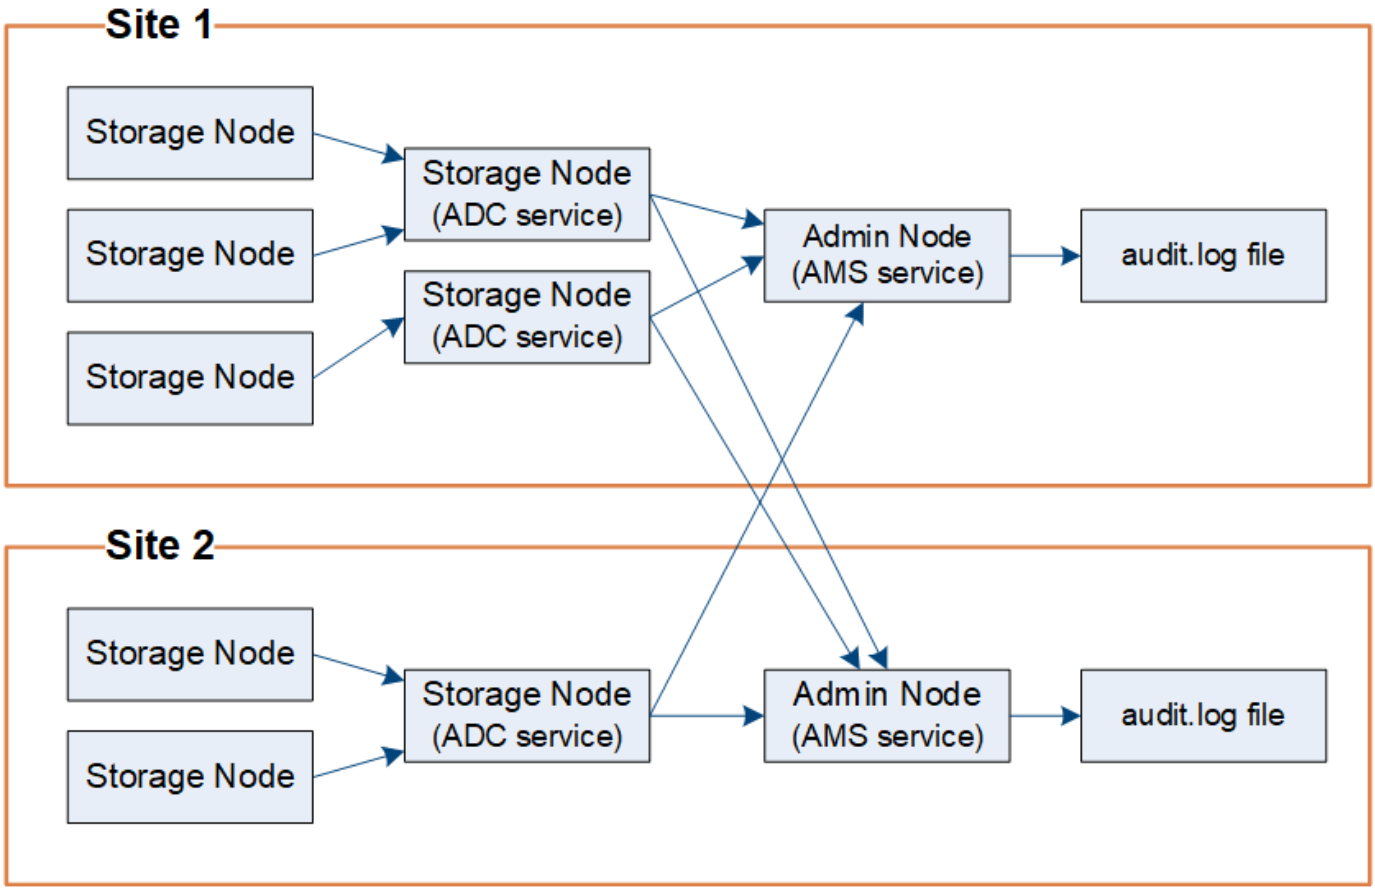

Audit message retention

StorageGRID uses a copy-and-delete process to ensure that no audit messages are lost before they can be written to the audit log.

When a node generates or relays an audit message, the message is stored in an audit message queue on the system disk of the grid node. A copy of the message is always held in an audit message queue until the message is written to the audit log file in the Admin Node's `/var/local/log` directory. This helps prevent loss of an audit message during transport.

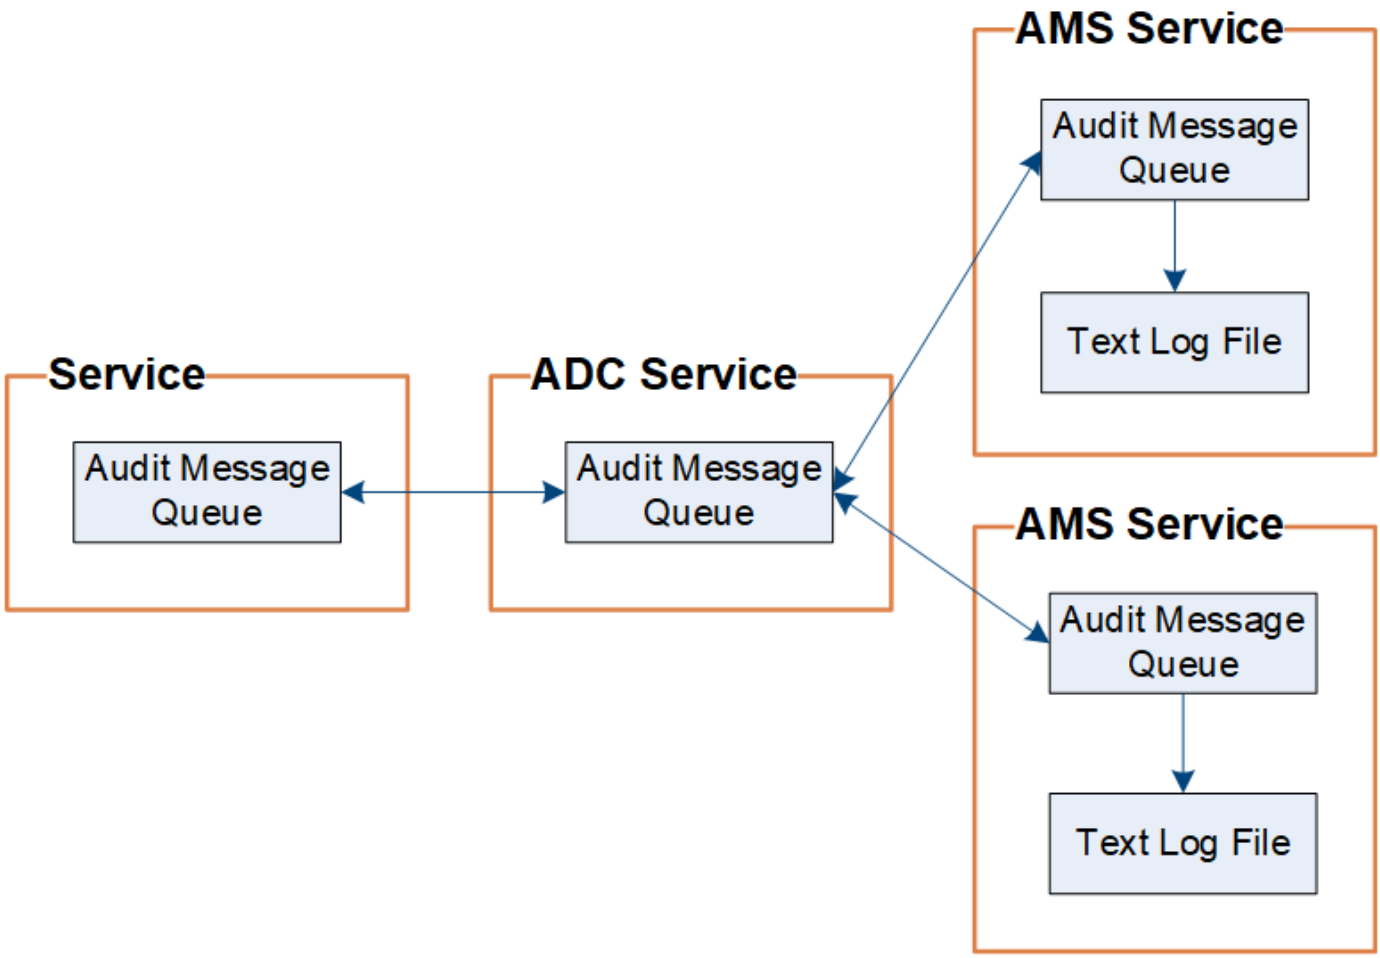

The audit message queue can temporarily increase due to network connectivity issues or insufficient audit capacity. As the queues increase, they consume more of the available space in each node's `/var/local/` directory. If the issue persists and a node's audit message directory becomes too full, the individual nodes will prioritize processing their backlog and become temporarily unavailable for new messages.

Specifically, you might see the following behaviors:

- If the `/var/local/log` directory used by an Admin Node becomes full, the Admin Node will be flagged as unavailable to new audit messages until the directory is no longer full. S3 and Swift client requests aren't affected. The XAMS (Unreachable Audit Repositories) alarm is triggered when an audit repository is unreachable.
- If the `/var/local/` directory used by a Storage Node with the ADC service becomes 92% full, the node will be flagged as unavailable to audit messages until the directory is only 87% full. S3 and Swift client requests to other nodes aren't affected. The NRLY (Available Audit Relays) alarm is triggered when audit relays are unreachable.

> If there are no available Storage Nodes with the ADC service, the Storage Nodes store the audit messages locally in the `/var/local/log/localaudit.log` file.

- If the `/var/local/` directory used by a Storage Node becomes 85% full, the node will start refusing S3 and Swift client requests with `503 Service Unavailable`.

The following types of issues can cause audit message queues to grow very large:

- The outage of an Admin Node or a Storage Node with the ADC service. If one of the system's nodes is down, the remaining nodes might become backlogged.
- A sustained activity rate that exceeds the audit capacity of the system.
- The `/var/local/` space on an ADC Storage Node becoming full for reasons unrelated to audit messages. When this happens, the node stops accepting new audit messages and prioritizes its current backlog, which can cause backlogs on other nodes.

Large audit queue alert and Audit Messages Queued (AMQS) alarm

To help you monitor the size of audit message queues over time, the **Large audit queue** alert and the legacy AMQS alarm are triggered when the number of messages in a Storage Node queue or Admin Node queue reaches certain thresholds.

If the **Large audit queue** alert or the legacy AMQS alarm is triggered, start by checking the load on the system—if there have been a significant number of recent transactions, the alert and the alarm should resolve over time and can be ignored.

If the alert or alarm persists and increases in severity, view a chart of the queue size. If the number is steadily increasing over hours or days, the audit load has likely exceeded the audit capacity of the system. Reduce the client operation rate or decrease the number of audit messages logged by changing the audit level for Client Writes and Client Reads to Error or Off. See Configure audit messages and log destinations.

Duplicate messages

The StorageGRID system takes a conservative approach if a network or node failure occurs. For this reason, duplicate messages might exist in the audit log.

Access audit log file

The audit share contains the active `audit.log` file and any compressed audit log files. You can access audit log files directly from the command line of the Admin Node.

Before you begin

- You have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of an Admin Node.

Steps

1. Log in to an Admin Node:
   a. Enter the following command: `ssh admin@`*`primary_Admin_Node_IP`*
   b. Enter the password listed in the `Passwords.txt` file.

c. Enter the following command to switch to root: `su -`

d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

2. Go to the directory containing the audit log files:

```
cd /var/local/log
```

3. View the current or a saved audit log file, as required.

Audit log file rotation

Audit logs files are saved to an Admin Node's `/var/local/log` directory. The active audit log files are named `audit.log`.

> Optionally, you can change the destination of audit logs and send audit information to an external syslog server. Local logs of audit records continue to be generated and stored when an external syslog server is configured. See Configure audit messages and log destinations.

Once a day, the active `audit.log` file is saved, and a new `audit.log` file is started. The name of the saved file indicates when it was saved, in the format *`yyyy-mm-dd.txt`*. If more than one audit log is created in a single day, the file names use the date the file was saved, appended by a number, in the format *`yyyy-mm-dd.txt.n`*. For example, `2018-04-15.txt` and `2018-04-15.txt.1` are the first and second log files created and saved on 15 April 2018.

