# Monitor StorageGRID system

StorageGRID 11.7

NetApp
March 05, 2024

This PDF was generated from https://docs.netapp.com/us-en/storagegrid-117/monitor/index.html on March 05, 2024. Always check docs.netapp.com for the latest.

# Table of Contents

Monitor StorageGRID system . . . 1
Monitor a StorageGRID system: Overview . . . 1
View and manage the dashboard . . . 1
View the Nodes page . . . 4
Information you should monitor regularly . . . 37
Alerts and alarms . . . 66
Log files reference . . . 150
Configure audit messages and log destinations . . . 168
Use an external syslog server . . . 172
Use SNMP monitoring . . . 183
Collect additional StorageGRID data . . . 198

# Monitor StorageGRID system

## Monitor a StorageGRID system: Overview

Use these instructions to monitor a StorageGRID system and to assess issues that might occur.

These instructions describe how to use the Grid Manager to monitor a StorageGRID system. You will learn which information you should monitor regularly, how to manage alerts and legacy alarms, how to use SNMP for monitoring, and how to obtain additional StorageGRID data, including metrics and diagnostics.

**Before you begin**

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

## 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**.

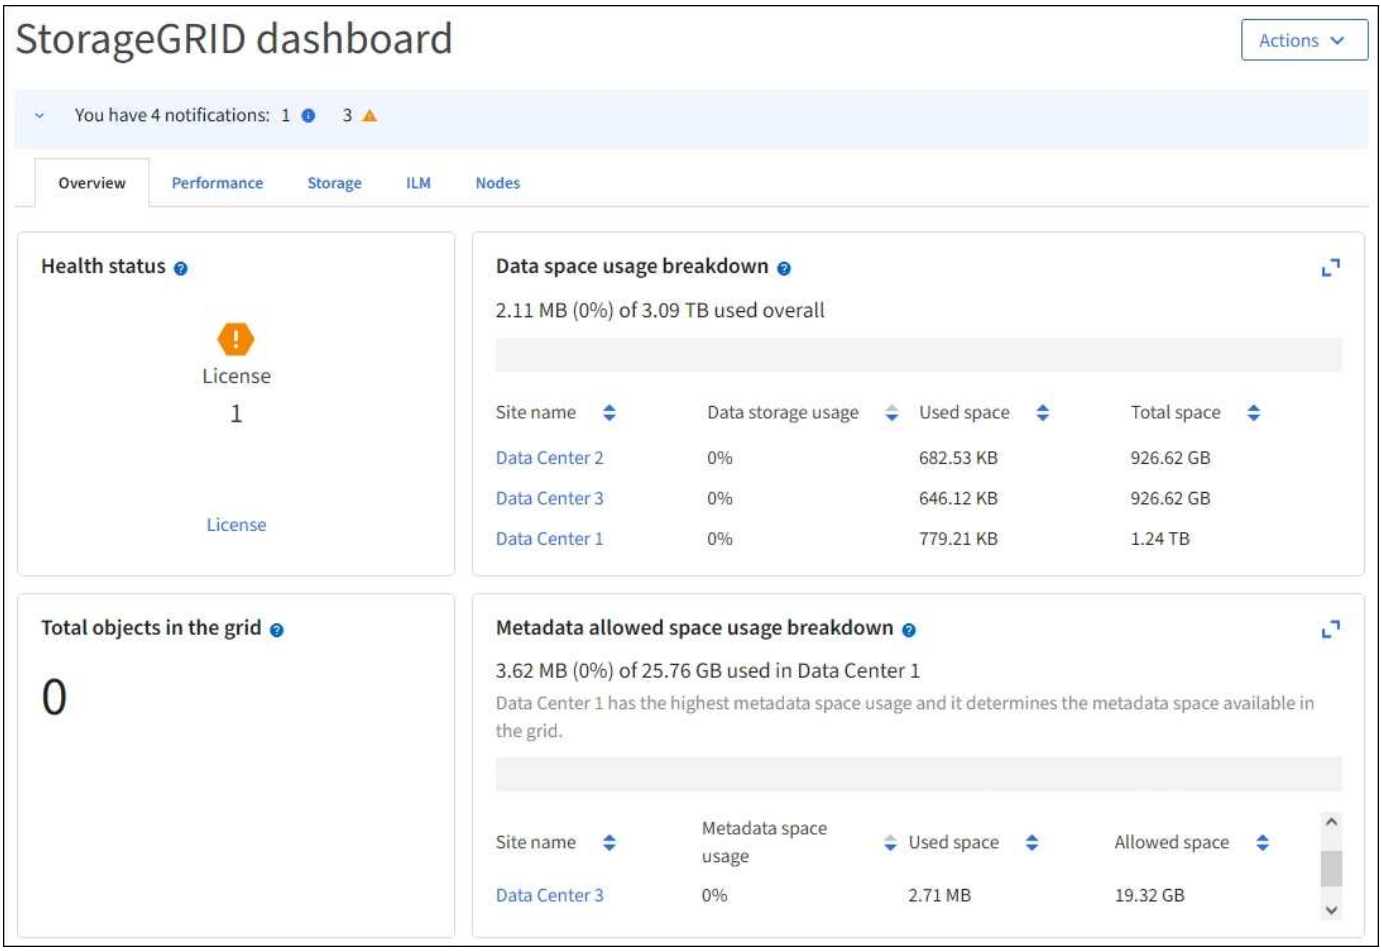


## View the dashboard

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. |
| 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. |

| Tab on system-provided dashboard | Contains |
| --- | --- |
| 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**.

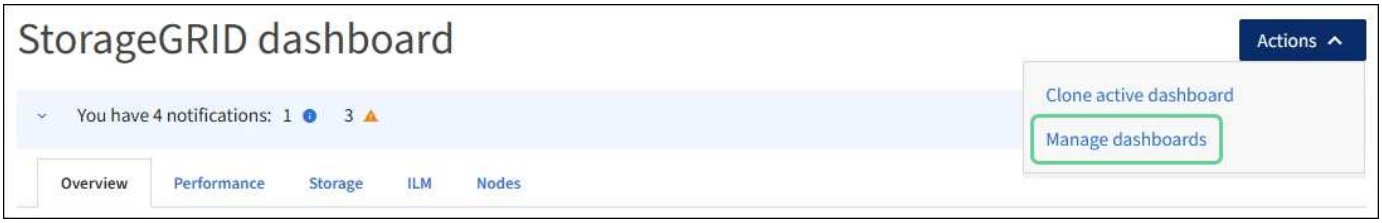


## Configure dashboards

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

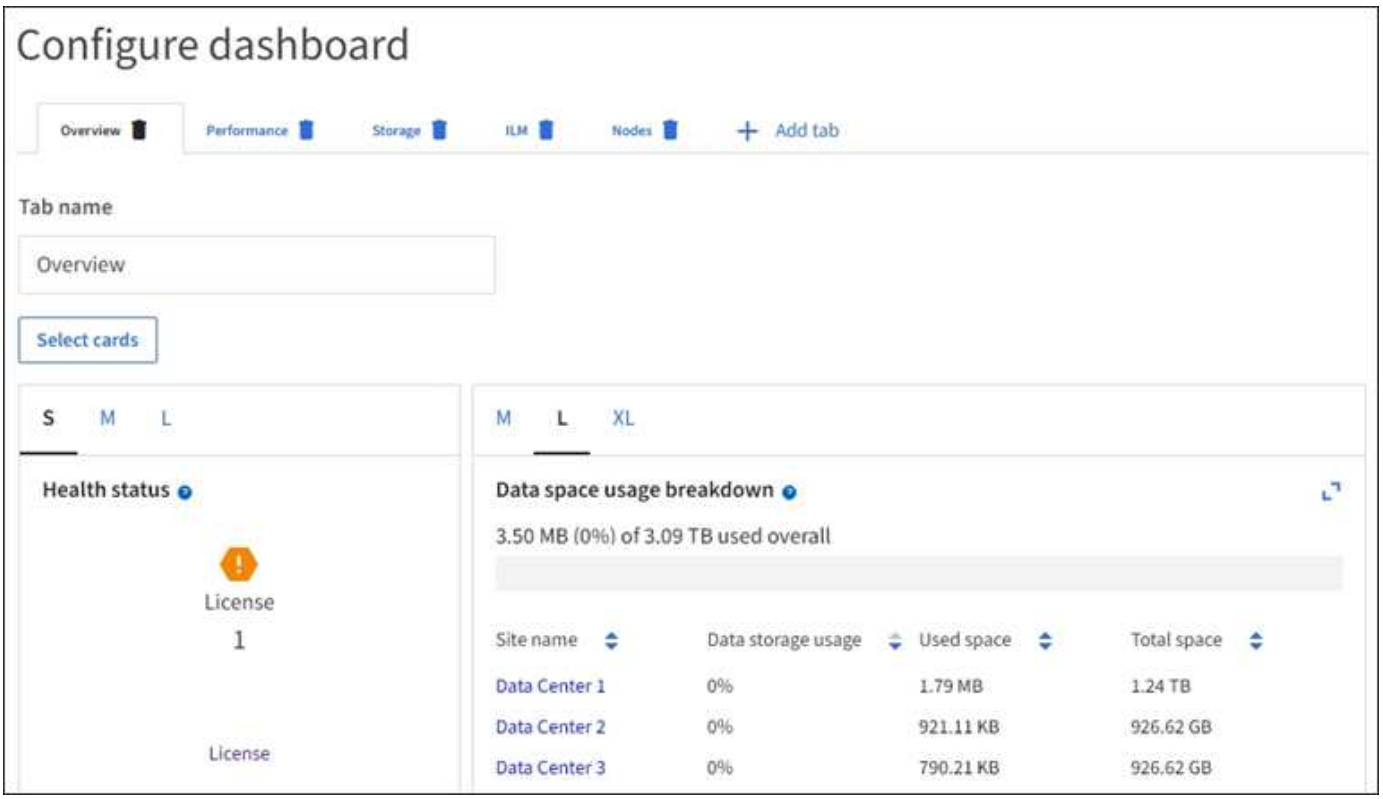


# View the Nodes page

## 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**.

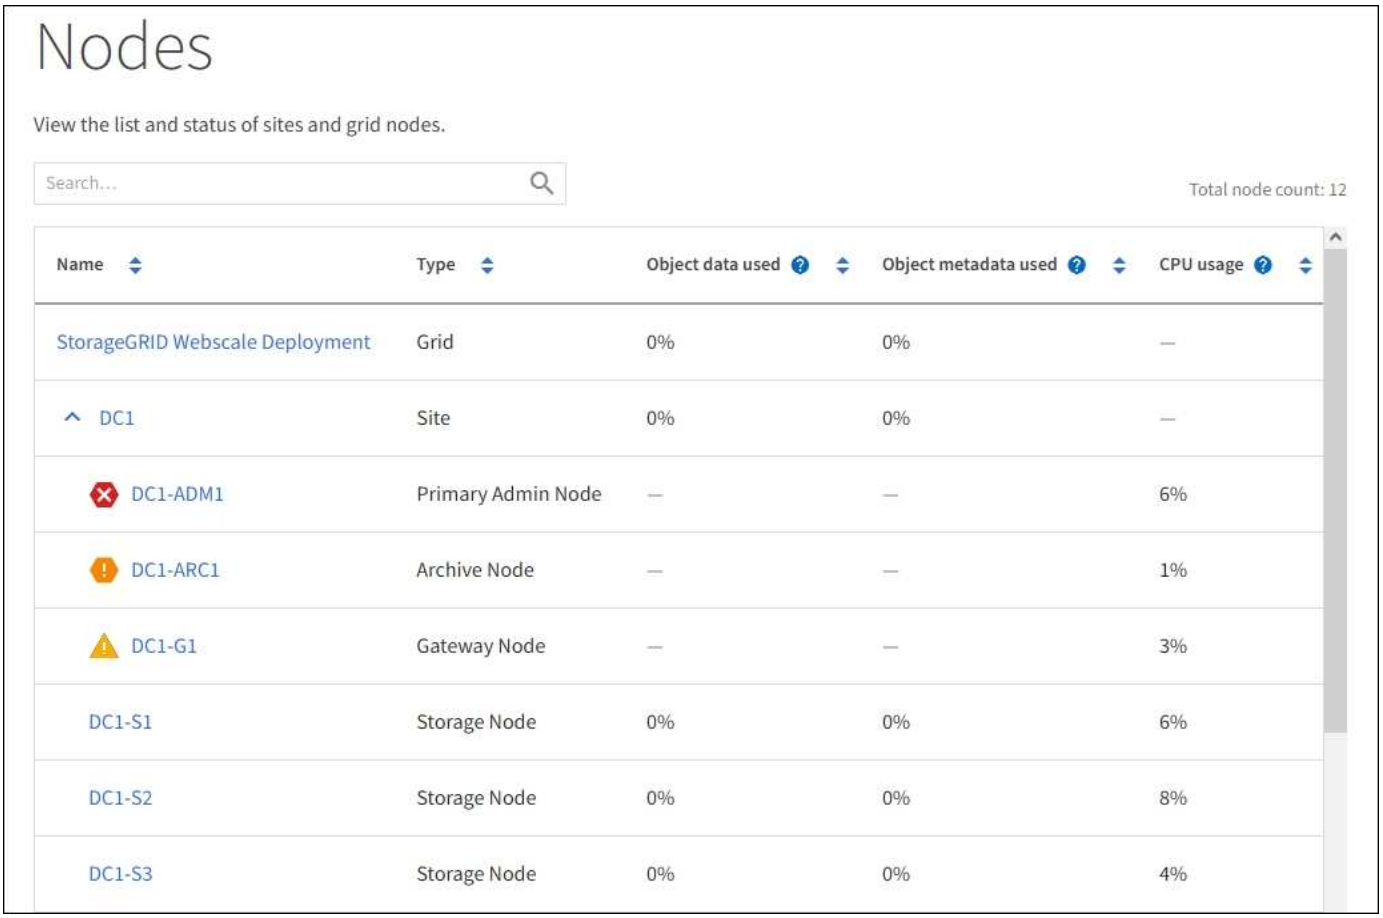


### Connection state icons

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**<br><br>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.<br><br>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.<br><br>For example, you might need to restart a service that has stopped or restart the host for the node.<br><br>**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**<br><br>For an expected reason, node is not connected to grid.<br><br>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.<br><br>Based on the underlying issue, these nodes often go back online with no intervention. | Determine if any alerts are affecting this node.<br><br>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.