After a day, the saved file is compressed and renamed, in the format *`yyyy-mm-dd.txt.gz`*, which preserves the original date. Over time, this results in the consumption of storage allocated for audit logs on the Admin Node. A script monitors the audit log space consumption and deletes log files as necessary to free space in the `/var/local/log` directory. Audit logs are deleted based on the date they were created, with the oldest being deleted first. You can monitor the script's actions in the following file: `/var/local/log/manage-audit.log`.

This example shows the active `audit.log` file, the previous day's file (`2018-04-15.txt`), and the compressed file for the prior day (`2018-04-14.txt.gz`).

```
audit.log
2018-04-15.txt
2018-04-14.txt.gz
```

Audit log file format

Audit log file format: Overview

The audit log files are found on every Admin Node and contain a collection of individual audit messages.

Each audit message contains the following:

- The Coordinated Universal Time (UTC) of the event that triggered the audit message (ATIM) in ISO 8601 format, followed by a space:

  YYYY-MM-DDTHH:MM:SS.UUUUUU, where *UUUUUU* are microseconds.

- The audit message itself, enclosed within square brackets and beginning with `AUDT`.

The following example shows three audit messages in an audit log file (line breaks added for readability). These messages were generated when a tenant created an S3 bucket and added two objects to that bucket.

```
2019-08-07T18:43:30.247711
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991681][TIME(UI64):73520][SAI
P(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCkNB8M3MTWNt-
PhoTDwB9JOk7PtyLkQmA=="][SUSR(CSTR):"urn:sgws:identity::175300642415970547
18:root"]
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"buc
ket1"][AVER(UI32):10][ATIM(UI64):1565203410247711]
[ATYP(FC32):SPUT][ANID(UI32):12454421][AMID(FC32):S3RQ][ATID(UI64):7074142
142472611085]]

2019-08-07T18:43:30.783597
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991696][TIME(UI64):120713][SA
IP(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCkNB8M3MTWNt-
PhoTDwB9JOk7PtyLkQmA=="][SUSR(CSTR):"urn:sgws:identity::175300642415970547
18:root"]
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"buc
ket1"][S3KY(CSTR):"fh-small-0"]
[CBID(UI64):0x779557A069B2C037][UUID(CSTR):"94BA6949-38E1-4B0C-BC80-
EB44FB4FCC7F"][CSIZ(UI64):1024][AVER(UI32):10]
[ATIM(UI64):1565203410783597][ATYP(FC32):SPUT][ANID(UI32):12454421][AMID(F
C32):S3RQ][ATID(UI64):8439606722108456022]]

2019-08-07T18:43:30.784558
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991693][TIME(UI64):121666][SA
IP(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCkNB8M3MTWNt-
PhoTDwB9JOk7PtyLkQmA=="][SUSR(CSTR):"urn:sgws:identity::175300642415970547
18:root"]
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"buc
ket1"][S3KY(CSTR):"fh-small-2000"]
[CBID(UI64):0x180CBD8E678EED17][UUID(CSTR):"19CE06D0-D2CF-4B03-9C38-
E578D66F7ADD"][CSIZ(UI64):1024][AVER(UI32):10]
[ATIM(UI64):1565203410784558][ATYP(FC32):SPUT][ANID(UI32):12454421][AMID(F
C32):S3RQ][ATID(UI64):13489590586043706682]]
```

In their default format, the audit messages in the audit log files aren't easy to read or interpret. You can use the audit-explain tool to obtain simplified summaries of the audit messages in the audit log. You can use the audit-sum tool to summarize how many write, read, and delete operations were logged and how long these operations took.

Use audit-explain tool

You can use the `audit-explain` tool to translate the audit messages in the audit log in to an easy-to-read format.

Before you begin

- You have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of the primary Admin Node.

About this task

The `audit-explain` tool, available on the primary Admin Node, provides simplified summaries of the audit messages in an audit log.

> The `audit-explain` tool is primarily intended for use by technical support during troubleshooting operations. Processing `audit-explain` queries can consume a large amount of CPU power, which might impact StorageGRID operations.

This example shows typical output from the `audit-explain` tool. These four SPUT audit messages were generated when the S3 tenant with account ID 92484777680322627870 used S3 PUT requests to create a bucket named "bucket1" and add three objects to that bucket.

```
SPUT S3 PUT bucket bucket1 account:92484777680322627870 usec:124673
SPUT S3 PUT object bucket1/part1.txt tenant:92484777680322627870
cbid:9DCB157394F99FE5 usec:101485
SPUT S3 PUT object bucket1/part2.txt tenant:92484777680322627870
cbid:3CFBB07AB3D32CA9 usec:102804
SPUT S3 PUT object bucket1/part3.txt tenant:92484777680322627870
cbid:5373D73831ECC743 usec:93874
```

The `audit-explain` tool can do the following:

- Process plain or compressed audit logs. For example:

  ```
  audit-explain audit.log
  ```

  ```
  audit-explain 2019-08-12.txt.gz
  ```

- Process multiple files simultaneously. For example:

  ```
  audit-explain audit.log 2019-08-12.txt.gz 2019-08-13.txt.gz
  ```

  ```
  audit-explain /var/local/log/*
  ```

- Accept input from a pipe, which allows you to filter and preprocess the input using the `grep` command or other means. For example:

  ```
  grep SPUT audit.log | audit-explain

  grep bucket-name audit.log | audit-explain
  ```

Because audit logs can be very large and slow to parse, you can save time by filtering parts that you want to look at and running `audit-explain` on the parts, instead of the entire file.

> The `audit-explain` tool does not accept compressed files as piped input. To process compressed files, provide their file names as command-line arguments, or use the `zcat` tool to decompress the files first. For example:
>
> ```
> zcat audit.log.gz | audit-explain
> ```

Use the `help (-h)` option to see the available options. For example:

```
$ audit-explain -h
```

Steps

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@`*`primary_Admin_Node_IP`*
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.

2. Enter the following command, where `/var/local/log/audit.log` represents the name and the location of the file or files you want to analyze:

   ```
   $ audit-explain /var/local/log/audit.log
   ```

   The `audit-explain` tool prints human-readable interpretations of all messages in the specified file or files.

   > To reduce line lengths and to aid readability, timestamps aren't shown by default. If you want to see the timestamps, use the timestamp (`-t`) option.

Use audit-sum tool

You can use the `audit-sum` tool to count the write, read, head, and delete audit messages and to see the minimum, maximum, and average time (or size) for each operation type.

Before you begin

- You have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of the primary Admin Node.

About this task

The `audit-sum` tool, available on the primary Admin Node, summarizes how many write, read, and delete operations were logged and how long these operations took.

> The `audit-sum` tool is primarily intended for use by technical support during troubleshooting operations. Processing `audit-sum` queries can consume a large amount of CPU power, which might impact StorageGRID operations.

This example shows typical output from the `audit-sum` tool. This example shows how long protocol operations took.

```
  message group           count     min(sec)        max(sec)
average(sec)
  =============           =====     ========        ========
============
  IDEL                      274
  SDEL                   213371        0.004          20.934
0.352
  SGET                   201906        0.010        1740.290
1.132
  SHEA                    22716        0.005           2.349
0.272
  SPUT                  1771398        0.011        1770.563
0.487
```

The `audit-sum` tool provides counts and times for the following S3, Swift, and ILM audit messages in an audit log:

Code	Description	Refer to
ARCT	Archive Retrieve from Cloud-Tier	ARCT: Archive Retrieve from Cloud-Tier
ASCT	Archive Store Cloud-Tier	ASCT: Archive Store Cloud-Tier
IDEL	ILM Initiated Delete: Logs when ILM starts the process of deleting an object.	IDEL: ILM Initiated Delete
SDEL	S3 DELETE: Logs a successful transaction to delete an object or bucket.	SDEL: S3 DELETE

Code	Description	Refer to
SGET	S3 GET: Logs a successful transaction to retrieve an object or list the objects in a bucket.	SGET: S3 GET
SHEA	S3 HEAD: Logs a successful transaction to check for the existence of an object or bucket.	SHEA: S3 HEAD
SPUT	S3 PUT: Logs a successful transaction to create a new object or bucket.	SPUT: S3 PUT
WDEL	Swift DELETE: Logs a successful transaction to delete an object or container.	WDEL: Swift DELETE
WGET	Swift GET: Logs a successful transaction to retrieve an object or list the objects in a container.	WGET: Swift GET
WHEA	Swift HEAD: Logs a successful transaction to check for the existence of an object or container.	WHEA: Swift HEAD
WPUT	Swift PUT: Logs a successful transaction to create a new object or container.	WPUT: Swift PUT

The `audit-sum` tool can do the following:

- Process plain or compressed audit logs. For example:

  ```
  audit-sum audit.log
  ```

  ```
  audit-sum 2019-08-12.txt.gz
  ```

- Process multiple files simultaneously. For example:

  ```
  audit-sum audit.log 2019-08-12.txt.gz 2019-08-13.txt.gz
  ```

  ```
  audit-sum /var/local/log/*
  ```

- Accept input from a pipe, which allows you to filter and preprocess the input using the `grep` command or other means. For example:

  ```
  grep WGET audit.log | audit-sum
  ```

  ```
  grep bucket1 audit.log | audit-sum
  ```

  ```
  grep SPUT audit.log | grep bucket1 | audit-sum
  ```

> This tool does not accept compressed files as piped input. To process compressed files, provide their file names as command-line arguments, or use the `zcat` tool to decompress the files first. For example:
>
> ```
> audit-sum audit.log.gz
>
> zcat audit.log.gz | audit-sum
> ```

You can use command-line options to summarize operations on buckets separately from operations on objects or to group message summaries by bucket name, by time period, or by target type. By default, the summaries show the minimum, maximum, and average operation time, but you can use the `size (-s)` option to look at object size instead.

Use the `help (-h)` option to see the available options. For example:

```
$ audit-sum -h
```

Steps

1. Log in to the primary Admin Node:
   a. Enter the following command: `ssh admin@`*`primary_Admin_Node_IP`*
   b. Enter the password listed in the `Passwords.txt` file.
   c. Enter the following command to switch to root: `su -`
   d. Enter the password listed in the `Passwords.txt` file.

      When you are logged in as root, the prompt changes from `$` to `#`.
2. If you want to analyze all messages related to write, read, head, and delete operations, follow these steps:
   a. Enter the following command, where `/var/local/log/audit.log` represents the name and the location of the file or files you want to analyze:

      ```
      $ audit-sum /var/local/log/audit.log
      ```

      This example shows typical output from the `audit-sum` tool. This example shows how long protocol operations took.

```
  message group           count     min(sec)        max(sec)
average(sec)
  =============           =====     ========        ========
============
  IDEL                      274
  SDEL                    213371        0.004          20.934
0.352
  SGET                    201906        0.010        1740.290
1.132
  SHEA                     22716        0.005           2.349
0.272
  SPUT                  1771398        0.011        1770.563
0.487
```

In this example, SGET (S3 GET) operations are the slowest on average at 1.13 seconds, but SGET and SPUT (S3 PUT) operations both show long worst-case times of about 1,770 seconds.

b. To show the slowest 10 retrieval operations, use the grep command to select only SGET messages and add the long output option (`-l`) to include object paths:

```
grep SGET audit.log | audit-sum -l
```

The results include the type (object or bucket) and path, which allows you to grep the audit log for other messages relating to these particular objects.

```
Total:          201906 operations
    Slowest:      1740.290 sec
    Average:         1.132 sec
    Fastest:         0.010 sec
    Slowest operations:
        time(usec)       source ip         type      size(B) path
        ========== =============== ============ ============ ====
        1740289662   10.96.101.125       object   5663711385
backup/r9O1OaQ8JB-1566861764-4519.iso
        1624414429   10.96.101.125       object   5375001556
backup/r9O1OaQ8JB-1566861764-6618.iso
        1533143793   10.96.101.125       object   5183661466
backup/r9O1OaQ8JB-1566861764-4518.iso
             70839   10.96.101.125       object        28338
bucket3/dat.1566861764-6619
             68487   10.96.101.125       object        27890
bucket3/dat.1566861764-6615
             67798   10.96.101.125       object        27671
bucket5/dat.1566861764-6617
             67027   10.96.101.125       object        27230
bucket5/dat.1566861764-4517
             60922   10.96.101.125       object        26118
bucket3/dat.1566861764-4520
             35588   10.96.101.125       object        11311
bucket3/dat.1566861764-6616
             23897   10.96.101.125       object        10692
bucket3/dat.1566861764-4516
```

From this example output, you can see that the three slowest S3 GET requests were for objects about 5 GB in size, which is much larger than the other objects. The large size accounts for the slow worst-case retrieval times.

3. If you want to determine what sizes of objects are being ingested into and retrieved from your grid, use the size option (`-s`):

   ```
   audit-sum -s audit.log
   ```

```
  message group           count       min(MB)          max(MB)
average(MB)
  =============           =====     ========        ========
============
  IDEL                      274        0.004        5000.000
1654.502
  SDEL                   213371        0.000          10.504
1.695
  SGET                   201906        0.000        5000.000
14.920
  SHEA                    22716        0.001          10.504
2.967
  SPUT                  1771398        0.000        5000.000
2.495
```

In this example, the average object size for SPUT is under 2.5 MB, but the average size for SGET is much larger. The number of SPUT messages is much higher than the number of SGET messages, indicating that most objects are never retrieved.

4. If you want to determine if retrievals were slow yesterday:

   a. Issue the command on the appropriate audit log and use the group-by-time option (`-gt`), followed by the time period (for example, 15M, 1H, 10S):

      ```
      grep SGET audit.log | audit-sum -gt 1H
      ```

```
  message group           count    min(sec)       max(sec)
average(sec)
  =============           =====     ========        ========
============
  2019-09-05T00            7591        0.010        1481.867
1.254
  2019-09-05T01            4173        0.011        1740.290
1.115
  2019-09-05T02           20142        0.011        1274.961
1.562
  2019-09-05T03           57591        0.010        1383.867
1.254
  2019-09-05T04          124171        0.013        1740.290
1.405
  2019-09-05T05          420182        0.021        1274.511
1.562
  2019-09-05T06         1220371        0.015        6274.961
5.562
  2019-09-05T07          527142        0.011        1974.228
2.002
  2019-09-05T08          384173        0.012        1740.290
1.105
  2019-09-05T09           27591        0.010        1481.867
1.354
```

These results show that S3 GET traffic spiked between 06:00 and 07:00. The max and average times are both considerably higher at these times as well, and they did not ramp up gradually as the count increased. This suggests that capacity was exceeded somewhere, perhaps in the network or in the grid's ability to process requests.

b. To determine what size objects were being retrieved each hour yesterday, add the size option (`-s`) to the command:

```
grep SGET audit.log | audit-sum -gt 1H -s
```

```
  message group           count       min(B)          max(B)
average(B)
  =============           =====     ========        ========
============
  2019-09-05T00            7591        0.040        1481.867
1.976
  2019-09-05T01            4173        0.043        1740.290
2.062
  2019-09-05T02           20142        0.083        1274.961
2.303
  2019-09-05T03           57591        0.912        1383.867
1.182
  2019-09-05T04          124171        0.730        1740.290
1.528
  2019-09-05T05          420182        0.875        4274.511
2.398
  2019-09-05T06         1220371        0.691  5663711385.961
51.328
  2019-09-05T07          527142        0.130        1974.228
2.147
  2019-09-05T08          384173        0.625        1740.290
1.878
  2019-09-05T09           27591        0.689        1481.867
1.354
```

These results indicate that some very large retrievals occurred when the overall retrieval traffic was at its maximum.

c. To see more detail, use the audit-explain tool to review all the SGET operations during that hour:

```
grep 2019-09-05T06 audit.log | grep SGET | audit-explain | less
```

If the output of the grep command is expected to be many lines, add the `less` command to show the contents of the audit log file one page (one screen) at a time.

5. If you want to determine if SPUT operations on buckets are slower than SPUT operations for objects:

a. Start by using the `-go` option, which groups messages for object and bucket operations separately:

```
grep SPUT sample.log | audit-sum -go
```

```
  message group           count     min(sec)        max(sec)
average(sec)
  =============           =====     ========        ========
============
  SPUT.bucket                 1        0.125           0.125
0.125
  SPUT.object                12        0.025           1.019
0.236
```

The results show that SPUT operations for buckets have different performance characteristics than SPUT operations for objects.

b. To determine which buckets have the slowest SPUT operations, use the `-gb` option, which groups messages by bucket:

```
grep SPUT audit.log | audit-sum -gb
```

```
  message group                  count     min(sec)        max(sec)
average(sec)
  =============                  =====     ========        ========
============
  SPUT.cho-non-versioning        71943        0.046        1770.563
1.571
  SPUT.cho-versioning            54277        0.047        1736.633
1.415
  SPUT.cho-west-region           80615        0.040          55.557
1.329
  SPUT.ldt002                  1564563        0.011          51.569
0.361
```

c. To determine which buckets have the largest SPUT object size, use both the `-gb` and the `-s` options:

```
grep SPUT audit.log | audit-sum -gb -s
```

```
  message group                    count       min(B)          max(B)
average(B)
  =============                    =====     ========        ========
============
  SPUT.cho-non-versioning          71943        2.097        5000.000
21.672
  SPUT.cho-versioning              54277        2.097        5000.000
21.120
  SPUT.cho-west-region             80615        2.097         800.000
14.433
  SPUT.ldt002                    1564563        0.000         999.972
0.352
```

Audit message format

Audit message format: Overview

Audit messages exchanged within the StorageGRID system include standard information common to all messages and specific content describing the event or activity being reported.

If the summary information provided by the audit-explain and audit-sum tools is insufficient, refer to this section to understand the general format of all audit messages.

The following is an example audit message as it might appear in the audit log file:

```
2014-07-17T03:50:47.484627
[AUDT:[RSLT(FC32):VRGN][AVER(UI32):10][ATIM(UI64):1405569047484627][ATYP(F
C32):SYSU][ANID(UI32):11627225][AMID(FC32):ARNI][ATID(UI64):94457363265006
03516]]
```

Each audit message contains a string of attribute elements. The entire string is enclosed in brackets (`[ ]`), and each attribute element in the string has the following characteristics:

- Enclosed in brackets `[ ]`
- Introduced by the string `AUDT`, which indicates an audit message
- Without delimiters (no commas or spaces) before or after
- Terminated by a line feed character `\n`

Each element includes an attribute code, a data type, and a value that are reported in this format:

```
[ATTR(type):value][ATTR(type):value]...
[ATTR(type):value]\n
```

The number of attribute elements in the message depends on the event type of the message. The attribute elements aren't listed in any particular order.

The following list describes the attribute elements:

- `ATTR` is a four-character code for the attribute being reported. There are some attributes that are common to all audit messages and others that are event-specific.
- `type` is a four-character identifier of the programming data type of the value, such as UI64, FC32, and so on. The type is enclosed in parentheses `( )`.
- `value` is the content of the attribute, typically a numeric or text value. Values always follow a colon (`:`). Values of data type CSTR are surrounded by double quotes " ".

Data types

Different data types are used to store information in audit messages.

Type	Description
UI32	Unsigned long integer (32 bits); it can store the numbers 0 to 4,294,967,295.
UI64	Unsigned double long integer (64 bits); it can store the numbers 0 to 18,446,744,073,709,551,615.
FC32	Four-character constant; a 32-bit unsigned integer value represented as four ASCII characters such as "ABCD."
IPAD	Used for IP addresses.
CSTR	A variable-length array of UTF-8 characters. Characters can be escaped with the following conventions: • Backslash is \\. • Carriage return is \r. • Double quotes is \". • Line feed (new line) is \n. • Characters can be replaced by their hexadecimal equivalents (in the format \xHH, where HH is the hexadecimal value representing the character).

Event-specific data

Each audit message in the audit log records data specific to a system event.

Following the opening `[AUDT:` container that identifies the message itself, the next set of attributes provide information about the event or action described by the audit message. These attributes are highlighted in the following example:

```
2018-12-05T08:24:45.921845 [AUDT:*\[RSLT\(FC32\):SUCS\]*
\[TIME\(UI64\):11454\]\[SAIP\(IPAD\):"10.224.0.100"\]\[S3AI\(CSTR\):"60025621595611246499"\]
\[SACC\(CSTR\):"account"\]\[S3AK\(CSTR\):"SGKH4_Nc8SO1H6w3w0nCOFCGgk__E6dYzKlumRs
KJA=="\]
\[SUSR\(CSTR\):"urn:sgws:identity::60025621595611246499:root"\]
\[SBAI\(CSTR\):"60025621595611246499"\]\[SBAC\(CSTR\):"account"\]\[S3BK\(CSTR\):"bucket"\]
\[S3KY\(CSTR\):"object"\]\[CBID\(UI64\):0xCC128B9B9E428347\]
\[UUID\(CSTR\):"B975D2CE-E4DA-4D14-8A23-
1CB4B83F2CD8"\]\[CSIZ\(UI64\):30720\][AVER(UI32):10]
\[ATIM(UI64):1543998285921845]\[ATYP\(FC32\):SHEA\][ANID(UI32):12281045][AMID(FC32):S3RQ]
\[ATID(UI64):15552417629170647261]]
```

The `ATYP` element (underlined in the example) identifies which event generated the message. This example message includes the SHEA message code ([ATYP(FC32):SHEA]), indicating it was generated by a successful S3 HEAD request.

Common elements in audit messages

All audit messages contain the common elements.

Code	Type	Description
AMID	FC32	Module ID: A four-character identifier of the module ID that generated the message. This indicates the code segment within which the audit message was generated.
ANID	UI32	Node ID: The grid node ID assigned to the service that generated the message. Each service is allocated a unique identifier at the time the StorageGRID system is configured and installed. This ID can't be changed.
ASES	UI64	Audit Session Identifier: In previous releases, this element indicated the time at which the audit system was initialized after the service started up. This time value was measured in microseconds since the operating system epoch (00:00:00 UTC on 1 January, 1970). **Note:** This element is obsolete and no longer appears in audit messages.
ASQN	UI64	Sequence Count: In previous releases, this counter was incremented for each generated audit message on the grid node (ANID) and reset to zero at service restart. **Note:** This element is obsolete and no longer appears in audit messages.
ATID	UI64	Trace ID: An identifier that is shared by the set of messages that were triggered by a single event.

Code	Type	Description
ATIM	UI64	Timestamp: The time the event was generated that triggered the audit message, measured in microseconds since the operating system epoch (00:00:00 UTC on 1 January, 1970). Note that most available tools for converting the timestamp to local date and time are based on milliseconds. Rounding or truncation of the logged timestamp might be required. The human-readable time that appears at the beginning of the audit message in the `audit.log` file is the ATIM attribute in ISO 8601 format. The date and time are represented as *`YYYY-MMDDTHH:MM:SS.UUUUUU`*, where the `T` is a literal string character indicating the beginning of the time segment of the date. *`UUUUUU`* are microseconds.
ATYP	FC32	Event Type: A four-character identifier of the event being logged. This governs the "payload" content of the message: the attributes that are included.
AVER	UI32	Version: The version of the audit message. As the StorageGRID software evolves, new versions of services might incorporate new features in audit reporting. This field enables backward compatibility in the AMS service to process messages from older versions of services.
RSLT	FC32	Result: The result of event, process, or transaction. If is not relevant for a message, NONE is used rather than SUCS so that the message is not accidentally filtered.

Audit message examples

You can find detailed information in each audit message. All audit messages use the same format.

The following is an example audit message as it might appear in the `audit.log` file:

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"][
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"][S3BK(CSTR):"s3small1"][S3K
Y(CSTR):"hello1"][CBID(UI64):0x50C4F7AC2BC8EDF7][CSIZ(UI64):0
][AVER(UI32):10][ATIM(UI64):1405631878959669][ATYP(FC32):SPUT
][ANID(UI32):12872812][AMID(FC32):S3RQ][ATID(UI64):1579224144
102530435]]
```

The audit message contains information about the event being recorded, as well as information about the audit message itself.

To identify which event is recorded by the audit message, look for the ATYP attribute (highlighted below):

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"][
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"][S3BK(CSTR):"s3small1"][S3K
Y(CSTR):"hello1"][CBID(UI64):0x50C4F7AC2BC8EDF7][CSIZ(UI64):0
][AVER(UI32):10][ATIM(UI64):1405631878959669][ATYP(FC32):SP
UT][ANID(UI32):12872812][AMID(FC32):S3RQ][ATID(UI64):1579224
144102530435]]
```

The value of the ATYP attribute is SPUT. SPUT represents an S3 PUT transaction, which logs the ingest of an object to a bucket.