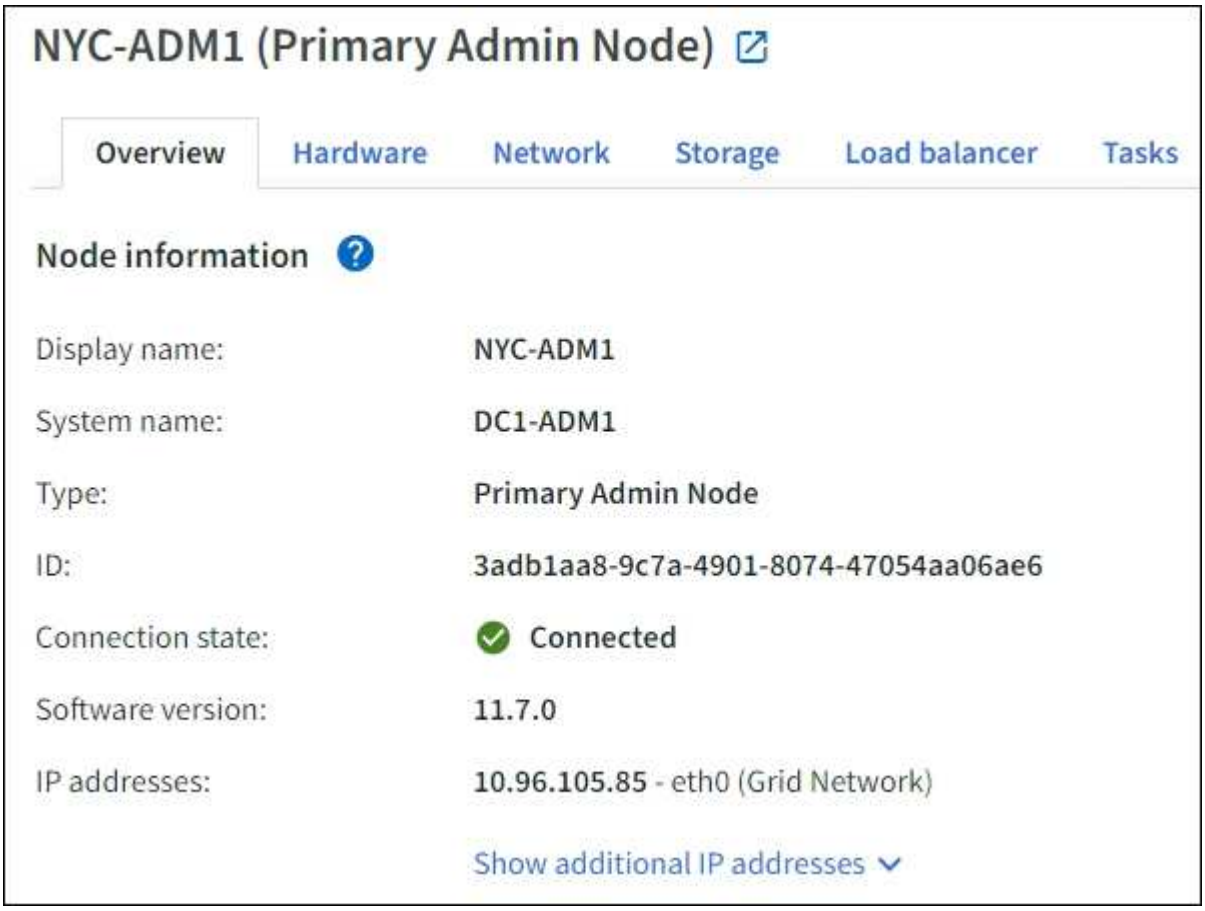


The overview information for a node includes the following:

- **Display name** (shown only if the node has been renamed): The current display name for the node. Use the Rename grid, site, 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 (for both archiving to the cloud using the S3 API and archiving to tape using TSM middleware) is deprecated and will be removed in a future release. Moving objects from an Archive Node to an external archival storage system 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 Hardware tab is shown for all nodes.

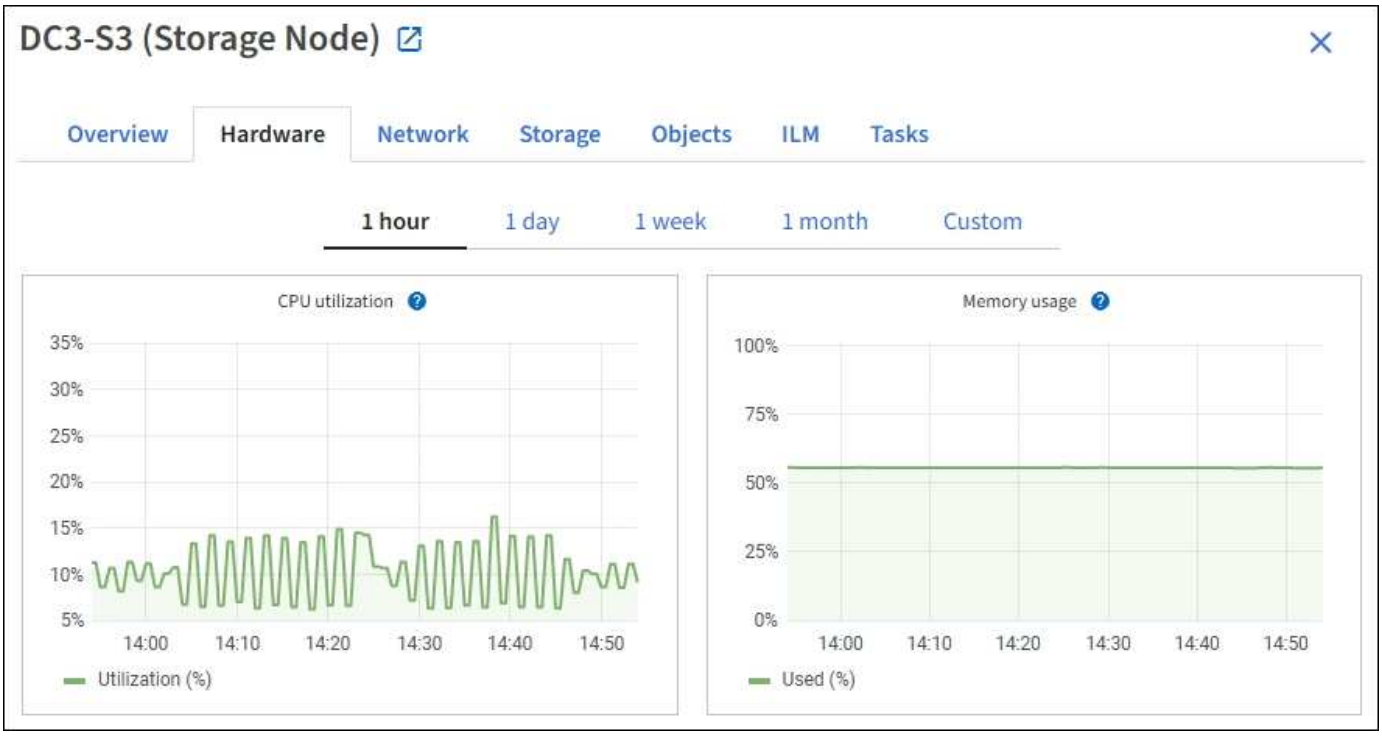


To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.

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

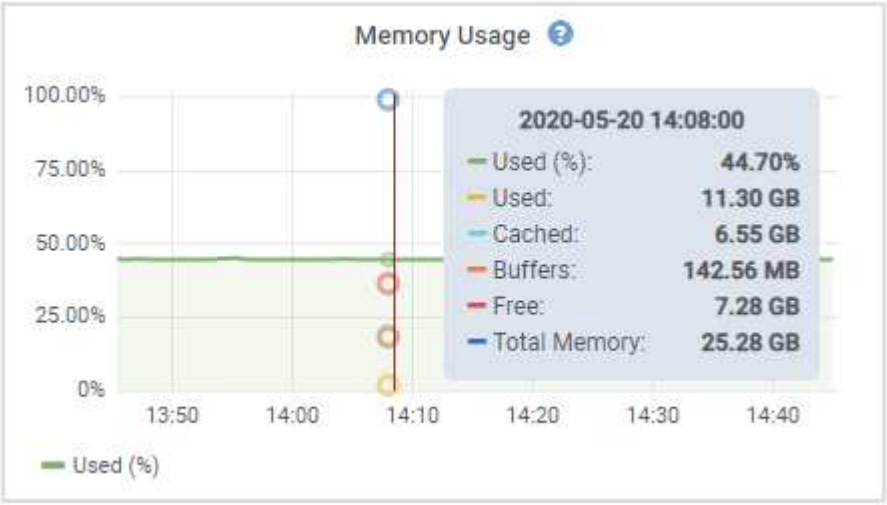


If the node is an appliance node, this tab also includes a section with more information about the appliance hardware.

### 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.

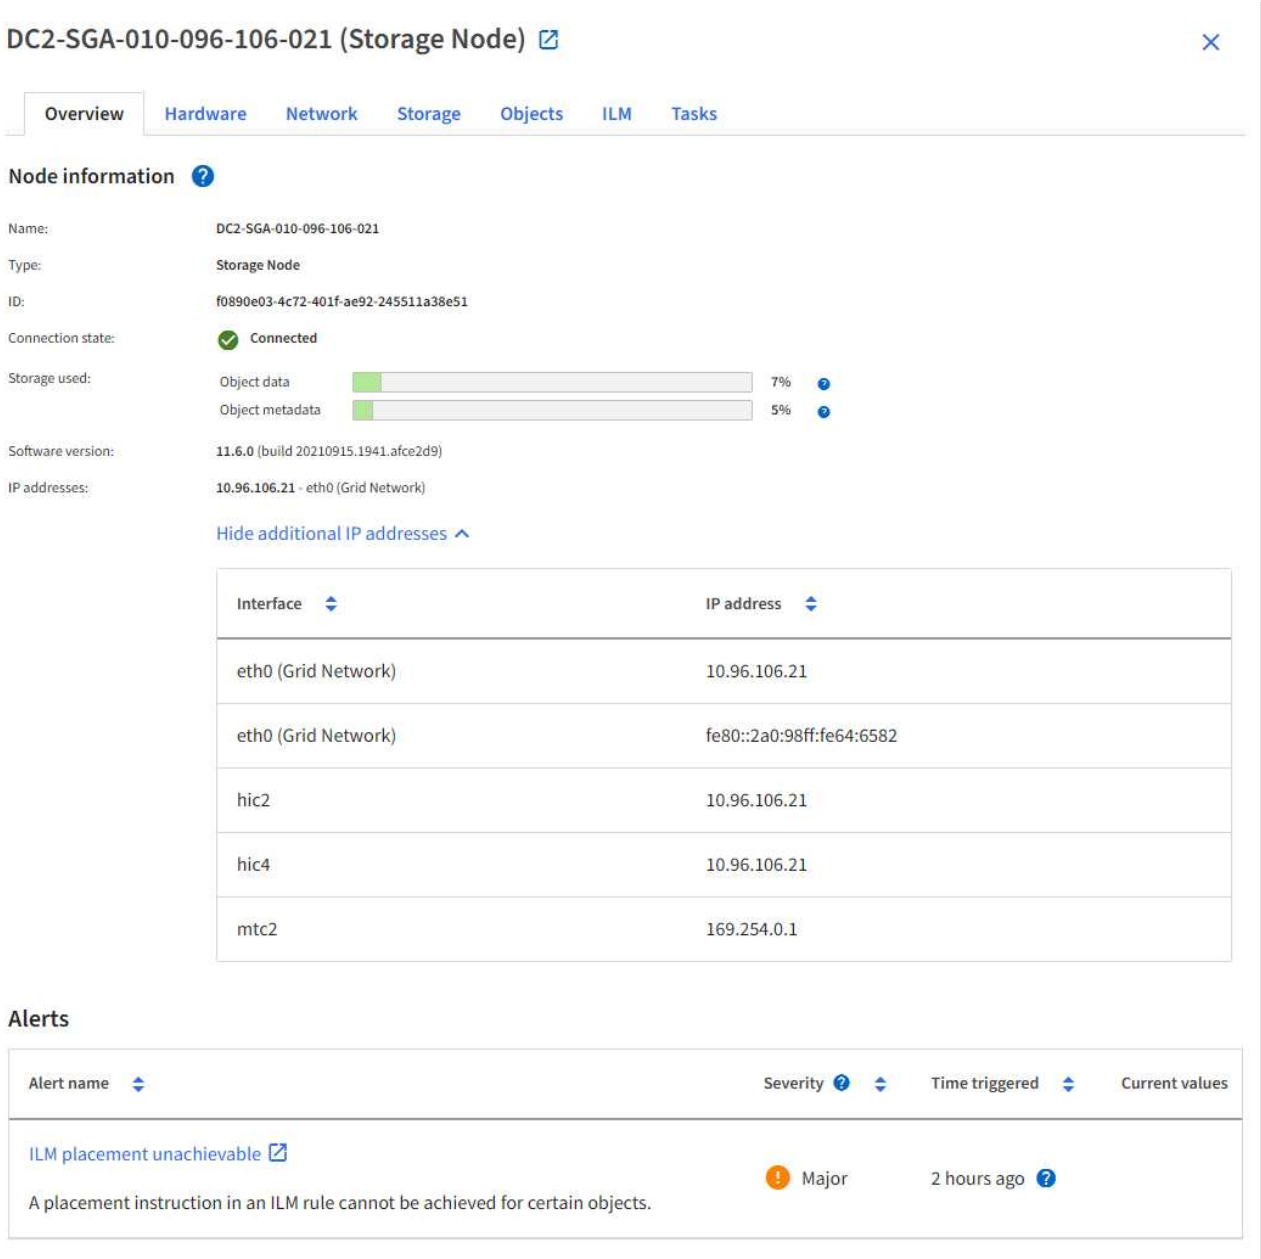


   The Alerts section of the Overview tab displays any active alerts for the node.