The following audit message also shows the bucket to which the object is associated:

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"][
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"][S3BK\(CSTR\):"s3small1"][S3
KY(CSTR):"hello1"][CBID(UI64):0x50C4F7AC2BC8EDF7][CSIZ(UI64):
0][AVER(UI32):10][ATIM(UI64):1405631878959669][ATYP(FC32):SPU
T][ANID(UI32):12872812][AMID(FC32):S3RQ][ATID(UI64):157922414
4102530435]]
```

To discover when the PUT event occurred, note the Universal Coordinated Time (UTC) timestamp at the beginning of the audit message. This value is a human-readable version of the ATIM attribute of the audit message itself:

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"][
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"][S3BK(CSTR):"s3small1"][S3K
Y(CSTR):"hello1"][CBID(UI64):0x50C4F7AC2BC8EDF7][CSIZ(UI64):0
][AVER(UI32):10][ATIM\(UI64\):1405631878959669][ATYP(FC32):SP
UT][ANID(UI32):12872812][AMID(FC32):S3RQ][ATID(UI64):15792241
44102530435]]
```

ATIM records the time, in microseconds, since the beginning of the UNIX epoch. In the example, the value `1405631878959669` translates to Thursday, 17-Jul-2014 21:17:59 UTC.

Audit messages and the object lifecycle

When are audit message generated?

Audit messages are generated each time an object is ingested, retrieved, or deleted. You

can identify these transactions in the audit log by locating API-specific (S3 or Swift) audit messages.

Audit messages are linked through identifiers specific to each protocol.

Protocol	Code
Linking S3 operations	S3BK (bucket), S3KY (key), or both
Linking Swift operations	WCON (container), WOBJ (object), or both
Linking internal operations	CBID (object's internal identifier)

Timing of audit messages

Because of factors such as timing differences between grid nodes, object size, and network delays, the order of audit messages generated by the different services can vary from that shown in the examples in this section.

Archive Nodes

The series of audit messages generated when an Archive Node sends object data to an external archival storage system is similar to that for Storage Nodes except that there is no SCMT (Store Object Commit) message, and the ATCE (Archive Object Store Begin) and ASCE (Archive Object Store End) messages are generated for each archived copy of object data.

The series of audit messages generated when an Archive Node retrieves object data from an external archival storage system is similar to that for Storage Nodes except that the ARCB (Archive Object Retrieve Begin) and ARCE (Archive Object Retrieve End) messages are generated for each retrieved copy of object data.

The series of audit messages generated when an Archive Node deletes object data from an external archival storage system is similar to that for Storage Nodes except that there is no SREM (Object Store Remove) message, and there is an AREM (Archive Object Remove) message for each delete request.

Object ingest transactions

You can identify client ingest transactions in the audit log by locating API-specific (S3 or Swift) audit messages.

Not all audit messages generated during an ingest transaction are listed in the following tables. Only the messages required to trace the ingest transaction are included.

S3 ingest audit messages

Code	Name	Description	Trace	See
SPUT	S3 PUT transaction	An S3 PUT ingest transaction has completed successfully.	CBID, S3BK, S3KY	SPUT: S3 PUT
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	ORLM: Object Rules Met

Swift ingest audit messages

Code	Name	Description	Trace	See
WPUT	Swift PUT transaction	A Swift PUT ingest transaction has successfully completed.	CBID, WCON, WOBJ	WPUT: Swift PUT
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	ORLM: Object Rules Met

Example: S3 object ingest

The series of audit messages below is an example of the audit messages generated and saved to the audit log when an S3 client ingests an object to a Storage Node (LDR service).

In this example, the active ILM policy includes the Make 2 Copies ILM rule.

> Not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 ingest transaction (SPUT) are listed.

This example assumes that an S3 bucket has been previously created.

SPUT: S3 PUT

The SPUT message is generated to indicate that an S3 PUT transaction has been issued to create an object in a specific bucket.

```
2017-07-
17T21:17:58.959669[AUDT:[RSLT(FC32):SUCS][TIME(UI64):25771][SAIP(IPAD):"10
.96.112.29"][S3AI(CSTR):"70899244468554783528"][SACC(CSTR):"test"][S3AK(CS
TR):"SGKHyalRU_5cLflqajtaFmxJn946lAWRJfBF33gAOg=="][SUSR(CSTR):"urn:sgws:i
dentity::70899244468554783528:root"][SBAI(CSTR):"70899244468554783528"][SB
AC(CSTR):"test"][S3BK(CSTR):"example"][S3KY(CSTR):"testobject-0-
3"][CBID\(UI64\):0x8EF52DF8025E63A8][CSIZ(UI64):30720][AVER(UI32):10][ATIM
(UI64):150032627859669][ATYP\(FC32\):SPUT][ANID(UI32):12086324][AMID(FC32)
:S3RQ][ATID(UI64):14399932238768197038]]
```

ORLM: Object Rules Met

The ORLM message indicates that the ILM policy has been satisfied for this object. The message includes the object's CBID and the name of the ILM rule that was applied.

For replicated objects, the LOCS field includes the LDR node ID and volume ID of the object locations.

```
2019-07-
17T21:18:31.230669[AUDT:[CBID\(UI64\):0x50C4F7AC2BC8EDF7][RULE(CSTR):"Make
2 Copies"][STAT(FC32):DONE][CSIZ(UI64):0][UUID(CSTR):"0B344E18-98ED-4F22-
A6C8-A93ED68F8D3F"][LOCS(CSTR):"CLDI 12828634 2148730112, CLDI 12745543
2147552014"][RSLT(FC32):SUCS][AVER(UI32):10][ATYP\(FC32\):ORLM][ATIM(UI64)
:1563398230669][ATID(UI64):15494889725796157557][ANID(UI32):13100453][AMID
(FC32):BCMS]]
```

For erasure-coded objects, the LOCS field includes the erasure-coding profile ID and the erasure coding group ID

```
2019-02-23T01:52:54.647537
[AUDT:[CBID(UI64):0xFA8ABE5B5001F7E2][RULE(CSTR):"EC_2_plus_1"][STAT(FC32)
:DONE][CSIZ(UI64):10000][UUID(CSTR):"E291E456-D11A-4701-8F51-
D2F7CC9AFECA"][LOCS(CSTR):"CLEC 1 A471E45D-A400-47C7-86AC-
12E77F229831"][RSLT(FC32):SUCS][AVER(UI32):10][ATIM(UI64):1550929974537]\[
ATYP\(FC32\):ORLM\][ANID(UI32):12355278][AMID(FC32):ILMX][ATID(UI64):41685
59046473725560]]
```

The PATH field includes S3 bucket and key information or Swift container and object information, depending on which API was used.

```
2019-09-15.txt:2018-01-24T13:52:54.131559
[AUDT:[CBID(UI64):0x82704DFA4C9674F4][RULE(CSTR):"Make 2
Copies"][STAT(FC32):DONE][CSIZ(UI64):3145729][UUID(CSTR):"8C1C9CAC-22BB-
4880-9115-
CE604F8CE687"][PATH(CSTR):"frisbee_Bucket1/GridDataTests151683676324774_1_
1vf9d"][LOCS(CSTR):"CLDI 12525468, CLDI
12222978"][RSLT(FC32):SUCS][AVER(UI32):10][ATIM(UI64):1568555574559][ATYP(
FC32):ORLM][ANID(UI32):12525468][AMID(FC32):OBDI][ATID(UI64):3448338865383
69336]]
```

Object delete transactions

You can identify object delete transactions in the audit log by locating API-specific (S3 and Swift) audit messages.

Not all audit messages generated during a delete transaction are listed in the following tables. Only messages required to trace the delete transaction are included.

S3 delete audit messages

Code	Name	Description	Trace	See
SDEL	S3 Delete	Request made to delete the object from a bucket.	CBID, S3KY	SDEL: S3 DELETE

Swift delete audit messages

Code	Name	Description	Trace	See
WDEL	Swift Delete	Request made to delete the object from a container, or the container.	CBID, WOBJ	WDEL: Swift DELETE

Example: S3 object deletion

When an S3 client deletes an object from a Storage Node (LDR service), an audit message is generated and saved to the audit log.

> Not all audit messages generated during a delete transaction are listed in the example below. Only those related to the S3 delete transaction (SDEL) are listed.

SDEL: S3 Delete

Object deletion begins when the client sends a DeleteObject request to an LDR service. The message contains the bucket from which to delete the object and the object's S3 Key, which is used to identify the object.

```
2017-07-
17T21:17:58.959669[AUDT:[RSLT(FC32):SUCS][TIME(UI64):14316][SAIP(IPAD):"10
.96.112.29"][S3AI(CSTR):"70899244468554783528"][SACC(CSTR):"test"][S3AK(CS
TR):"SGKHyalRU_5cLflqajtaFmxJn946lAWRJfBF33gAOg=="][SUSR(CSTR):"urn:sgws:i
dentity::70899244468554783528:root"][SBAI(CSTR):"70899244468554783528"][SB
AC(CSTR):"test"]\[S3BK\(CSTR\):"example"\]\[S3KY\(CSTR\):"testobject-0-
7"\][CBID\(UI64\):0x339F21C5A6964D89][CSIZ(UI64):30720][AVER(UI32):10][ATI
M(UI64):150032627859669][ATYP\(FC32\):SDEL][ANID(UI32):12086324][AMID(FC32
):S3RQ][ATID(UI64):4727861330952970593]]
```

Object retrieve transactions

You can identify object retrieve transactions in the audit log by locating API-specific (S3 and Swift) audit messages.

Not all audit messages generated during a retrieve transaction are listed in the following tables. Only messages required to trace the retrieve transaction are included.

S3 retrieval audit messages

Code	Name	Description	Trace	See
SGET	S3 GET	Request made to retrieve an object from a bucket.	CBID, S3BK, S3KY	SGET: S3 GET

Swift retrieval audit messages

Code	Name	Description	Trace	See
WGET	Swift GET	Request made to retrieve an object from a container.	CBID, WCON, WOBJ	WGET: Swift GET

Example: S3 object retrieval

When an S3 client retrieves an object from a Storage Node (LDR service), an audit message is generated and saved to the audit log.

Note that not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 retrieval transaction (SGET) are listed.

SGET: S3 GET

Object retrieval begins when the client sends a GetObject request to an LDR service. The message contains the bucket from which to retrieve the object and the object's S3 Key, which is used to identify the object.

```
2017-09-20T22:53:08.782605
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):47807][SAIP(IPAD):"10.96.112.26"][S3AI(
CSTR):"43979298178977966408"][SACC(CSTR):"s3-account-
a"][S3AK(CSTR):"SGKHt7GzEcu0yXhFhT_rL5mep4nJt1w75GBh-
O_FEw=="][SUSR(CSTR):"urn:sgws:identity::43979298178977966408:root"][SBAI(
CSTR):"43979298178977966408"][SBAC(CSTR):"s3-account-
a"]\[S3BK\(CSTR\):"bucket-
anonymous"\]\[S3KY\(CSTR\):"Hello.txt"\][CBID(UI64):0x83D70C6F1F662B02][CS
IZ(UI64):12][AVER(UI32):10][ATIM(UI64):1505947988782605]\[ATYP\(FC32\):SGE
T\][ANID(UI32):12272050][AMID(FC32):S3RQ][ATID(UI64):17742374343649889669]
]
```

If the bucket policy allows, a client can anonymously retrieve objects, or can retrieve objects from a bucket that is owned by a different tenant account. The audit message contains information about the bucket owner's tenant account so that you can track these anonymous and cross-account requests.

In the following example message, the client sends a GetObject request for an object stored in a bucket that they don't own. The values for SBAI and SBAC record the bucket owner's tenant account ID and name, which differs from the tenant account ID and name of the client recorded in S3AI and SACC.

```
2017-09-20T22:53:15.876415
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):53244][SAIP(IPAD):"10.96.112.26"]\[S3AI
\(CSTR\):"17915054115450519830"\]\[SACC\(CSTR\):"s3-account-
b"\][S3AK(CSTR):"SGKHpoblWlP_kBkqSCbTi754Ls8lBUog67I2LlSiUg=="][SUSR(CSTR)
:"urn:sgws:identity::17915054115450519830:root"]\[SBAI\(CSTR\):"4397929817
8977966408"\]\[SBAC\(CSTR\):"s3-account-a"\][S3BK(CSTR):"bucket-
anonymous"][S3KY(CSTR):"Hello.txt"][CBID(UI64):0x83D70C6F1F662B02][CSIZ(UI
64):12][AVER(UI32):10][ATIM(UI64):1505947995876415][ATYP(FC32):SGET][ANID(
UI32):12272050][AMID(FC32):S3RQ][ATID(UI64):6888780247515624902]]
```

Example: S3 Select on an object

When an S3 client issues an S3 Select query on an object, audit messages are generated and saved to the audit log.

Note that not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 Select transaction (SelectObjectContent) are listed.

Each query results in two audit messages: one that performs the authorization of the S3 Select request (the S3SR field is set to "select") and a subsequent standard GET operation that retrieves the data from storage during processing.

```
2021-11-08T15:35:30.750038
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1636385730715700][TIME(UI64):29173][SAI
P(IPAD):"192.168.7.44"][S3AI(CSTR):"63147909414576125820"][SACC(CSTR):"Ten
ant1636027116"][S3AK(CSTR):"AUFD1XNVZ905F3TW7KSU"][SUSR(CSTR):"urn:sgws:id
entity::63147909414576125820:root"][SBAI(CSTR):"63147909414576125820"][SBA
C(CSTR):"Tenant1636027116"][S3BK(CSTR):"619c0755-9e38-42e0-a614-
05064f74126d"][S3KY(CSTR):"SUB-
EST2020_ALL.csv"][CBID(UI64):0x0496F0408A721171][UUID(CSTR):"D64B1A4A-
9F01-4EE7-B133-
08842A099628"][CSIZ(UI64):0][S3SR(CSTR):"select"][AVER(UI32):10][ATIM(UI64
):1636385730750038][ATYP(FC32):SPOS][ANID(UI32):12601166][AMID(FC32):S3RQ]
[ATID(UI64):13630097093968959 85]]
```

```
2021-11-08T15:35:32.604886
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1636383069486504][TIME(UI64):430690][SA
IP(IPAD):"192.168.7.44"][HTRH(CSTR):"{\"x-forwarded-
for\":\"unix:\"}"][S3AI(CSTR):"63147909414576125820"][SACC(CSTR):"Tenant16
36027116"][S3AK(CSTR):"AUFD1XNVZ905F3TW7KSU"][SUSR(CSTR):"urn:sgws:identit
y::63147909414576125820:root"][SBAI(CSTR):"63147909414576125820"][SBAC(CST
R):"Tenant1636027116"][S3BK(CSTR):"619c0755-9e38-42e0-a614-
05064f74126d"][S3KY(CSTR):"SUB-
EST2020_ALL.csv"][CBID(UI64):0x0496F0408A721171][UUID(CSTR):"D64B1A4A-
9F01-4EE7-B133-
08842A099628"][CSIZ(UI64):10185581][MTME(UI64):1636380348695262][AVER(UI32
):10][ATIM(UI64):1636385732604886][ATYP(FC32):SGET][ANID(UI32):12733063][A
MID(FC32):S3RQ][ATID(UI64):16562288121152341130]]
```

Metadata update messages

Audit messages are generated when an S3 client updates an object's metadata.

S3 metadata update audit messages

Code	Name	Description	Trace	See
SUPD	S3 Metadata Updated	Generated when an S3 client updates the metadata for an ingested object.	CBID, S3KY, HTRH	SUPD: S3 Metadata Updated

Example: S3 metadata update

The example shows a successful transaction to update the metadata for an existing S3 object.

SUPD: S3 Metadata Update

The S3 client makes a request (SUPD) to update the specified metadata (`x-amz-meta-\*`) for the S3 object (S3KY). In this example, request headers are included in the field HTRH because it has been configured as an audit protocol header (**CONFIGURATION** > **Monitoring** > **Audit and syslog server**). See Configure audit messages and log destinations.

```
2017-07-11T21:54:03.157462
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):17631][SAIP(IPAD):"10.96.100.254"]
[HTRH(CSTR):"{\"accept-encoding\":\"identity\",\"authorization\":\"AWS
LIUF17FGJARQHPY2E761:jul/hnZs/uNY+aVvV0lTSYhEGts=\",
\"content-length\":\"0\",\"date\":\"Tue, 11 Jul 2017 21:54:03
GMT\",\"host\":\"10.96.99.163:18082\",
\"user-agent\":\"aws-cli/1.9.20 Python/2.7.6 Linux/3.13.0-119-generic
botocore/1.3.20\",
\"x-amz-copy-source\":\"/testbkt1/testobj1\",\"x-amz-metadata-
directive\":\"REPLACE\",\"x-amz-meta-city\":\"Vancouver\"}"]
[S3AI(CSTR):"20956855414285633225"][SACC(CSTR):"acct1"][S3AK(CSTR):"SGKHyy
v9ZQqWRbJSQc5vI7mgioJwrdplShE02AUaww=="]
[SUSR(CSTR):"urn:sgws:identity::20956855414285633225:root"]
[SBAI(CSTR):"20956855414285633225"][SBAC(CSTR):"acct1"][S3BK(CSTR):"testbk
t1"]
[S3KY(CSTR):"testobj1"][CBID(UI64):0xCB1D5C213434DD48][CSIZ(UI64):10][AVER
(UI32):10]
[ATIM(UI64):1499810043157462][ATYP(FC32):SUPD][ANID(UI32):12258396][AMID(F
C32):S3RQ]
[ATID(UI64):8987436599021955788]]
```

Audit messages

Audit messages: Overview

Detailed descriptions of audit messages returned by the system are listed in the following sections. Each audit message is first listed in a table that groups related messages by the class of activity that the message represents. These groupings are useful both for understanding the types of activities that are audited, and for selecting the desired type of audit message filtering.

The audit messages are also listed alphabetically by their four-character codes. This alphabetic list enables you to find information about specific messages.

The four-character codes used throughout this chapter are the ATYP values found in the audit messages as shown in the following example message:

```
2014-07-17T03:50:47.484627
\[AUDT:[RSLT(FC32):VRGN][AVER(UI32):10][ATIM(UI64):1405569047484627][ATYP\
(FC32\):SYSU][ANID(UI32):11627225][AMID(FC32):ARNI][ATID(UI64):94457363265
00603516]]
```

For information about setting audit message levels, changing log destinations, and using an external syslog server for your audit information, see Configure audit messages and log destinations

Audit message categories

System audit messages

The audit messages belonging to the system audit category are used for events related to the auditing system itself, grid node states, system-wide task activity (grid tasks), and service backup operations.

Code	Message title and description	See
ECMC	Missing Erasure-Coded Data Fragment: Indicates that a missing erasure-coded data fragment has been detected.	ECMC: Missing Erasure-Coded Data Fragment
ECOC	Corrupt Erasure-Coded Data Fragment: Indicates that a corrupt erasure-coded data fragment has been detected.	ECOC: Corrupt Erasure-Coded Data Fragment
ETAF	Security Authentication Failed: A connection attempt using Transport Layer Security (TLS) failed.	ETAF: Security Authentication Failed
GNRG	GNDS Registration: A service updated or registered information about itself in the StorageGRID system.	GNRG: GNDS Registration
GNUR	GNDS Unregistration: A service has unregistered itself from the StorageGRID system.	GNUR: GNDS Unregistration
GTED	Grid Task Ended: The CMN service finished processing the grid task.	GTED: Grid Task Ended
GTST	Grid Task Started: The CMN service started to process the grid task.	GTST: Grid Task Started
GTSU	Grid Task Submitted: A grid task was submitted to the CMN service.	GTSU: Grid Task Submitted
LLST	Location Lost: This audit message is generated when a location is lost.	LLST: Location Lost
OLST	Object Lost: A requested object cannot be located within the StorageGRID system.	OLST: System Detected Lost Object
SADD	Security Audit Disable: Audit message logging was turned off.	SADD: Security Audit Disable
SADE	Security Audit Enable: Audit message logging has been restored.	SADE: Security Audit Enable

Code	Message title and description	See
SVRF	Object Store Verify Fail: A content block failed verification checks.	SVRF: Object Store Verify Fail
SVRU	Object Store Verify Unknown: Unexpected object data detected in the object store.	SVRU: Object Store Verify Unknown
SYSD	Node Stop: A shutdown was requested.	SYSD: Node Stop
SYST	Node Stopping: A service initiated a graceful stop.	SYST: Node Stopping
SYSU	Node Start: A service started; the nature of the previous shutdown is indicated in the message.	SYSU: Node Start

Object storage audit messages

The audit messages belonging to the object storage audit category are used for events related to the storage and management of objects within the StorageGRID system. These include object storage and retrievals, grid-node to grid-node transfers, and verifications.

Code	Description	See
APCT	Archive Purge from Cloud-Tier: Archived object data is deleted from an external archival storage system, which connects to the StorageGRID through the S3 API.	APCT: Archive Purge from Cloud-Tier
ARCB	Archive Object Retrieve Begin: The ARC service begins the retrieval of object data from the external archival storage system.	ARCB: Archive Object Retrieve Begin
ARCE	Archive Object Retrieve End: Object data has been retrieved from an external archival storage system, and the ARC service reports the status of the retrieval operation.	ARCE: Archive Object Retrieve End
ARCT	Archive Retrieve from Cloud-Tier: Archived object data is retrieved from an external archival storage system, which connects to the StorageGRID through the S3 API.	ARCT: Archive Retrieve from Cloud-Tier
AREM	Archive Object Remove: A content block was successfully or unsuccessfully deleted from the external archival storage system.	AREM: Archive Object Remove

Code	Description	See
ASCE	Archive Object Store End: A content block has been written to the external archival storage system, and the ARC service reports the status of the write operation.	ASCE: Archive Object Store End
ASCT	Archive Store Cloud-Tier: Object data is stored to an external archival storage system, which connects to the StorageGRID through the S3 API.	ASCT: Archive Store Cloud-Tier
ATCE	Archive Object Store Begin: Writing a content block to an external archival storage has started.	ATCE: Archive Object Store Begin
AVCC	Archive Validate Cloud-Tier Configuration: The account and bucket settings provided were successfully or unsuccessfully validated.	AVCC: Archive Validate Cloud-Tier Configuration
BROR	Bucket Read Only Request: A bucket entered or exited read-only mode.	BROR: Bucket Read Only Request
CBSE	Object Send End: The source entity completed a grid-node to grid-node data transfer operation.	CBSE: Object Send End
CBRE	Object Receive End: The destination entity completed a grid-node to grid-node data transfer operation.	CBRE: Object Receive End
CGRR	Cross-Grid Replication Request: StorageGRID attempted a cross-grid replication operation to replicate objects between buckets in a grid federation connection.	CGRR: Cross-Grid Replication Request
EBDL	Empty Bucket Delete: The ILM scanner deleted an object in a bucket that is deleting all objects (performing an empty bucket operation).	EBDL: Empty Bucket Delete
EBKR	Empty Bucket Request: A user sent a request to turn empty bucket on or off (that is, to delete bucket objects or to stop deleting objects).	EBKR: Empty Bucket Request
SCMT	Object Store Commit: A content block was completely stored and verified, and can now be requested.	SCMT: Object Store Commit Request
SREM	Object Store Remove: A content block was deleted from a grid node, and can no longer be requested directly.	SREM: Object Store Remove

Client read audit messages

Client read audit messages are logged when an S3 or Swift client application makes a request to retrieve an object.

Code	Description	Used by	See
S3SL	S3 Select request: Logs a completion after an S3 Select request has been returned to the client. The S3SL message can include error message and error code details. The request might not have been successful.	S3 client	S3SL: S3 Select request
SGET	S3 GET: Logs a successful transaction to retrieve an object or list the objects in a bucket. **Note:** If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	SGET: S3 GET
SHEA	S3 HEAD: Logs a successful transaction to check for the existence of an object or bucket.	S3 client	SHEA: S3 HEAD
WGET	Swift GET: Logs a successful transaction to retrieve an object or list the objects in a container.	Swift client	WGET: Swift GET
WHEA	Swift HEAD: Logs a successful transaction to check for the existence of an object or container.	Swift client	WHEA: Swift HEAD

Client write audit messages

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

Code	Description	Used by	See
OVWR	Object Overwrite: Logs a transaction to overwrite one object with another object.	S3 and Swift clients	OVWR: Object Overwrite
SDEL	S3 DELETE: Logs a successful transaction to delete an object or bucket. **Note:** If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	SDEL: S3 DELETE
SPOS	S3 POST: Logs a successful transaction to restore an object from AWS Glacier storage to a Cloud Storage Pool.	S3 client	SPOS: S3 POST

Code	Description	Used by	See
SPUT	S3 PUT: Logs a successful transaction to create a new object or bucket. **Note:** If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	SPUT: S3 PUT
SUPD	S3 Metadata Updated: Logs a successful transaction to update the metadata for an existing object or bucket.	S3 client	SUPD: S3 Metadata Updated
WDEL	Swift DELETE: Logs a successful transaction to delete an object or container.	Swift client	WDEL: Swift DELETE
WPUT	Swift PUT: Logs a successful transaction to create a new object or container.	Swift client	WPUT: Swift PUT

Management audit message

The Management category logs user requests to the Management API.

Code	Message title and description	See
MGAU	Management API audit message: A log of user requests.	MGAU: Management audit message

ILM audit messages

The audit messages belonging to the ILM audit category are used for events related to information lifecycle management (ILM) operations.