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.

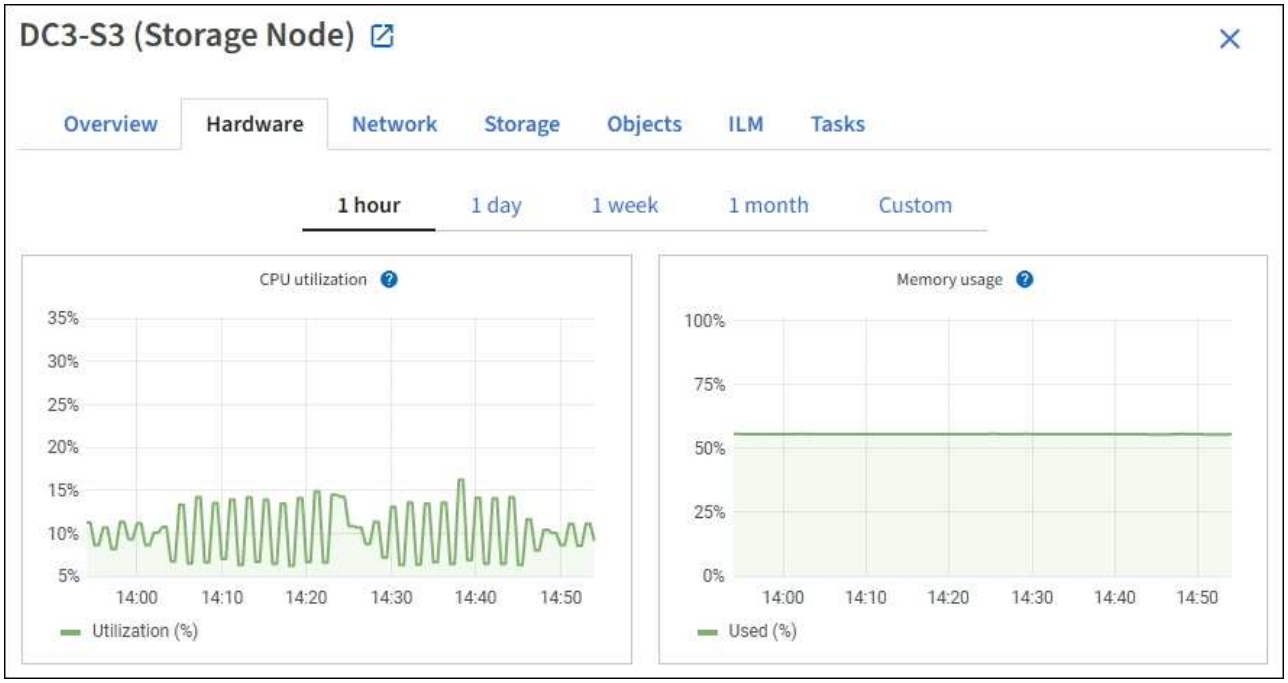


   b. Scroll down to view the table of components for the appliance. This table contains information such as the model name of the appliance; controller names, serial numbers, and IP addresses; and the status of each component.

> 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.

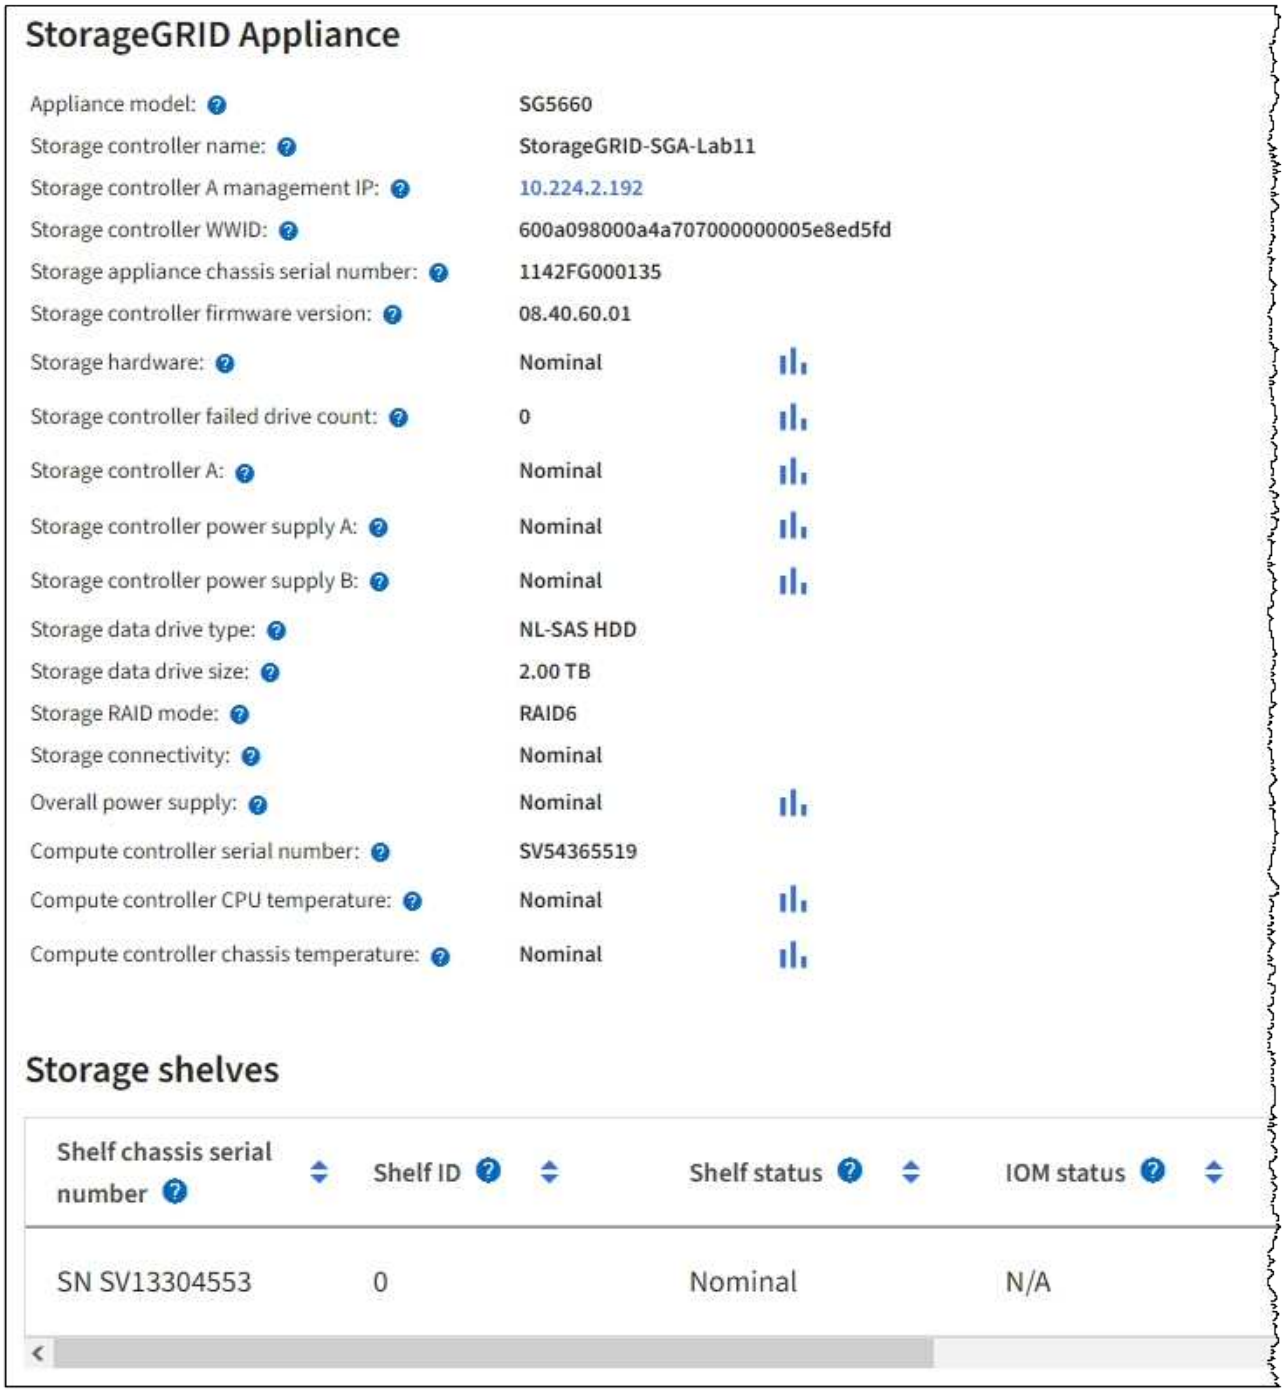


| Field in the Appliance table | Description |
|---|---|
| Appliance model | The model number for this StorageGRID appliance shown in SANtricity OS. |
| Storage controller name | The name for this StorageGRID appliance shown in SANtricity OS. |
| Storage controller A management IP | IP address for management 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.<br><br>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.<br><br>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.<br><br>**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.<br><br>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.<br><br>• 99: Storage controller shelf<br>• 0: First expansion shelf<br>• 1: Second expansion shelf<br><br>**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.

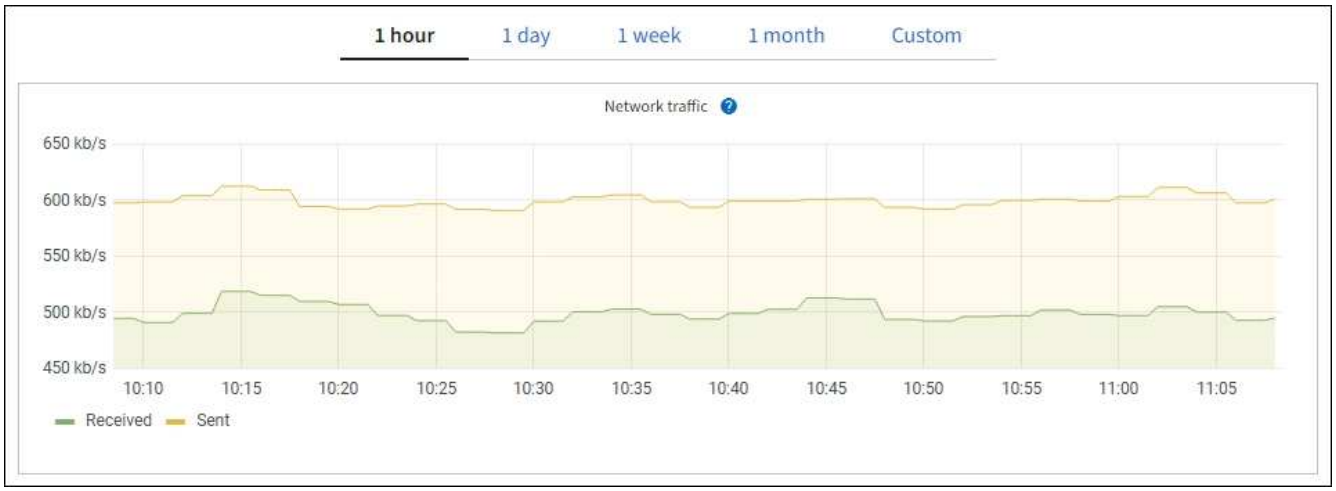


   a. Review the Network Interfaces section.

**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.

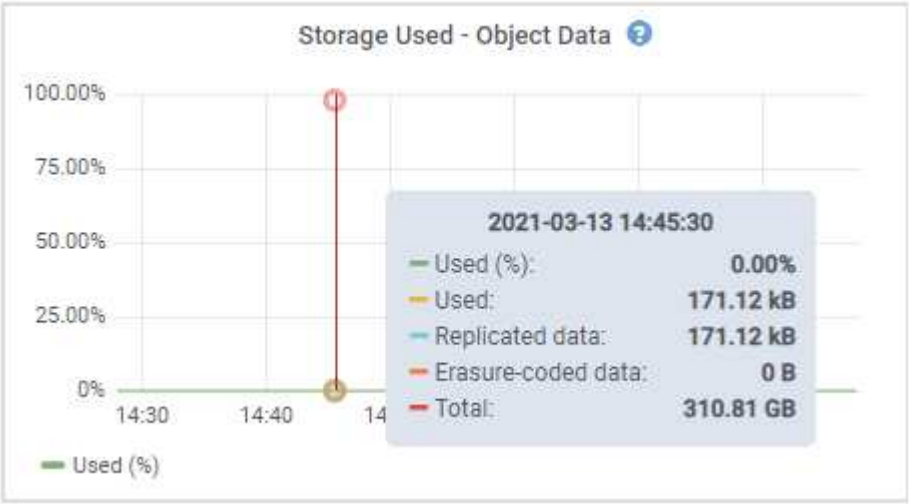


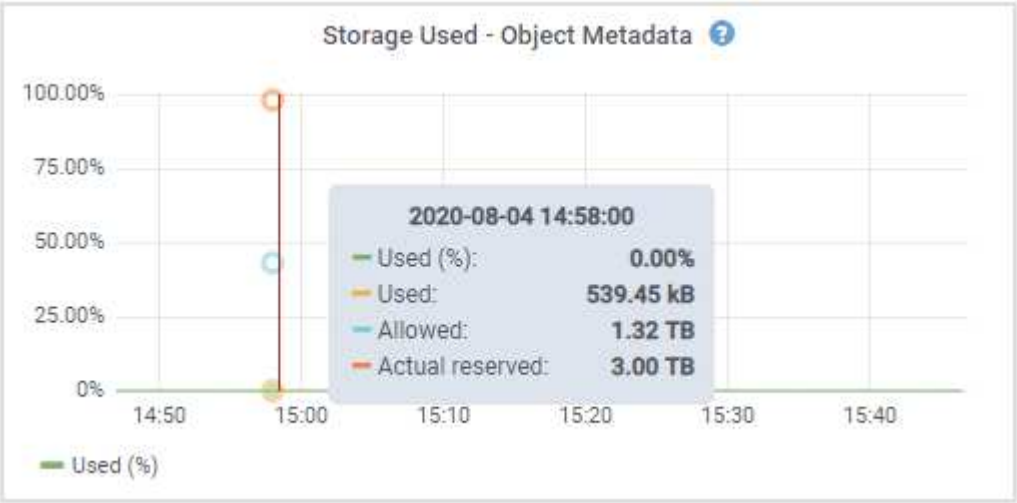


a. Scroll down to view the amounts of available storage for each volume and object store.

   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.

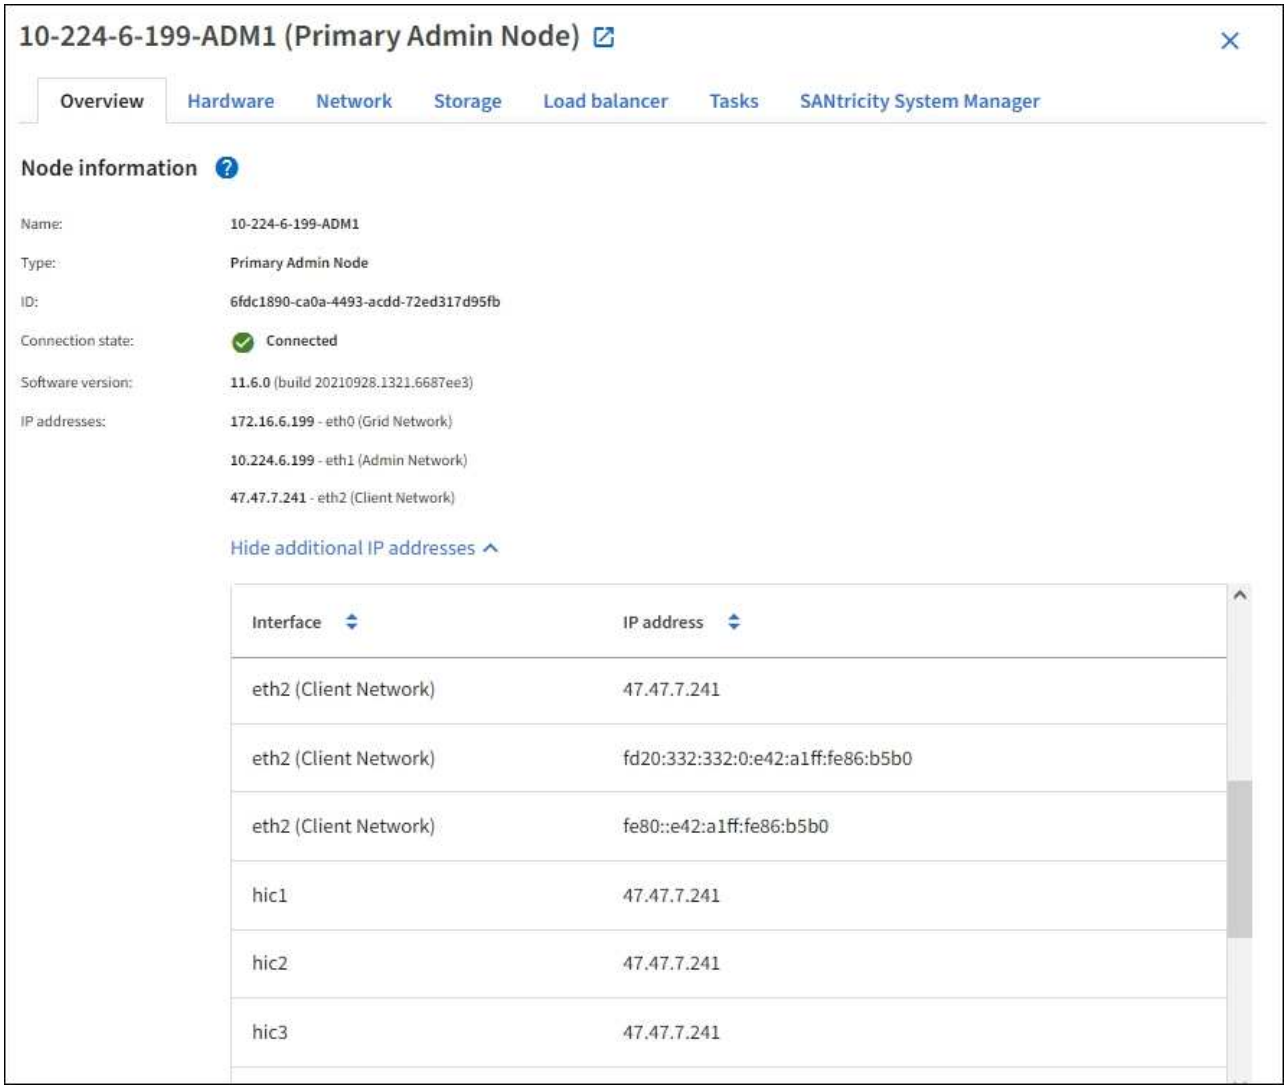


   The Alerts section of the Overview tab displays any active alerts for the node.

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.

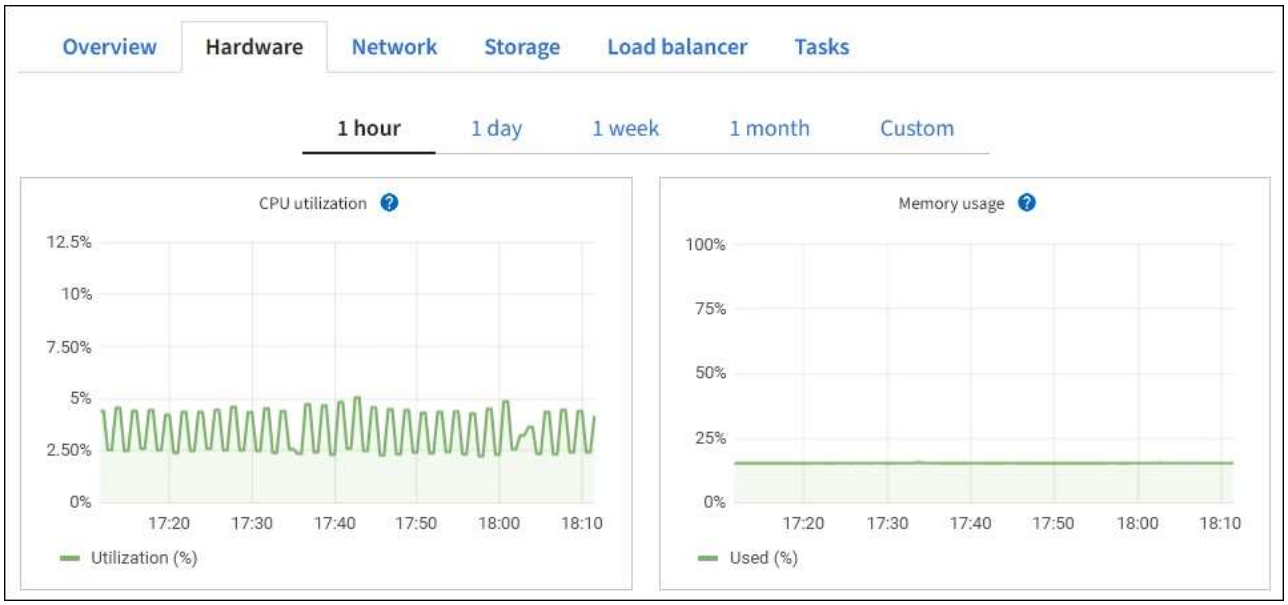


b. Scroll down to view the table of components for the appliance. This table contains information such as the model name, serial number, controller firmware version, and the status of each component.

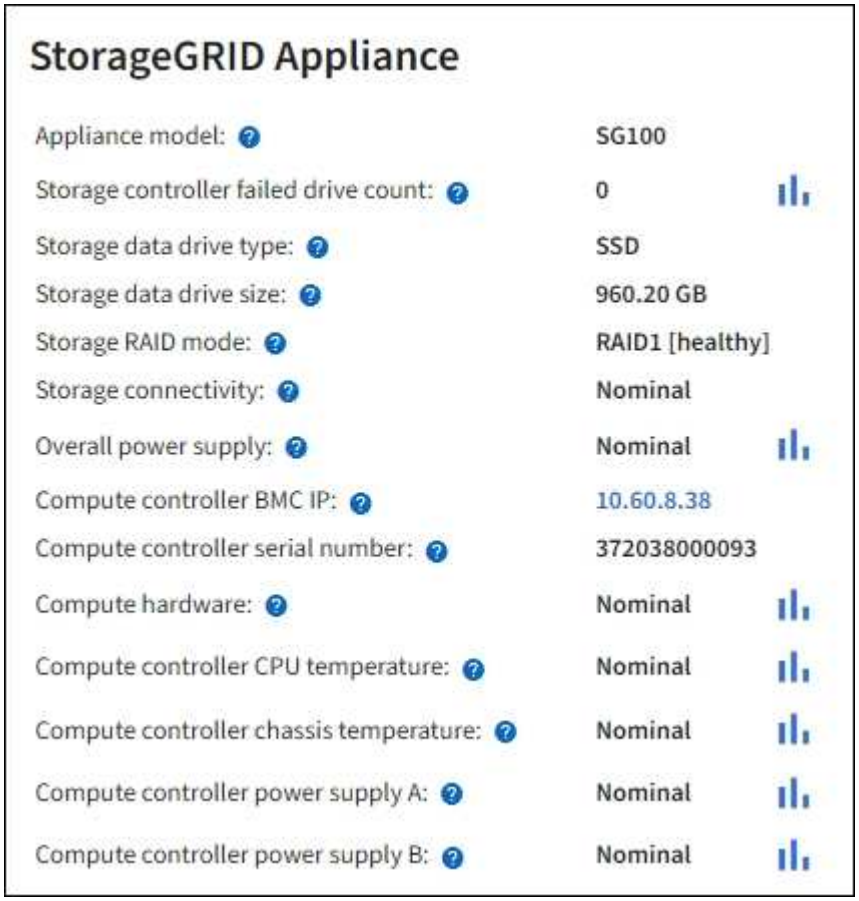


| Field in the Appliance table | Description |
|---|---|
| Appliance model | The model number for this StorageGRID appliance. |

| 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.<br><br>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.

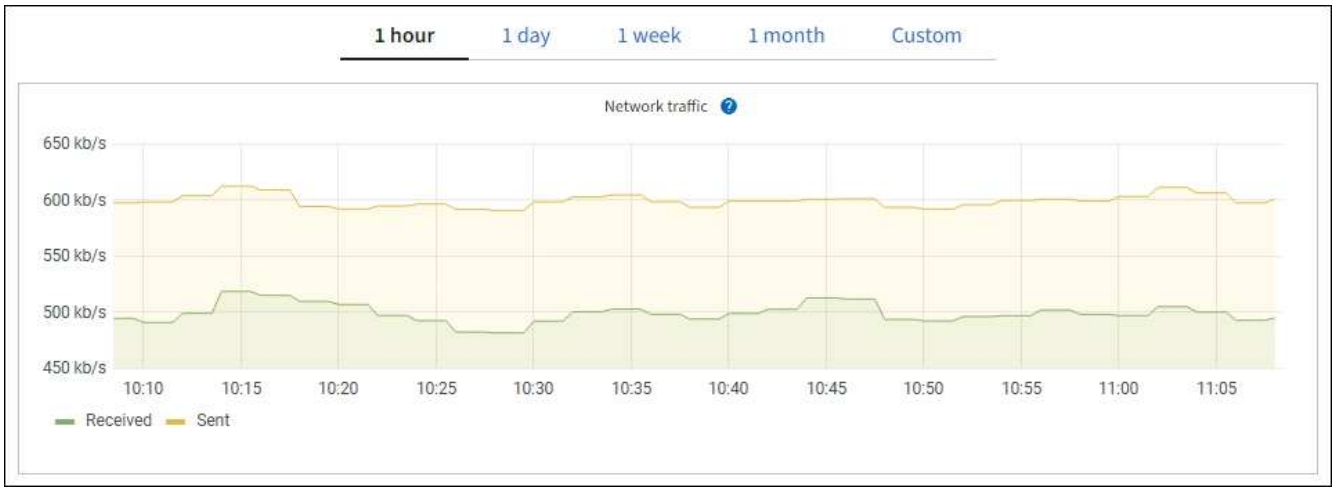


a. Review the Network Interfaces section.