Code	Message title and description	See
IDEL	ILM Initiated Delete: This audit message is generated when ILM starts the process of deleting an object.	IDEL: ILM Initiated Delete
LKCU	Overwritten Object Cleanup. This audit message is generated when an overwritten object is automatically removed to free up storage space.	LKCU: Overwritten Object Cleanup
ORLM	Object Rules Met: This audit message is generated when object data is stored as specified by the ILM rules.	ORLM: Object Rules Met

Audit message reference

APCT: Archive Purge from Cloud-Tier

This message is generated when archived object data is deleted from an external archival storage system, which connects to the StorageGRID through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was deleted.
CSIZ	Content Size	The size of the object in bytes. Always returns 0.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the cloud-tier from which the object was deleted.

ARCB: Archive Object Retrieve Begin

This message is generated when a request is made to retrieve archived object data and the retrieval process begins. Retrieval requests are processed immediately, but can be reordered to improve efficiency of retrieval from linear media such as tape.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival storage system.
RSLT	Result	Indicates the result of starting the archive retrieval process. Currently defined value is:SUCS: The content request was received and queued for retrieval.

This audit message marks the time of an archive retrieval. It allows you to match the message with a corresponding ARCE end message to determine the duration of archive retrieval, and whether the operation was successful.

ARCE: Archive Object Retrieve End

This message is generated when an attempt by the Archive Node to retrieve object data from an external archival storage system completes. If successful, the message indicates that the requested object data has been completely read from the archive location, and was successfully verified. After the object data has been retrieved and verified, it is delivered to the requesting service.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival storage system.

Code	Field	Description
VLID	Volume Identifier	The identifier of the volume on which the data was archived. If an archive location for the content is not found, a Volume ID of 0 is returned.
RSLT	Retrieval Result	The completion status of the archive retrieval process: • SUCS: successful • VRFL: failed (object verification failure) • ARUN: failed (external archival storage system unavailable) • CANC: failed (retrieval operation canceled) • GERR: failed (general error)

Matching this message with the corresponding ARCB message can indicate the time taken to perform the archive retrieval. This message indicates whether the retrieval was successful, and in the case of failure, the cause of the failure to retrieve the content block.

ARCT: Archive Retrieve from Cloud-Tier

This message is generated when archived object data is retrieved from an external archival storage system, which connects to the StorageGRID through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was retrieved.
CSIZ	Content Size	The size of the object in bytes. The value is only accurate for successful retrieves.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the external archival storage system.
TIME	Time	Total processing time for the request in microseconds.

AREM: Archive Object Remove

The Archive Object Remove audit message indicates that a content block was successfully or unsuccessfully deleted from an Archive Node. If the result is successful, the Archive Node has successfully informed the external archival storage system that StorageGRID has released an object location. Whether the object is removed from the external archive storage system depends on the type of system and its configuration.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival media system.
VLID	Volume Identifier	The identifier of the volume on which the object data was archived.
RSLT	Result	The completion status of the archive removal process: • SUCS: successful • ARUN: failed (external archival storage system unavailable) • GERR: failed (general error)

ASCE: Archive Object Store End

This message indicates that writing a content block to an external archival storage system has ended.

Code	Field	Description
CBID	Content Block Identifier	The identifier of the content block stored on the external archival storage system.
VLID	Volume Identifier	The unique identifier of the archive volume to which the object data is written.
VREN	Verification Enabled	Indicates if verification is performed for content blocks. Currently defined values are: • VENA: verification is enabled • VDSA: verification is disabled
MCLS	Management Class	A string identifying the TSM Management Class to which the content block is assigned if applicable.
RSLT	Result	Indicates the result of the archive process. Currently defined values are: • SUCS: successful (archiving process succeeded) • OFFL: failed (archiving is offline) • VRFL: failed (object verification failed) • ARUN: failed (external archival storage system unavailable) • GERR: failed (general error)

This audit message means that the specified content block has been written to the external archival storage system. If the write fails, the result provides basic troubleshooting information about where the failure occurred. More detailed information about archive failures can be found by examining Archive Node attributes in the StorageGRID system.

ASCT: Archive Store Cloud-Tier

This message is generated when archived object data is stored to an external archival storage system, which connects to StorageGRID through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was retrieved.
CSIZ	Content Size	The size of the object in bytes.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the cloud-tier the content was stored to.
TIME	Time	Total processing time for the request in microseconds.

ATCE: Archive Object Store Begin

This message indicates that writing a content block to an external archival storage has started.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the content block to be archived.
VLID	Volume Identifier	The unique identifier of the volume to which the content block is written. If the operation fails, a volume ID of 0 is returned.
RSLT	Result	Indicates the result of the transfer of the content block. Currently defined values are: • SUCS: success (content block stored successfully) • EXIS: ignored (content block was already stored) • ISFD: failed (insufficient disk space) • STER: failed (error storing the CBID) • OFFL: failed (archiving is offline) • GERR: failed (general error)

AVCC: Archive Validate Cloud-Tier Configuration

This message is generated when the configuration settings are validated for a Cloud Tiering - Simple Storage Service (S3) target type.

Code	Field	Description
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	UUID associated with the external archival storage system being validated.

BROR: Bucket Read Only Request

The LDR service generates this audit message when a bucket enters or exits read-only mode. For example, a bucket enters read-only mode while all objects are being deleted.

Code	Field	Description
BKHD	Bucket UUID	The bucket ID.
BROV	Bucket read-only request value	Whether the bucket is being made read-only or is leaving the read-only state (1 = read-only, 0 = not-read-only).
BROS	Bucket read-only reason	The reason the bucket is being made read-only or leaving the read-only state. For example, emptyBucket.
S3AI	S3 tenant account ID	The ID of the tenant account that sent the request. An empty value indicates anonymous access.
S3BK	S3 bucket	The S3 bucket name.

CBRB: Object Receive Begin

During normal system operations, content blocks are continuously transferred between different nodes as data is accessed, replicated and retained. When transfer of a content block from one node to another is initiated, this message is issued by the destination entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.

Code	Field	Description
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the first sequence count requested. If successful, the transfer begins from this sequence count.
CTES	Expected End Sequence Count	Indicates the last sequence count requested. If successful, the transfer is considered complete when this sequence count has been received.
RSLT	Transfer Start Status	Status at the time the transfer was started: SUCS: Transfer started successfully.

This audit message means a node-to-node data transfer operation was initiated on a single piece of content, as identified by its Content Block Identifier. The operation requests data from "Start Sequence Count" to "Expected End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow, and when combined with storage audit messages, to verify replica counts.

CBRE: Object Receive End

When transfer of a content block from one node to another is completed, this message is issued by the destination entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.

Code	Field	Description
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.
RSLT	Transfer Result	The result of the transfer operation (from the perspective of the sending entity): SUCS: transfer successfully completed; all requested sequence counts were sent. CONL: connection lost during transfer CTMO: connection timed-out during establishment or transfer UNRE: destination node ID unreachable CRPT: transfer ended due to reception of corrupt or invalid data

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from "Start Sequence Count" to "Actual End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

CBSB: Object Send Begin

During normal system operations, content blocks are continuously transferred between different nodes as data is accessed, replicated and retained. When transfer of a content block from one node to another is initiated, this message is issued by the source entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.

Code	Field	Description
CTSS	Start Sequence Count	Indicates the first sequence count requested. If successful, the transfer begins from this sequence count.
CTES	Expected End Sequence Count	Indicates the last sequence count requested. If successful, the transfer is considered complete when this sequence count has been received.
RSLT	Transfer Start Status	Status at the time the transfer was started: SUCS: transfer started successfully.

This audit message means a node-to-node data transfer operation was initiated on a single piece of content, as identified by its Content Block Identifier. The operation requests data from "Start Sequence Count" to "Expected End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow, and when combined with storage audit messages, to verify replica counts.

CBSE: Object Send End

When transfer of a content block from one node to another is completed, this message is issued by the source entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.

Code	Field	Description
RSLT	Transfer Result	The result of the transfer operation (from the perspective of the sending entity): SUCS: Transfer successfully completed; all requested sequence counts were sent. CONL: connection lost during transfer CTMO: connection timed-out during establishment or transfer UNRE: destination node ID unreachable CRPT: transfer ended due to reception of corrupt or invalid data

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from "Start Sequence Count" to "Actual End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

CGRR: Cross-Grid Replication Request

This message is generated when StorageGRID attempts a cross-grid replication operation to replicate objects between buckets in a grid federation connection.

Code	Field	Description
CSIZ	Object Size	The size of the object in bytes. The CSIZ attribute was introduced in StorageGRID 11.8. As a result, cross-grid replication requests spanning a StorageGRID 11.7 to 11.8 upgrade might have an inaccurate total object size.
S3AI	S3 tenant account ID	The ID of the tenant account that owns the bucket from which the object is being replicated.
GFID	Grid federation connection ID	The ID of the grid federation connection being used for cross-grid replication.
OPER	CGR operation	The type of cross-grid replication operation that was attempted: • 0 = Replicate object • 1 = Replicate multipart object • 2 = Replicate delete marker
S3BK	S3 bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name.

Code	Field	Description
VSID	Version ID	The version ID of the specific version of an object that was being replicated.
RSLT	Result Code	Returns successful (SUCS) or general error (GERR).

EBDL: Empty Bucket Delete

The ILM scanner deleted an object in a bucket that is deleting all objects (performing an empty bucket operation).

Code	Field	Description
CSIZ	Object Size	The size of the object in bytes.
PATH	S3 Bucket/Key	The S3 bucket name and S3 key name.
SEGC	Container UUID	UUID of the container for the segmented object. This value is available only if the object is segmented.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
RSLT	Result of the delete operation	The result of event, process, or transaction. If is not relevant for a message, NONE is used rather than SUCS so that the message is not accidentally filtered.

EBKR: Empty Bucket Request

This message indicates a user sent a request to turn empty bucket on or off (that is, to delete bucket objects or to stop deleting objects).

Code	Field	Description
BUID	Bucket UUID	The bucket ID.
EBJS	Empty Bucket JSON Configuration	Contains the JSON representing the current Empty Bucket configuration.
S3AI	S3 tenant account ID	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.

ECMC: Missing Erasure-Coded Data Fragment

This audit message indicates that the system has detected a missing erasure-coded data fragment.

Code	Field	Description
VCMC	VCS ID	The name of the VCS that contains the missing chunk.
MCID	Chunk ID	The identifier of the missing erasure-coded fragment.
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this particular message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

ECOC: Corrupt Erasure-Coded Data Fragment

This audit message indicates that the system has detected a corrupt erasure-coded data fragment.

Code	Field	Description
VCCO	VCS ID	The name of the VCS that contains the corrupt chunk.
VLID	Volume ID	The RangeDB Volume that contains the corrupt erasure-coded fragment.
CCID	Chunk ID	The identifier of the corrupt erasure-coded fragment.
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this particular message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

ETAF: Security Authentication Failed

This message is generated when a connection attempt using Transport Layer Security (TLS) has failed.

Code	Field	Description
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection over which the authentication failed.
RUID	User Identity	A service dependent identifier representing the identity of the remote user.

Code	Field	Description
RSLT	Reason Code	The reason for the failure: SCNI: Secure connection establishment failed. CERM: Certificate was missing. CERT: Certificate was invalid. CERE: Certificate was expired. CERR: Certificate was revoked. CSGN: Certificate signature was invalid. CSGU: Certificate signer was unknown. UCRM: User credentials were missing. UCRI: User credentials were invalid. UCRU: User credentials were disallowed. TOUT: Authentication timed out.

When a connection is established to a secure service that uses TLS, the credentials of the remote entity are verified using the TLS profile and additional logic built into the service. If this authentication fails due to invalid, unexpected, or disallowed certificates or credentials, an audit message is logged. This enables queries for unauthorized access attempts and other security-related connection problems.

The message could result from a remote entity having an incorrect configuration, or from attempts to present invalid or disallowed credentials to the system. This audit message should be monitored to detect attempts to gain unauthorized access to the system.

GNRG: GNDS Registration

The CMN service generates this audit message when a service has updated or registered information about itself in the StorageGRID system.

Code	Field	Description
RSLT	Result	The result of the update request: • SUCS: Successful • SUNV: Service Unavailable • GERR: Other failure
GNID	Node ID	The node ID of the service that initiated the update request.
GNTP	Device Type	The grid node’s device type (for example, BLDR for an LDR service).