**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.

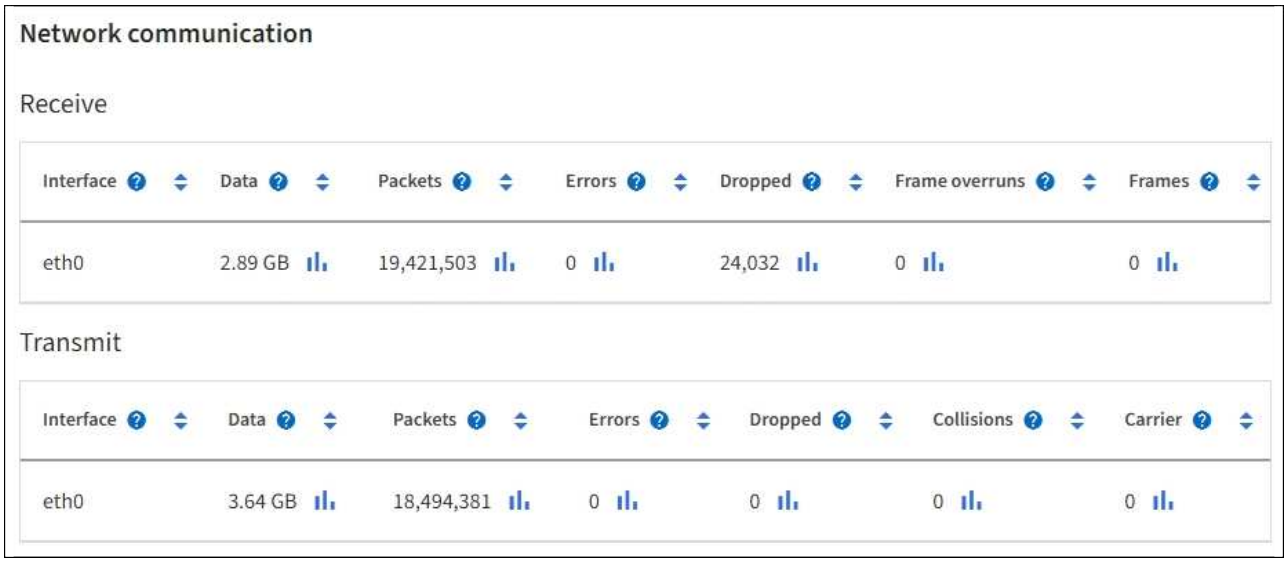

**Network communication**

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 |

**Related information**

SG100 and SG1000 services appliances

## 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.

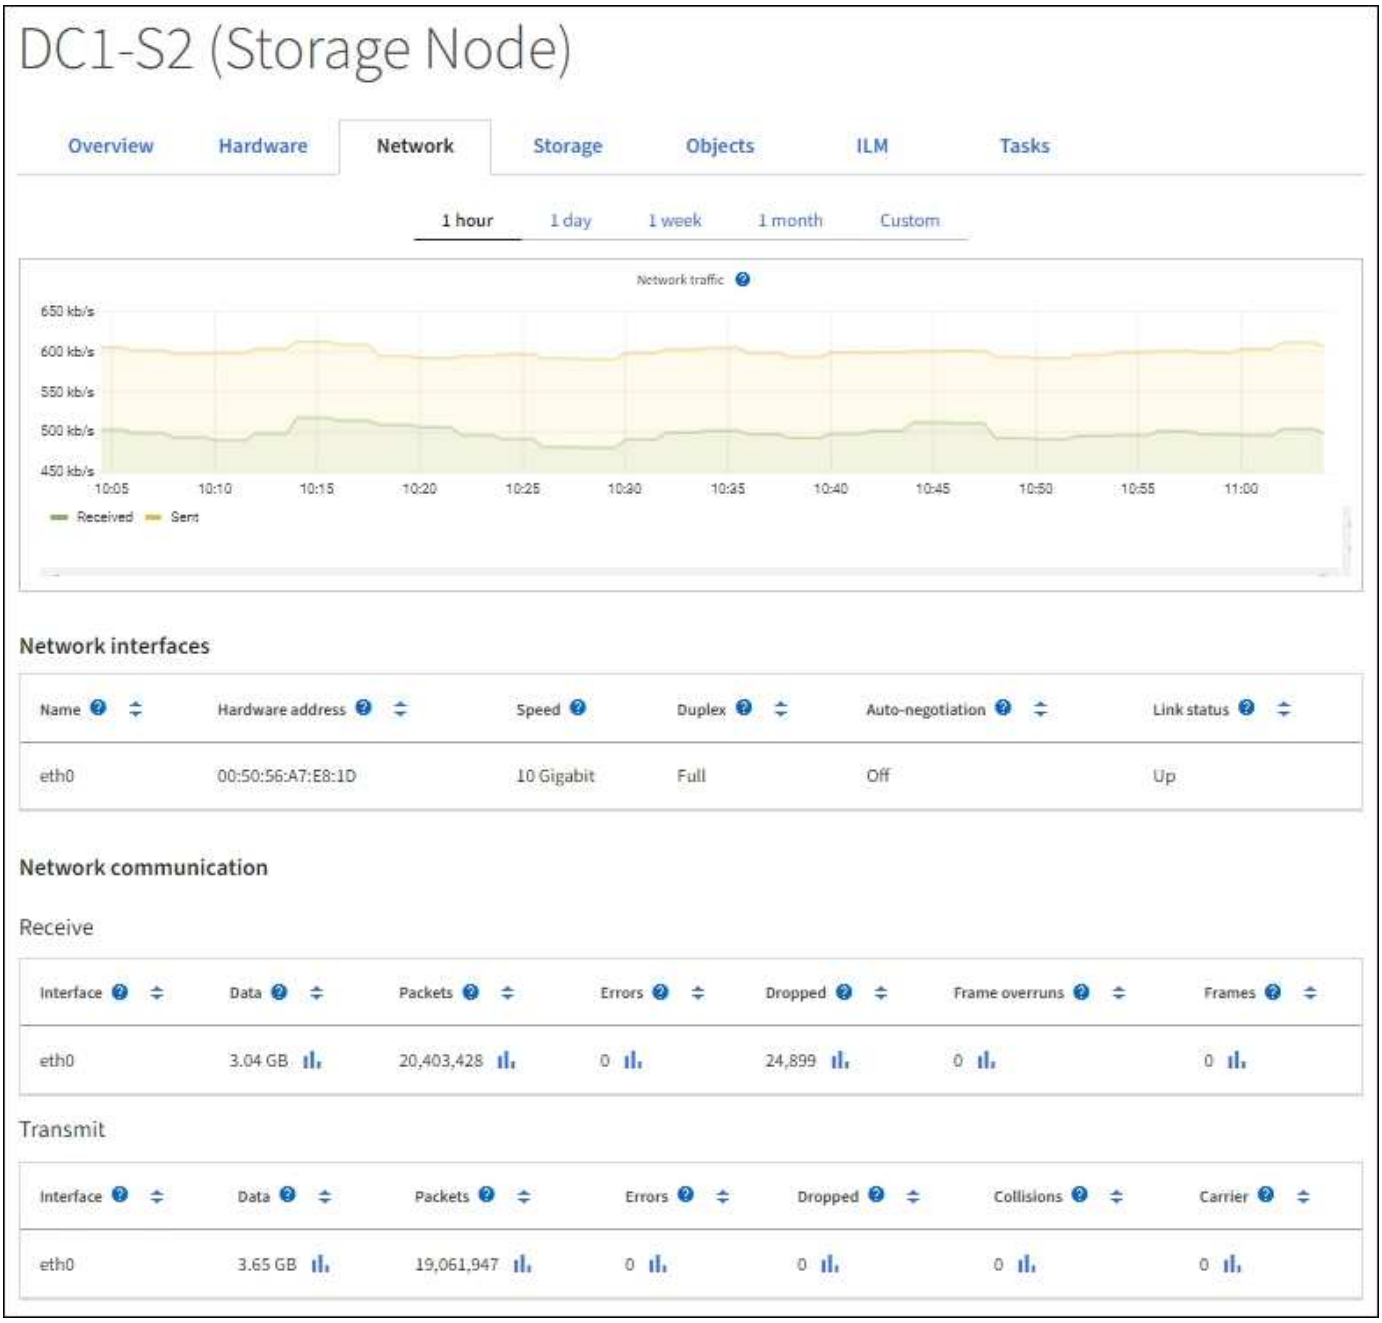


**Related information**

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.

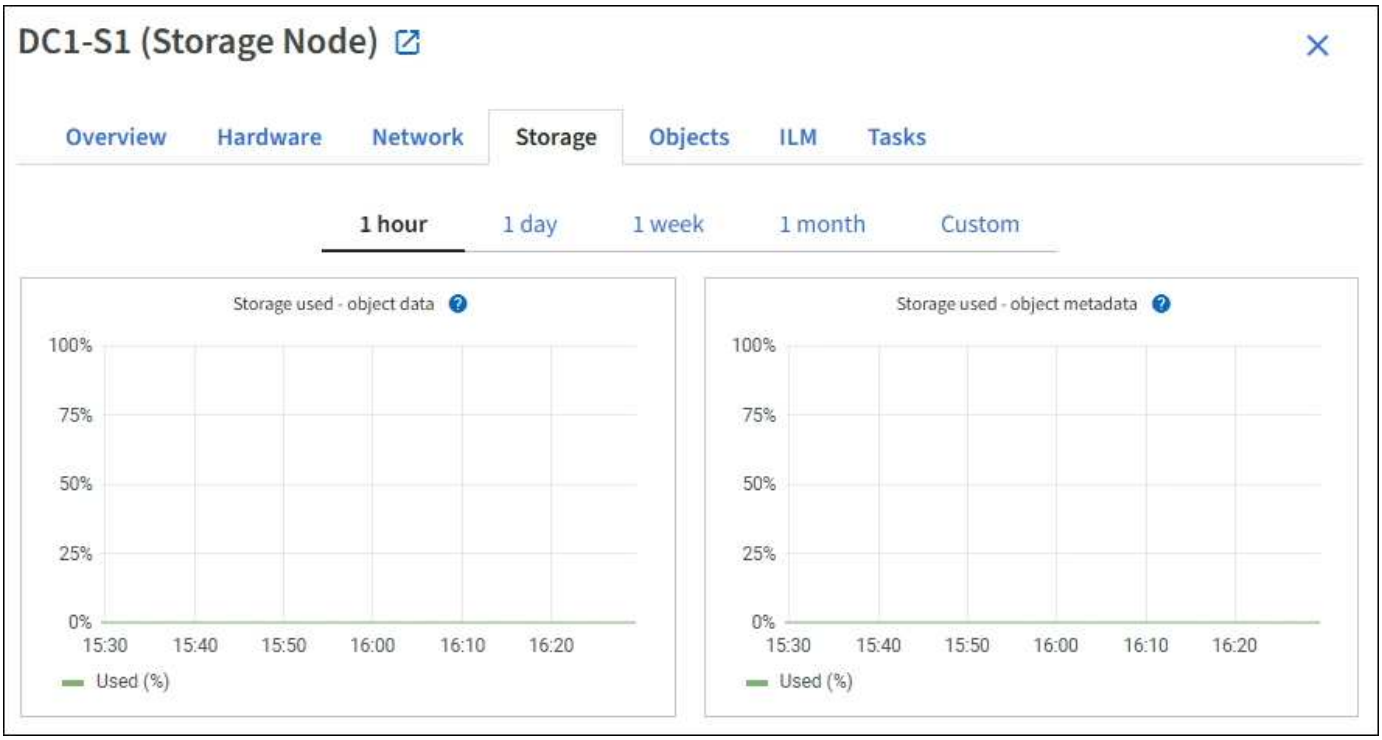


**Disk devices, Volumes, and Object stores tables**

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

## Use the Task tab to reboot a grid node

The Task tab allows you to reboot the selected node. The Task tab is shown for all nodes.

**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 the provisioning passphrase.

**About this task**

You can use the Task tab to reboot a node. For appliance nodes, you can also use the Task tab to place the appliance into maintenance mode.

- Rebooting a grid node from the Task tab issues the reboot command on the target node. When you reboot a node, the node shuts down and restarts. All services are restarted automatically.

  If you plan to reboot a Storage Node, note the following:

  - If an ILM rule specifies an ingest behavior of Dual commit or the rule specifies Balanced and it is not possible to immediately create all required copies, StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site and evaluates ILM later. If you want to reboot two or more Storage Nodes on a given site, you might not be able to access these objects for the duration of the reboot.
  - To ensure you can access all objects while a Storage Node is rebooting, stop ingesting objects at a site for approximately one hour before rebooting the node.
- You might need to put a StorageGRID appliance into maintenance mode to perform certain procedures, such as changing the link configuration or replacing a storage controller. For instructions, see Place appliance into maintenance mode.

  > In rare instances, placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

**Steps**

1. Select **NODES**.
2. Select the grid node you want to reboot.
3. Select the **Tasks** tab.

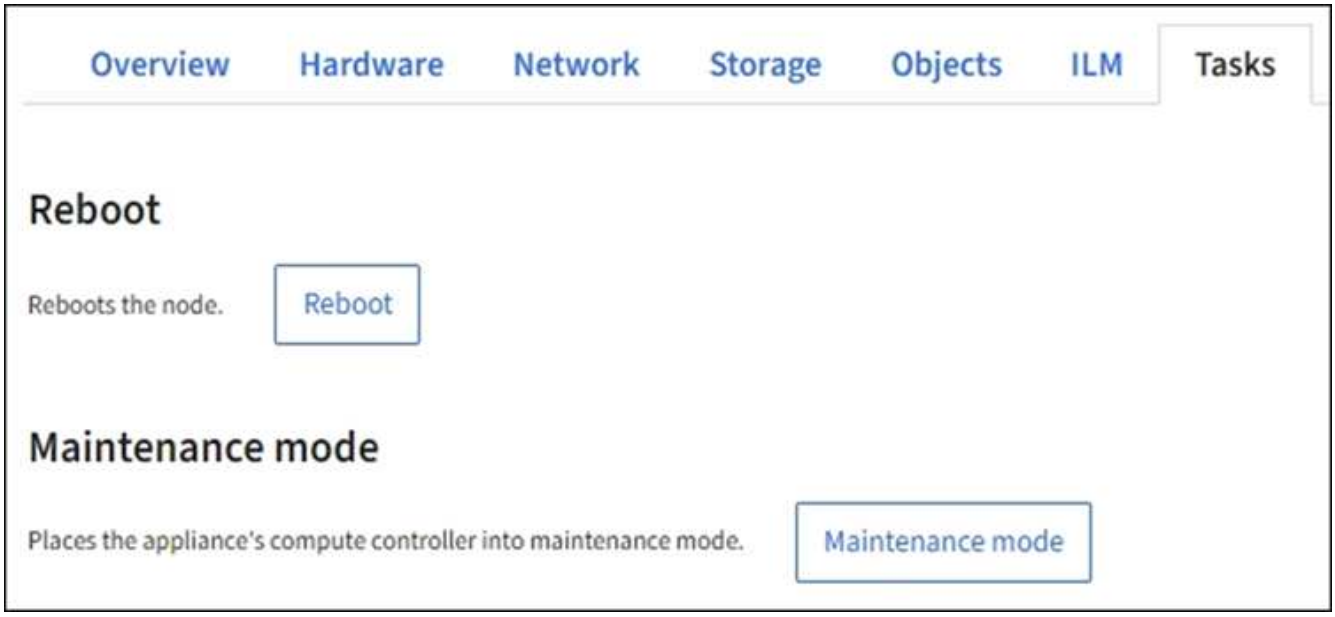


4. Select **Reboot**.

   A confirmation dialog box appears.

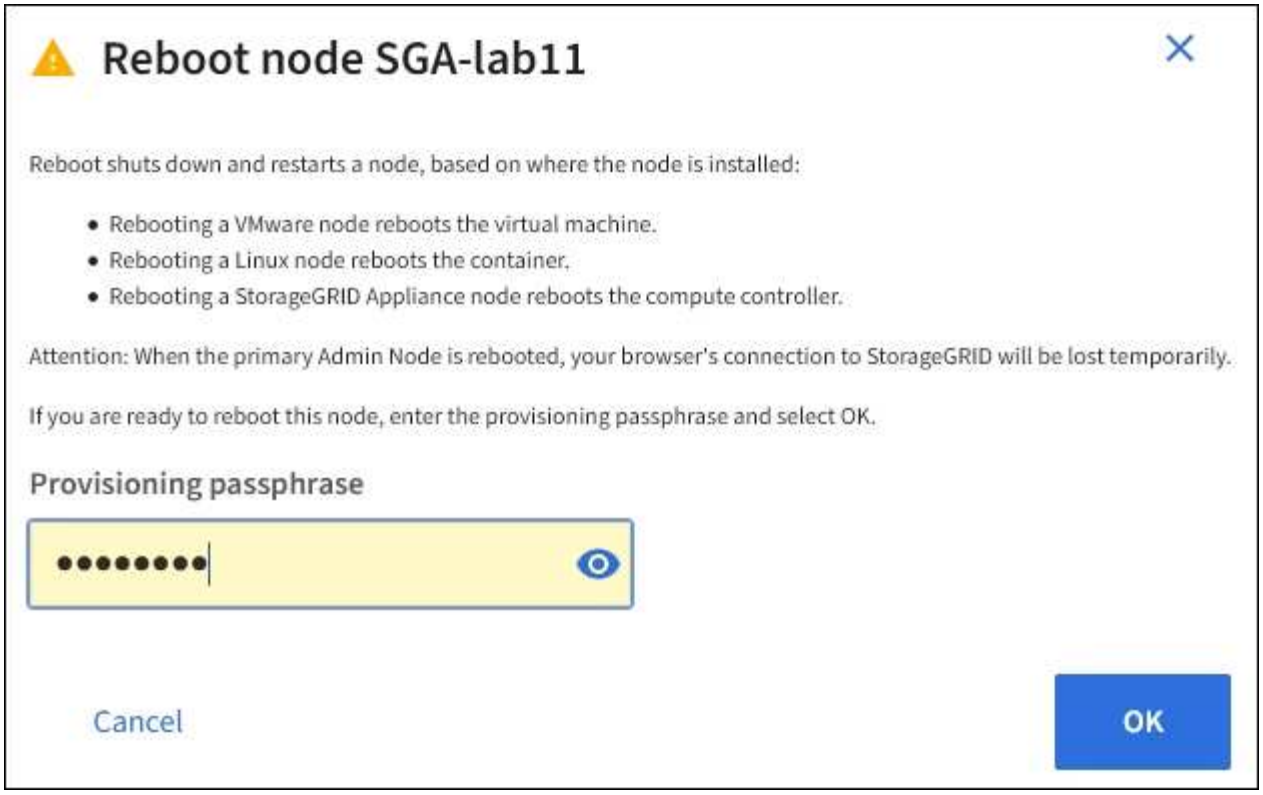


> If you are rebooting the primary Admin Node, the confirmation dialog box reminds you that your browser's connection to the Grid Manager will be lost temporarily when services are stopped.

5. Enter the provisioning passphrase, and select **OK**.
6. Wait for the node to reboot.

   It might take some time for services to shut down.

   When the node is rebooting, the gray icon (Administratively Down) appears on the left side of the **Nodes** page. When all services have started again and the node is successfully connected to the grid, the **Nodes** page should display a normal status (no icons to the left of the node name), indicating that no alerts are active and the node is connected to the grid.

## 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.

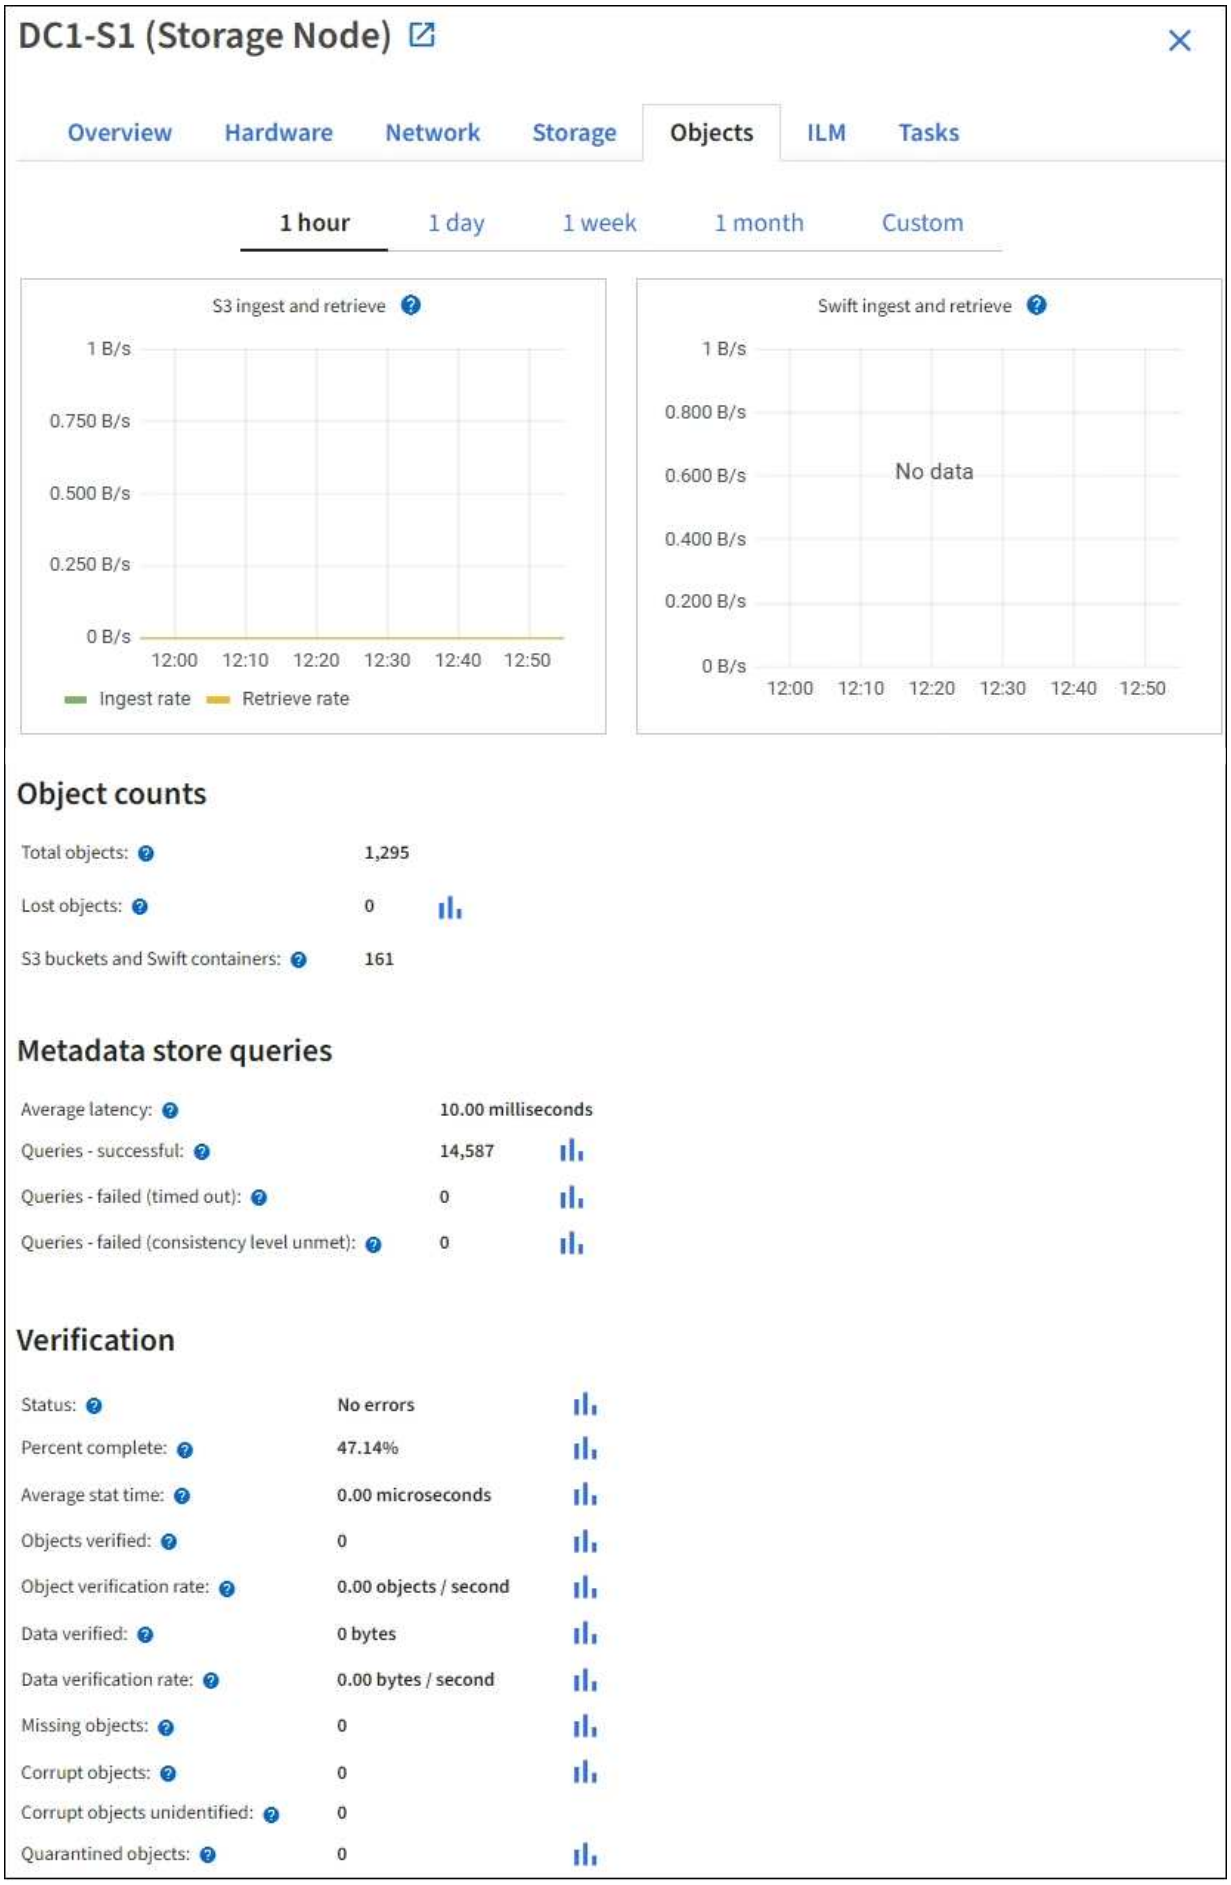
DC1-S1 (Storage Node)
Overview
Hardware
Network
Storage
Objects
ILM
Tasks
1 hour
1 day
1 week
1 month
Custom
S3 ingest and retrieve
1 B/s
0.750 B/s
0.500 B/s
0.250 B/s
0 B/s
12:00
12:10
12:20
12:30
12:40
12:50
Ingest rate
Retrieve rate
Swift ingest and retrieve
1 B/s
0.800 B/s
0.600 B/s
0.400 B/s
0.200 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.