Code	Field	Description
GNDV	Device Model version	The string identifying the grid node's device model version in the DMDL bundle.
GNGP	Group	The group to which the grid node belongs (in the context of link costs and service-query ranking).
GNIA	IP Address	The grid node's IP address.

This message is generated whenever a grid node updates its entry in the Grid Nodes Bundle.

GNUR: GNDS Unregistration

The CMN service generates this audit message when a service has unregistered information about itself from the StorageGRID system.

Code	Field	Description
RSLT	Result	The result of the update request: • SUCS: Successful • SUNV: Service Unavailable • GERR: Other failure
GNID	Node ID	The node ID of the service that initiated the update request.

GTED: Grid Task Ended

This audit message indicates that the CMN service has finished processing the specified grid task and has moved the task to the Historical table. If the result is SUCS, ABRT, or ROLF, there will be a corresponding Grid Task Started audit message. The other results indicate that processing of this grid task never started.

Code	Field	Description
TSID	Task ID	This field uniquely identifies a generated grid task and allows the grid task to be managed over its lifecycle. **Note:** The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.

Code	Field	Description
RSLT	Result	The final status result of the grid task: • SUCS: The grid task completed successfully. • ABRT: The grid task was terminated without a rollback error. • ROLF: The grid task was terminated and was unable to complete the rollback process. • CANC: The grid task was canceled by the user before it was started. • EXPR: The grid task expired before it was started. • IVLD: The grid task was invalid. • AUTH: The grid task was unauthorized. • DUPL: The grid task was rejected as a duplicate.

GTST: Grid Task Started

This audit message indicates that the CMN service has started to process the specified grid task. The audit message immediately follows the Grid Task Submitted message for grid tasks initiated by the internal Grid Task Submission service and selected for automatic activation. For grid tasks submitted into the Pending table, this message is generated when the user starts the grid task.

Code	Field	Description
TSID	Task ID	This field uniquely identifies a generated grid task and allows the task to be managed over its lifecycle. **Note:** The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.
RSLT	Result	The result. This field has only one value: • SUCS: The grid task was started successfully.

GTSU: Grid Task Submitted

This audit message indicates that a grid task has been submitted to the CMN service.

Code	Field	Description
TSID	Task ID	Uniquely identifies a generated grid task and allows the task to be managed over its lifecycle. **Note:** The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.
TTYP	Task Type	The type of grid task.
TVER	Task Version	A number indicating the version of the grid task.
TDSC	Task Description	A human-readable description of the grid task.
VATS	Valid After Timestamp	The earliest time (UINT64 microseconds from January 1, 1970 - UNIX time) at which the grid task is valid.
VBTS	Valid Before Timestamp	The latest time (UINT64 microseconds from January 1, 1970 - UNIX time) at which the grid task is valid.
TSRC	Source	The source of the task: • TXTB: The grid task was submitted through the StorageGRID system as a signed text block. • GRID: The grid task was submitted through the internal Grid Task Submission Service.
ACTV	Activation Type	The type of activation: • AUTO: The grid task was submitted for automatic activation. • PEND: The grid task was submitted into the pending table. This is the only possibility for the TXTB source.
RSLT	Result	The result of the submission: • SUCS: The grid task was submitted successfully. • FAIL: The task has been moved directly to the historical table.

IDEL: ILM Initiated Delete

This message is generated when ILM starts the process of deleting an object.

The IDEL message is generated in either of these situations:

- **For objects in compliant S3 buckets**: This message is generated when ILM starts the process of auto-deleting an object because its retention period has expired (assuming the auto-delete setting is enabled

and legal hold is off).

- **For objects in non-compliant S3 buckets or Swift containers**. This message is generated when ILM starts the process of deleting an object because no placement instructions in the active ILM policies currently apply to the object.

Code	Field	Description
CBID	Content Block Identifier	The CBID of the object.
CMPA	Compliance: Auto delete	For objects in compliant S3 buckets only. 0 (false) or 1 (true), indicating whether a compliant object should be deleted automatically when its retention period ends, unless the bucket is under a legal hold.
CMPL	Compliance: Legal hold	For objects in compliant S3 buckets only. 0 (false) or 1 (true), indicating whether the bucket is currently under a legal hold.
CMPR	Compliance: Retention period	For objects in compliant S3 buckets only. The length of the object's retention period in minutes.
CTME	Compliance: Ingest time	For objects in compliant S3 buckets only. The object's ingest time. You can add the retention period in minutes to this value to determine when the object can be deleted from the bucket.
DMRK	Delete Marker Version ID	The version ID of the delete marker created when deleting an object from a versioned bucket. Operations on buckets don't include this field.
CSIZ	Content size	The size of the object in bytes.
LOCS	Locations	The storage location of object data within the StorageGRID system. The value for LOCS is "" if the object has no locations (for example, it has been deleted). CLEC: for erasure-coded objects, the erasure-coding profile ID and the erasure coding group ID that is applied to the object's data. CLDI: for replicated objects, the LDR node ID and the volume ID of the object's location. CLNL: ARC node ID of the object's location if the object data is archived.
PATH	S3 Bucket/Key or Swift Container/Object ID	The S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
RSLT	Result	The result of the ILM operation. SUCS: The ILM operation was successful.

Code	Field	Description
RULE	Rules Label	• If an object in a compliant S3 bucket is being deleted automatically because its retention period has expired, this field is blank. • If the object is being deleted because there are no more placement instructions that currently apply to the object, this field shows the human-readable label of the last ILM rule that applied to the object.
SGRP	Site (Group)	If present, the object was deleted at the site specified, which is not the site where the object was ingested.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was deleted. Operations on buckets and objects in unversioned buckets don't include this field.

LKCU: Overwritten Object Cleanup

This message is generated when StorageGRID removes an overwritten object that previously required cleanup to free up storage space. An object is overwritten when an S3 or Swift client writes an object to a path already containing a object. The removal process occurs automatically and in the background.

Code	Field	Description
CSIZ	Content size	The size of the object in bytes.
LTYP	Type of cleanup	*Internal use only.*
LUID	Removed Object UUID	The identifier of the object that was removed.
PATH	S3 Bucket/Key or Swift Container/Object ID	The S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
SEGC	Container UUID	UUID of the container for the segmented object. This value is available only if the object is segmented.
UUID	Universally Unique Identifier	The identifier of the object that still exists. This value is available only if the object has not been deleted.

LLST: Location Lost

This message is generated whenever a location for an object copy (replicated or erasure-

coded) can't be found.

Code	Field	Description
CBIL	CBID	The affected CBID.
ECPR	Erasure-Coding Profile	For erasure-coded object data. The ID of the erasure-coding profile used.
LTYP	Location Type	CLDI (Online): For replicated object data CLEC (Online): For erasure-coded object data CLNL (Nearline): For archived replicated object data
NOID	Source Node ID	The node ID on which the locations were lost.
PCLD	Path to replicated object	The complete path to the disk location of the lost object data. Only returned when LTYP has a value of CLDI (that is, for replicated objects). Takes the form `/var/local/rangedb/2/p/13/13/00oJs6X%{h{U)SeUFxE@`
RSLT	Result	Always NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.
TSRC	Triggering Source	USER: User triggered SYST: System triggered
UUID	Universally Unique ID	The identifier of the affected object in the StorageGRID system.

MGAU: Management audit message

The Management category logs user requests to the Management API. Every request that is not a GET or HEAD request to the API logs a response with the username, IP, and type of request to the API.

Code	Field	Description
MDIP	Destination IP Address	The server (destination) IP address.
MDNA	Domain name	The host domain name.
MPAT	Request PATH	The request path.

Code	Field	Description
MPQP	Request query parameters	The query parameters for the request.
MRBD	Request body	The content of the request body. While the response body is logged by default, the request body is logged in certain cases when the response body is empty. Because the following information is not available in the response body, it is taken from the request body for the following POST methods: • Username and account ID in **POST authorize** • New subnets configuration in **POST /grid/grid-networks/update** • New NTP servers in **POST /grid/ntp-servers/update** • Decommissioned server IDs in **POST /grid/servers/decommission** **Note:** Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).
MRMD	Request method	The HTTP request method: • POST • PUT • DELETE • PATCH
MRSC	Response code	The response code.
MRSP	Response body	The content of the response (the response body) is logged by default. **Note:** Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).
MSIP	Source IP address	The client (source) IP address.
MUUN	User URN	The URN (uniform resource name) of the user who sent the request.
RSLT	Result	Returns successful (SUCS) or the error reported by the backend.

OLST: System Detected Lost Object

This message is generated when the DDS service can’t locate any copies of an object within the StorageGRID system.

Code	Field	Description
CBID	Content Block Identifier	The CBID of the lost object.
NOID	Node ID	If available, the last known direct or near-line location of the lost object. It is possible to have just the Node ID without a Volume ID if the volume information is not available.
PATH	S3 Bucket/Key or Swift Container/Object ID	If available, the S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.
UUID	Universally Unique ID	The identifier of the lost object within the StorageGRID system.
VOLI	Volume ID	If available, the Volume ID of the Storage Node or Archive Node for the last known location of the lost object.

ORLM: Object Rules Met

This message is generated when the object is successfully stored and copied as specified by the ILM rules.

> The ORLM message is not generated when an object is successfully stored by the default Make 2 Copies rule if another rule in the policy uses the Object Size advanced filter.

Code	Field	Description
BUID	Bucket Header	Bucket ID field. Used for internal operations. Appears only if STAT is PRGD.
CBID	Content Block Identifier	The CBID of the object.
CSIZ	Content size	The size of the object in bytes.

Code	Field	Description
LOCS	Locations	The storage location of object data within the StorageGRID system. The value for LOCS is "" if the object has no locations (for example, it has been deleted). CLEC: for erasure-coded objects, the erasure-coding profile ID and the erasure coding group ID that is applied to the object's data. CLDI: for replicated objects, the LDR node ID and the volume ID of the object's location. CLNL: ARC node ID of the object's location if the object data is archived.
PATH	S3 Bucket/Key or Swift Container/Object ID	The S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
RSLT	Result	The result of the ILM operation. SUCS: The ILM operation was successful.
RULE	Rules Label	The human-readable label given to the ILM rule applied to this object.
SEGC	Container UUID	UUID of the container for the segmented object. This value is available only if the object is segmented.
SGCB	Container CBID	CBID of the container for the segmented object. This value is available only for segmented and multipart objects.
STAT	Status	The status of ILM operation. DONE: ILM operations against the object have completed. DFER: The object has been marked for future ILM re-evaluation. PRGD: The object has been deleted from the StorageGRID system. NLOC: The object data can no longer be found in the StorageGRID system. This status might indicate that all copies of object data are missing or damaged.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of a new object created in a versioned bucket. Operations on buckets and objects in unversioned buckets don't include this field.

The ORLM audit message can be issued more than once for a single object. For instance, it is issued

whenever one of the following events occur:

- ILM rules for the object are satisfied forever.
- ILM rules for the object are satisfied for this epoch.
- ILM rules have deleted the object.
- The background verification process detects that a copy of replicated object data is corrupt. The StorageGRID system performs an ILM evaluation to replace the corrupt object.

Related information

- Object ingest transactions
- Object delete transactions

OVWR: Object Overwrite

This message is generated when an external (client-requested) operation causes one object to be overwritten by another object.

Code	Field	Description
CBID	Content Block Identifier (new)	The CBID for the new object.
CSIZ	Previous Object Size	The size, in bytes, of the object being overwritten.
OCBD	Content Block Identifier (previous)	The CBID for the previous object.
UUID	Universally Unique ID (new)	The identifier of the new object within the StorageGRID system.
OUID	Universally Unique ID (previous)	The identifier for the previous object within the StorageGRID system.
PATH	S3 or Swift Object Path	The S3 or Swift object path used for both the previous and new object
RSLT	Result Code	Result of the Object Overwrite transaction. Result is always: SUCS: Successful
SGRP	Site (Group)	If present, the overwritten object was deleted at the site specified, which is not the site where the overwritten object was ingested.

S3SL: S3 Select request

This message logs a completion after an S3 Select request has been returned to the client. The S3SL message can include error message and error code details. The request might not have been successful.

Code	Field	Description
BYSC	Bytes Scanned	Number of bytes scanned (received) from Storage Nodes. BYSC and BYPR are likely to be different if the object is compressed. If the object is compressed BYSC would have the compressed byte count and BYPR would be the bytes after decompression.
BYPR	Bytes Processed	Number of bytes processed. Indicates how many bytes of "Bytes Scanned" were actually processed or acted upon by an S3 Select job.
BYRT	Bytes Returned	Number of bytes that an S3 Select job returned to the client.
REPR	Records Processed	Number of records or rows that an S3 Select job received from Storage Nodes.
RERT	Records Returned	Number of records or rows an S3 Select job returned to the client.
JOFI	Job Finished	Indicates if the S3 Select job finished processing or not. If this is false, then the job failed to finish and the error fields will likely have data in them. The client might have received partial results, or no results at all.
REID	Request ID	Identifier for the S3 Select request.
EXTM	Execution Time	The time, in seconds, it took for the S3 Select Job to complete.
ERMG	Error Message	Error message that the S3 Select job generated.
ERTY	Error Type	Error type that the S3 Select job generated.
ERST	Error Stacktrace	Error Stacktrace that the S3 Select job generated.
S3BK	S3 bucket	The S3 bucket name.
S3AK	S3 Access Key ID (request sender)	The S3 access key ID for the user that sent the request.
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request.

Code	Field	Description
S3KY	S3 Key	The S3 key name, not including the bucket name.

SADD: Security Audit Disable

This message indicates that the originating service (node ID) has turned off audit message logging; audit messages are no longer being collected or delivered.

Code	Field	Description
AETM	Enable Method	The method used to disable the audit.
AEUN	User Name	The user name that executed the command to disable audit logging.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.

The message implies that logging was previously enabled, but has now been disabled. This is typically used only during bulk ingest to improve system performance. Following the bulk activity, auditing is restored (SADE) and the capability to disable auditing is then permanently blocked.

SADE: Security Audit Enable

This message indicates that the originating service (node ID) has restored audit message logging; audit messages are again being collected and delivered.

Code	Field	Description
AETM	Enable Method	The method used to enable the audit.
AEUN	User Name	The user name that executed the command to enable audit logging.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.

The message implies that logging was previously disabled (SADD), but has now been restored. This is typically only used during bulk ingest to improve system performance. Following the bulk activity, auditing is restored and the capability to disable auditing is then permanently blocked.

SCMT: Object Store Commit

Grid content is not made available or recognized as stored until it has been committed (meaning it has been stored persistently). Persistently stored content has been completely written to disk, and has passed related integrity checks. This message is issued when a content block is committed to storage.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block committed to permanent storage.
RSLT	Result Code	Status at the time the object was stored to disk: SUCS: Object successfully stored.

This message means a given content block has been completely stored and verified, and can now be requested. It can be used to track data flow within the system.

SDEL: S3 DELETE

When an S3 client issues a DELETE transaction, a request is made to remove the specified object or bucket, or to remove a bucket/object subresource. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don’t include this field.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the deleted object in bytes. Operations on buckets don’t include this field.
DMRK	Delete Marker Version ID	The version ID of the delete marker created when deleting an object from a versioned bucket. Operations on buckets don’t include this field.
GFID	Grid Federation Connection ID	The connection ID of the grid federation connection associated with a cross-grid replication delete request. Only included in audit logs on the destination grid.
GFSA	Grid Federation Source Account ID	The account ID of the tenant on the source grid for a cross-grid replication delete request. Only included in audit logs on the destination grid.

Code	Field	Description
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field). `x-amz-bypass-governance-retention` is automatically included if it is present in the request.
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	Result of the DELETE transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don’t include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.

Code	Field	Description
SGRP	Site (Group)	If present, the object was deleted at the site specified, which is not the site where the object was ingested.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: `urn:sgws:identity::03393893651506583485:root` Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUDM	Universally Unique Identifier for a Delete Marker	The identifier of a delete marker. Audit log messages specify either UUDM or UUID, where UUDM indicates a delete marker created as a result of an object delete request, and UUID indicates an object.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was deleted. Operations on buckets and objects in unversioned buckets don't include this field.

SGET: S3 GET

When an S3 client issues a GET transaction, a request is made to retrieve an object or list the objects in a bucket, or to remove a bucket/object subresource. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don't include this field.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.

Code	Field	Description
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field).
LITY	ListObjectsV2	A *v2 format* response was requested. For details, see AWS ListObjectsV2. For GET bucket operations only.
NCHD	Number of Children	Includes keys and common prefixes. For GET bucket operations only.
RANG	Range Read	For range read operations only. Indicates the range of bytes that was read by this request. The value after the slash (/) shows the size of the entire object.
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.

Code	Field	Description
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: `urn:sgws:identity::03393893651506583485:root` Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
TRNC	Truncated or Not Truncated	Set to false if all results were returned. Set to true if more results are available to return. For GET bucket operations only.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets don't include this field.

SHEA: S3 HEAD

When an S3 client issues a HEAD transaction, a request is made to check for the existence of an object or bucket and retrieve the metadata about an object. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don't include this field.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.

Code	Field	Description
CSIZ	Content Size	The size of the checked object in bytes. Operations on buckets don’t include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field).
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don’t include this field.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: `urn:sgws:identity::03393893651506583485:root` Empty for anonymous requests.

Code	Field	Description
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets don't include this field.

SPOS: S3 POST

When an S3 client issues a POST Object request, this message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field). (Not expected for SPOS).
RSLT	Result Code	Result of the RestoreObject request. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.

Code	Field	Description
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don’t include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable. Set to "select" for an S3 Select operation.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SRCF	Subresource Configuration	Restore information.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: `urn:sgws:identity::03393893651506583485:root` Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.

Code	Field	Description
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets don't include this field.

SPUT: S3 PUT

When an S3 client issues a PUT transaction, a request is made to create a new object or bucket, or to remove a bucket/object subresource. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don't include this field.
CMPS	Compliance Settings	The compliance settings used when creating the bucket, if present in the request (truncated to the first 1024 characters).
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets don't include this field.
GFID	Grid Federation Connection ID	The connection ID of the grid federation connection associated with a cross-grid replication PUT request. Only included in audit logs on the destination grid.
GFSA	Grid Federation Source Account ID	The account ID of the tenant on the source grid for a cross-grid replication PUT request. Only included in audit logs on the destination grid.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field). `x-amz-bypass-governance-retention` is automatically included if it is present in the request.

Code	Field	Description
LKEN	Object Lock Enabled	Value of the request header `x-amz-bucket-object-lock-enabled`, if present in the request.
LKLH	Object Lock Legal Hold	Value of the request header `x-amz-object-lock-legal-hold`, if present in the PutObject request.
LKMD	Object Lock Retention Mode	Value of the request header `x-amz-object-lock-mode`, if present in the PutObject request.
LKRU	Object Lock Retain Until Date	Value of the request header `x-amz-object-lock-retain-until-date`, if present in the PutObject request.
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	Result of the PUT transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don’t include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.

Code	Field	Description
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SRCF	Subresource Configuration	The new subresource configuration (truncated to the first 1024 characters).
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: `urn:sgws:identity::03393893651506583485:root` Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
ULID	Upload ID	Included only in SPUT messages for CompleteMultipartUpload operations. Indicates that all parts have been uploaded and assembled.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of a new object created in a versioned bucket. Operations on buckets and objects in unversioned buckets don’t include this field.
VSST	Versioning State	The new versioning state of a bucket. Two states are used: "enabled" or "suspended." Operations on objects don’t include this field.

SREM: Object Store Remove

This message is issued when content is removed from persistent storage and is no longer accessible through regular APIs.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block deleted from permanent storage.
RSLT	Result Code	Indicates the result of the content removal operations. The only defined value is: SUCS: Content removed from persistent storage

This audit message means a given content block has been deleted from a node and can no longer be requested directly. The message can be used to track the flow of deleted content within the system.

SUPD: S3 Metadata Updated

This message is generated by the S3 API when an S3 client updates the metadata for an ingested object. The message is issued by the server if the metadata update is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don't include this field.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request, when updating a bucket's compliance settings.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field).
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.

Code	Field	Description
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: `urn:sgws:identity::03393893651506583485:root` Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object whose metadata was updated. Operations on buckets and objects in unversioned buckets don’t include this field.

SVRF: Object Store Verify Fail

This message is issued whenever a content block fails the verification process. Each time replicated object data is read from or written to disk, several verification and integrity checks are performed to ensure the data sent to the requesting user is identical to the data originally ingested into the system. If any of these checks fail, the system automatically quarantines the corrupt replicated object data to prevent it from being retrieved again.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block which failed verification.
RSLT	Result Code	Verification failure type: CRCF: Cyclic redundancy check (CRC) failed. HMAC: Hash-based message authentication code (HMAC) check failed. EHSH: Unexpected encrypted content hash. PHSH: Unexpected original content hash. SEQC: Incorrect data sequence on disk. PERR: Invalid structure of disk file. DERR: Disk error. FNAM: Bad file name.

This message should be monitored closely. Content verification failures can indicate impending hardware failures.

To determine what operation triggered the message, see the value of the AMID (Module ID) field. For example, an SVFY value indicates that the message was generated by the Storage Verifier module, that is, background verification, and STOR indicates that the message was triggered by content retrieval.

SVRU: Object Store Verify Unknown

The LDR service's Storage component continuously scans all copies of replicated object data in the object store. This message is issued when an unknown or unexpected copy of replicated object data is detected in the object store and moved to the quarantine directory.

Code	Field	Description
FPTH	File Path	The file path of the unexpected object copy.
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

The SVRU: Object Store Verify Unknown audit message should be monitored closely. It means unexpected copies of object data were detected in the object store. This situation should be investigated immediately to determine how theses copies were created, because it can indicate impending hardware failures.

SYSD: Node Stop

When a service is stopped gracefully, this message is generated to indicate the shutdown was requested. Typically this message is sent only after a subsequent restart, because the audit message queue is not cleared before shutdown. Look for the SYST message, sent at the beginning of the shutdown sequence, if the service has not restarted.

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shutdown.

The message does not indicate if the host server is being stopped, only the reporting service. The RSLT of a SYSD can’t indicate a "dirty" shutdown, because the message is generated only by "clean" shutdowns.

SYST: Node Stopping

When a service is gracefully stopped, this message is generated to indicate the shutdown was requested and that the service has initiated its shutdown sequence. SYST can be used to determine if the shutdown was requested, before the service is restarted (unlike SYSD, which is typically sent after the service restarts.)

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shutdown.

The message does not indicate if the host server is being stopped, only the reporting service. The RSLT code of a SYST message can’t indicate a "dirty" shutdown, because the message is generated only by "clean" shutdowns.

SYSU: Node Start

When a service is restarted, this message is generated to indicate if the previous shutdown was clean (commanded) or disorderly (unexpected).

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shut down. DSDN: System was not cleanly shut down. VRGN: System was started for the first time after server installation (or re-installation).

The message does not indicate if the host server was started, only the reporting service. This message can be

used to:

- Detect discontinuity in the audit trail.
- Determine if a service is failing during operation (as the distributed nature of the StorageGRID system can mask these failures). Server Manager restarts a failed service automatically.

WDEL: Swift DELETE

When a Swift client issues a DELETE transaction, a request is made to remove the specified object or container. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on containers don't include this field.
CSIZ	Content Size	The size of the deleted object in bytes. Operations on containers don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field).
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	Result of the DELETE transaction. Result is always: SUCS: Successful
SAIP	IP address of requesting client	The IP address of the client application that made the request.
SGRP	Site (Group)	If present, the object was deleted at the site specified, which is not the site where the object was ingested.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.

Code	Field	Description
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers don’t include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

WGET: Swift GET

When a Swift client issues a GET transaction, a request is made to retrieve an object, list the objects in a container, or list the containers in an account. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on accounts and containers don’t include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers don’t include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field).
RSLT	Result Code	Result of the GET transaction. Result is always SUCS: successful
SAIP	IP address of requesting client	The IP address of the client application that made the request.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.

Code	Field	Description
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name. Operations on accounts don't include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers don't include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

WHEA: Swift HEAD

When a Swift client issues a HEAD transaction, a request is made to check for the existence of an account, container, or object, and retrieve any relevant metadata. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on accounts and containers don't include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field).
RSLT	Result Code	Result of the HEAD transaction. Result is always: SUCS: successful
SAIP	IP address of requesting client	The IP address of the client application that made the request.
TIME	Time	Total processing time for the request in microseconds.

Code	Field	Description
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name. Operations on accounts don't include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers don't include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

WPUT: Swift PUT

When a Swift client issues a PUT transaction, a request is made to create a new object or container. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on containers don't include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on containers don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. `X-Forwarded-For` is automatically included if it is present in the request and if the `X-Forwarded-For` value is different from the request sender IP address (SAIP audit field).
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	Result of the PUT transaction. Result is always: SUCS: successful

Code	Field	Description
SAIP	IP address of requesting client	The IP address of the client application that made the request.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers don’t include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.