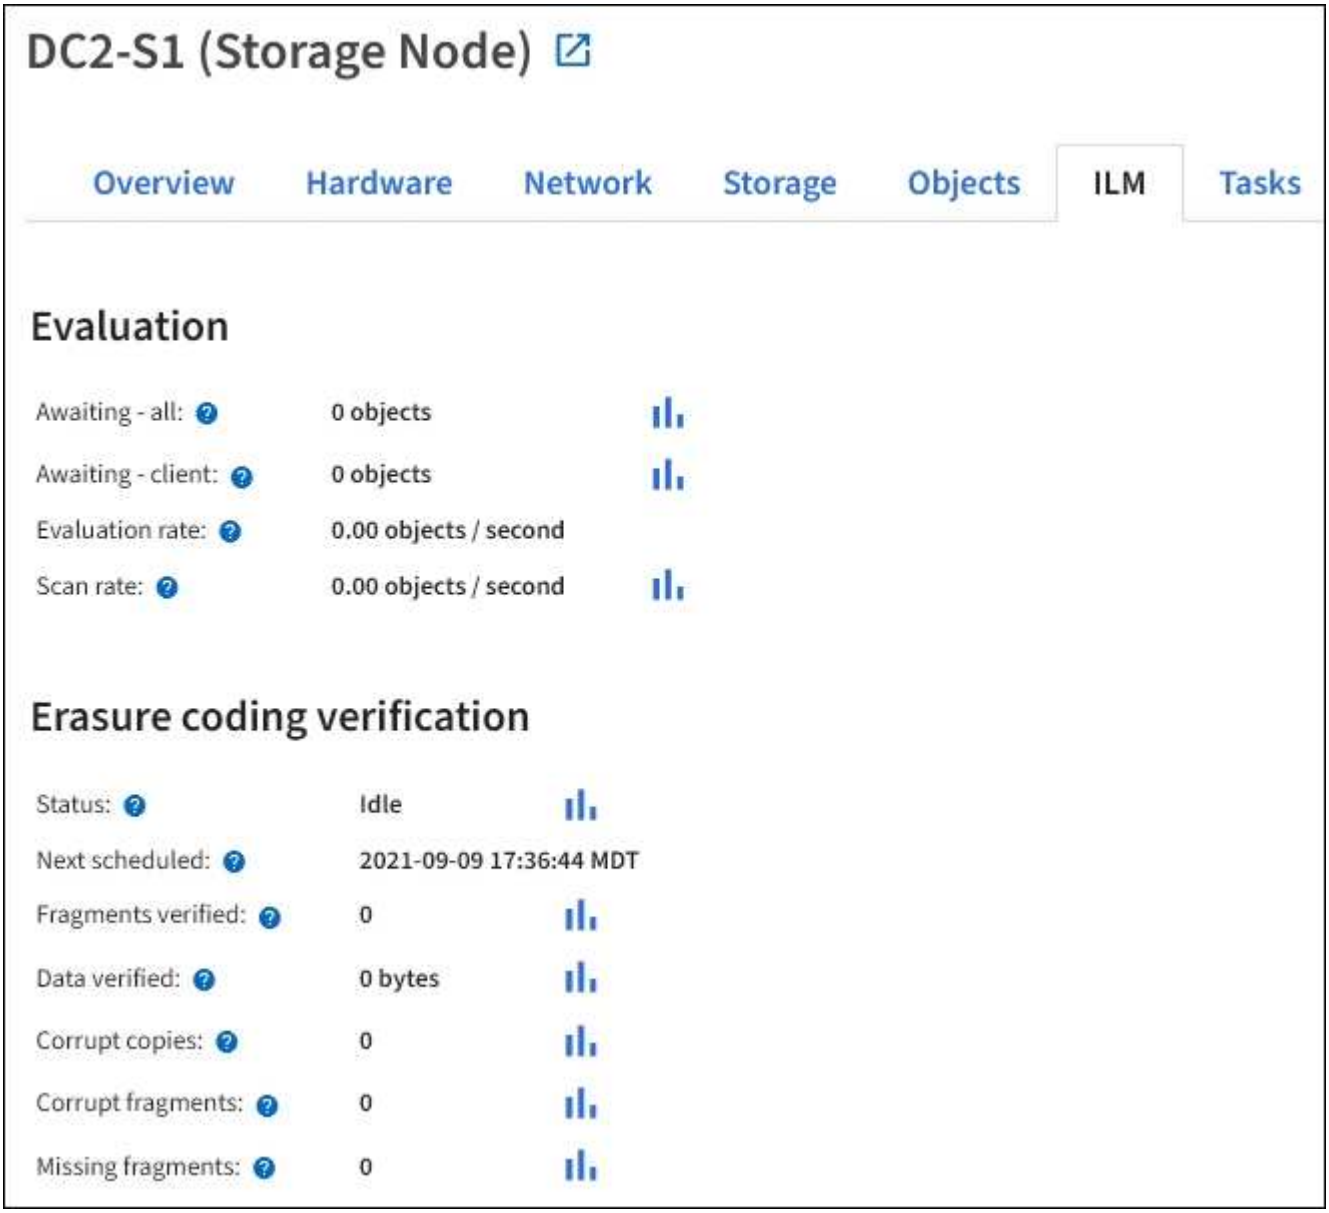


**Related information**

Monitor information lifecycle management

Administer StorageGRID

## 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.

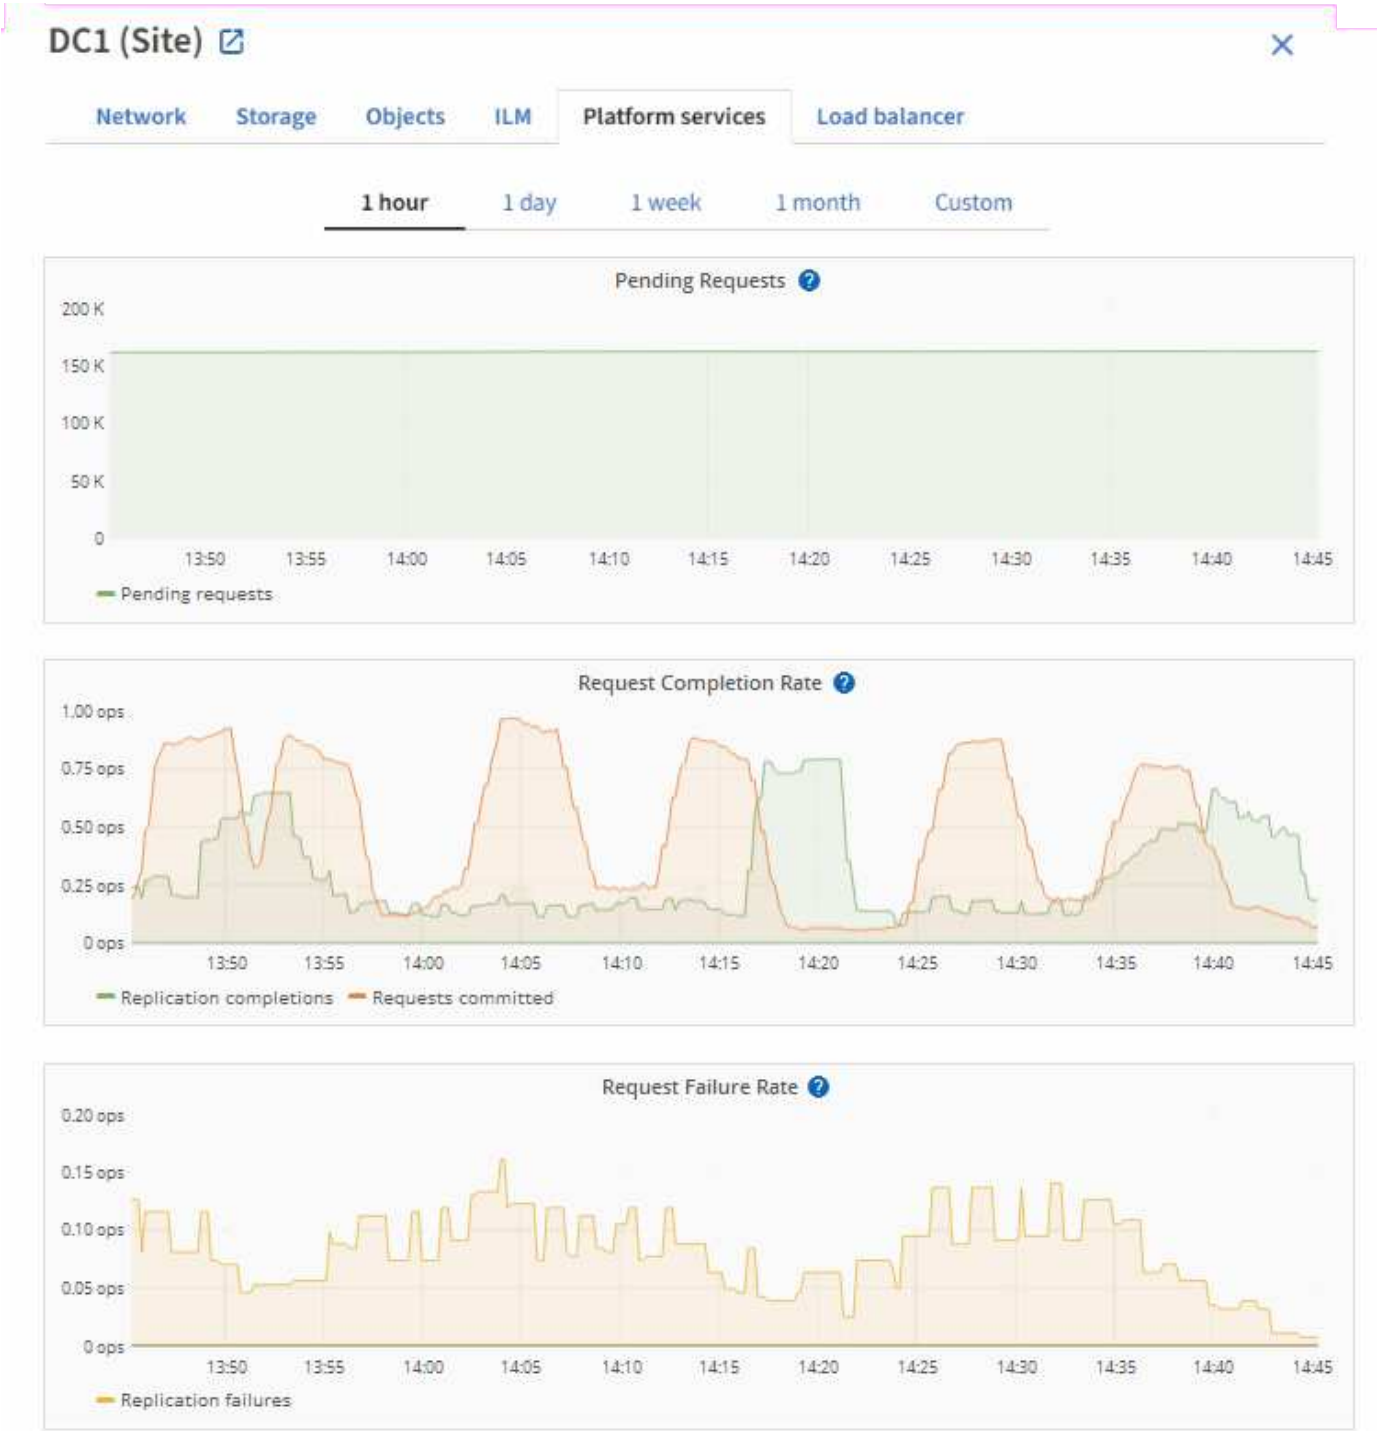


For more information about S3 platform services, including troubleshooting details, see the instructions for administering StorageGRID.

### View the SANtricity System Manager tab

The SANtricity System Manager tab on the Nodes page in Grid Manager 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 messages through StorageGRID.

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

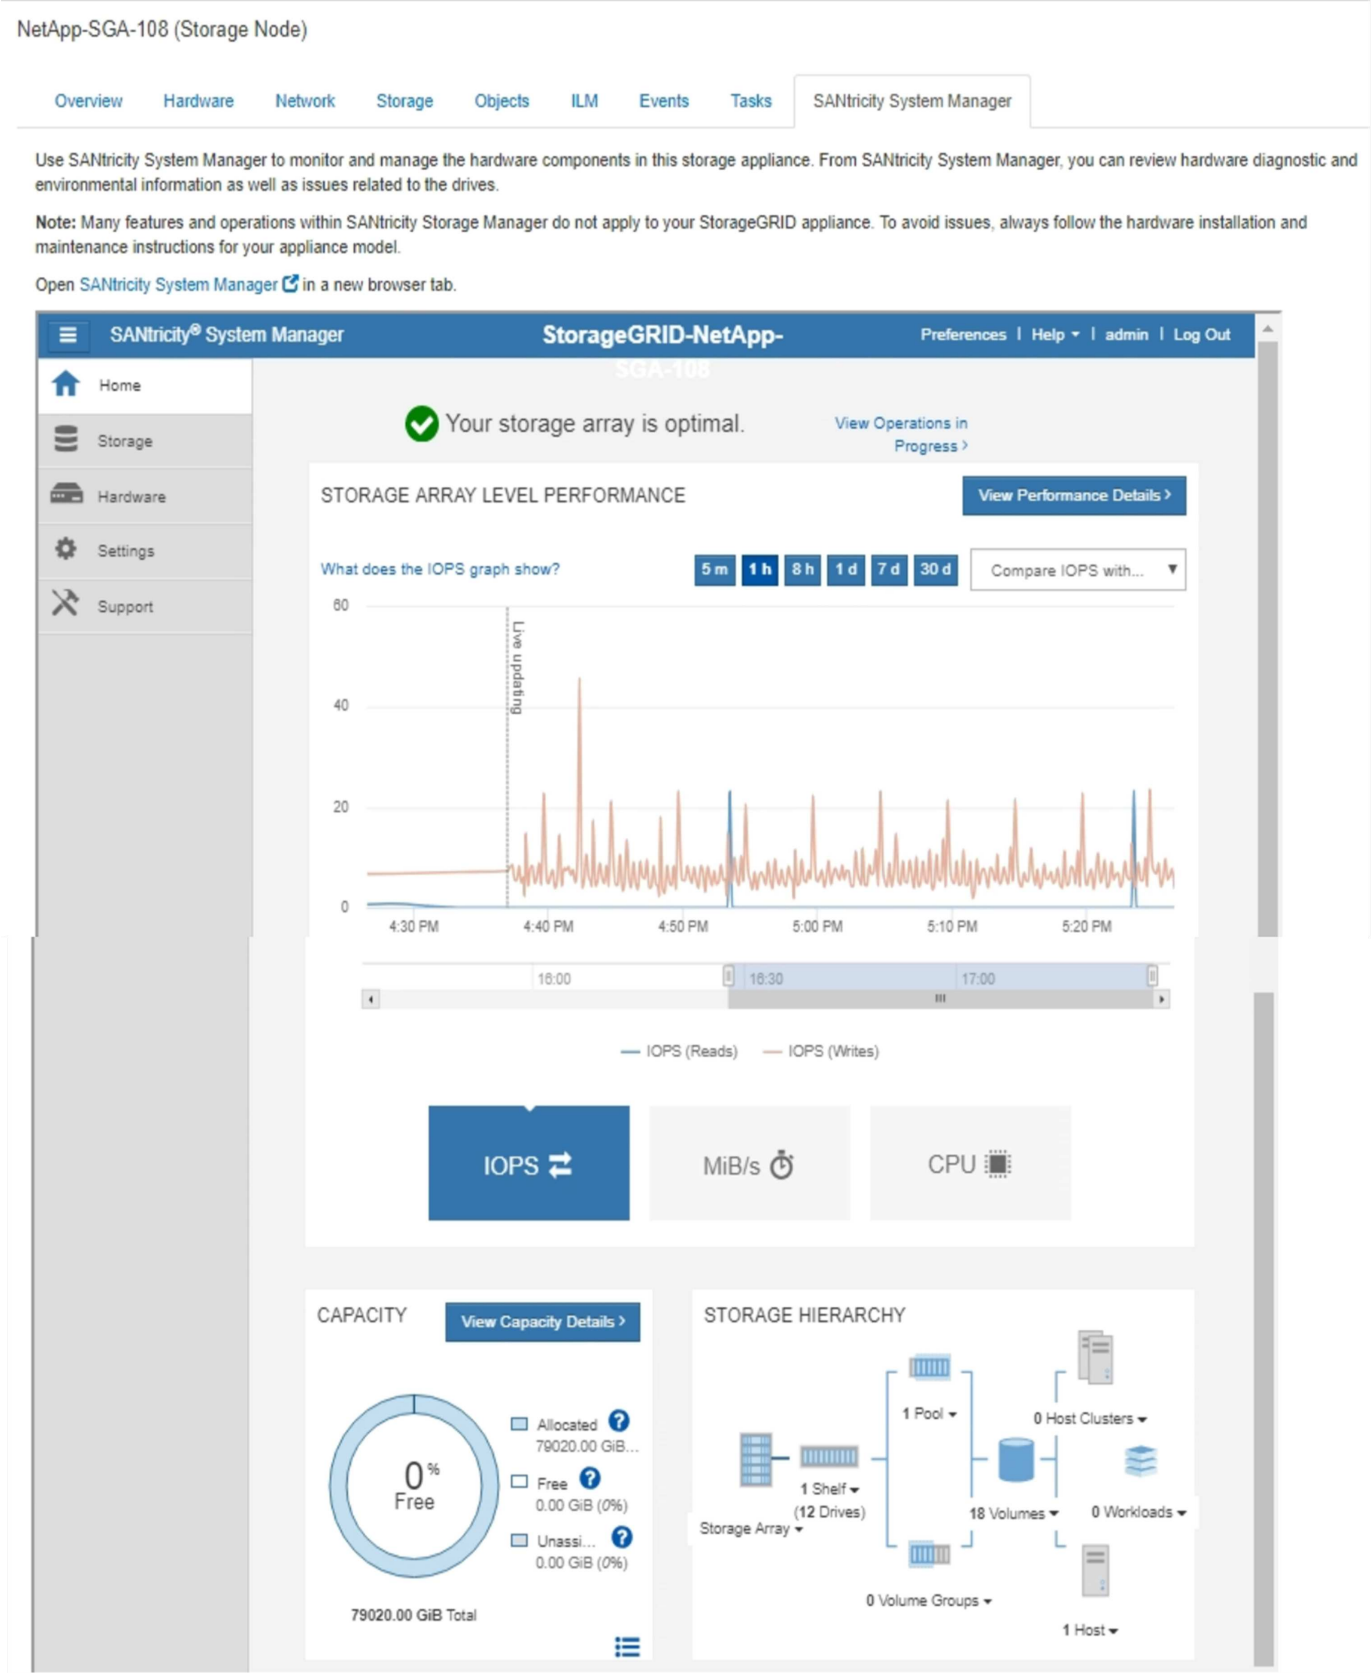


You can use the SANtricity System Manager link to open the SANtricity System Manager in a new browser window for easier viewing.

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.

**Related information**

- Maintain SG6000 appliance
- Maintain SG5700 appliance

# Information you should 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 |
| Load balancing operations | After the initial configuration and after any configuration changes |
| Grid federation connections | Weekly |
| Availability of software hotfixes and software upgrades | Monthly |
| 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.

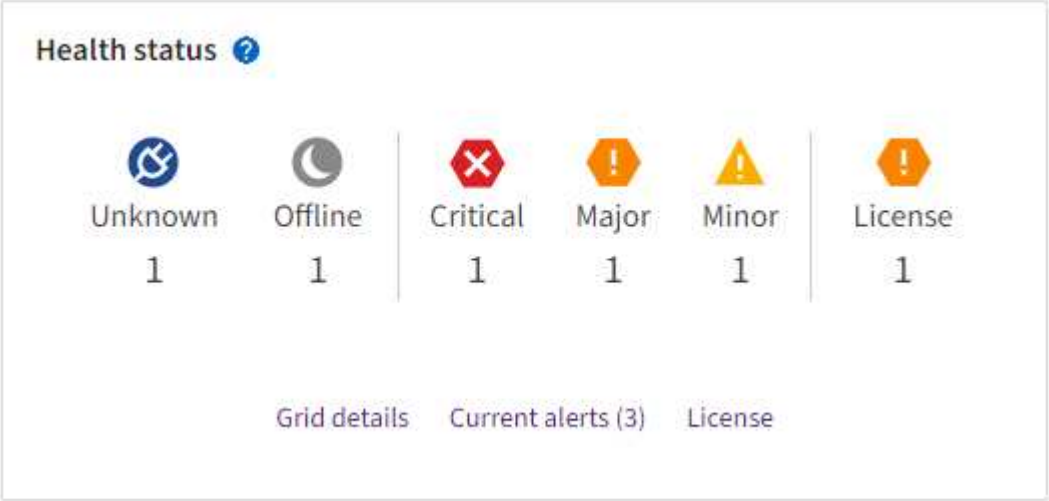


When issues exist, links appear that allow you to view additional details:

| 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**<br><br>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.<br><br>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.<br><br>For example, you might need to restart a service that has stopped or restart the host for the node.<br><br>**Note**: A node might appear as Unknown during managed shutdown operations. You can ignore the Unknown state in these cases. |
| | **Not connected - Administratively down**<br><br>For an expected reason, node is not connected to grid.<br><br>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.<br><br>Based on the underlying issue, these nodes often go back online with no intervention. | Determine if any alerts are affecting this node.<br><br>If one or more alerts are active, select each alert and follow the recommended actions. |
| | **Connected**<br><br>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 11.7

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.<br><br>**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.<br><br>**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.<br><br>**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*).<br><br>**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. |
| 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.<br><br>**Note:** If multiple current alerts are grouped, current values aren't shown in the title row. |

| Column header | Description |
| --- | --- |
| 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.

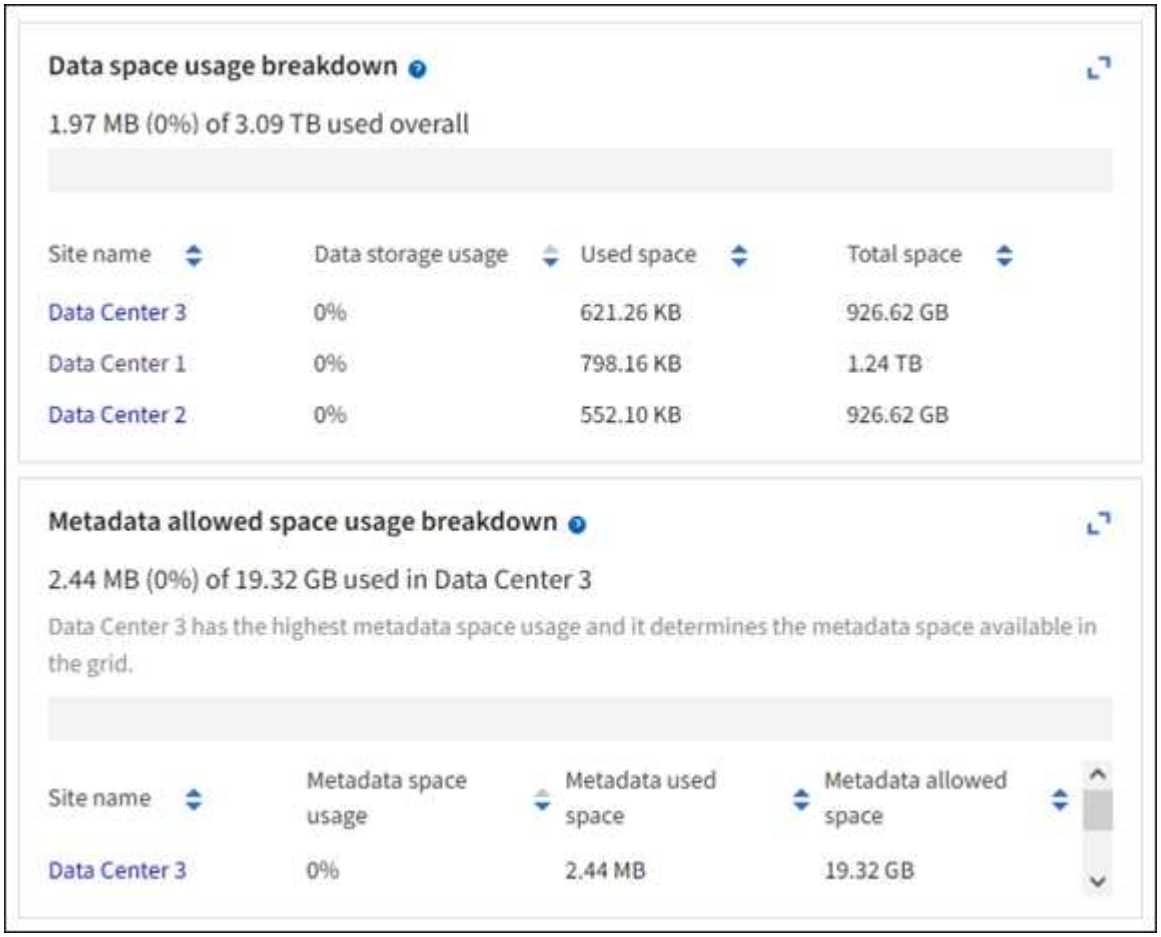


c. Note the chart on the Storage over time card. Use the time period drop-down to help you determine how quickly storage is consumed.

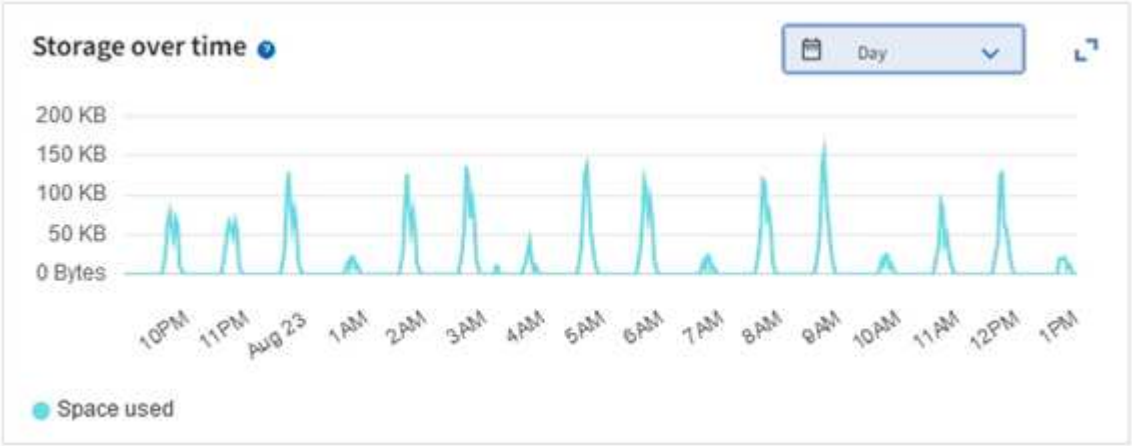


2. Use the Nodes page for additional details on how much storage has been used and how much storage remains available on the grid for object data and object metadata.
   a. Select **NODES**.
   b. Select ***grid*** > **Storage**.

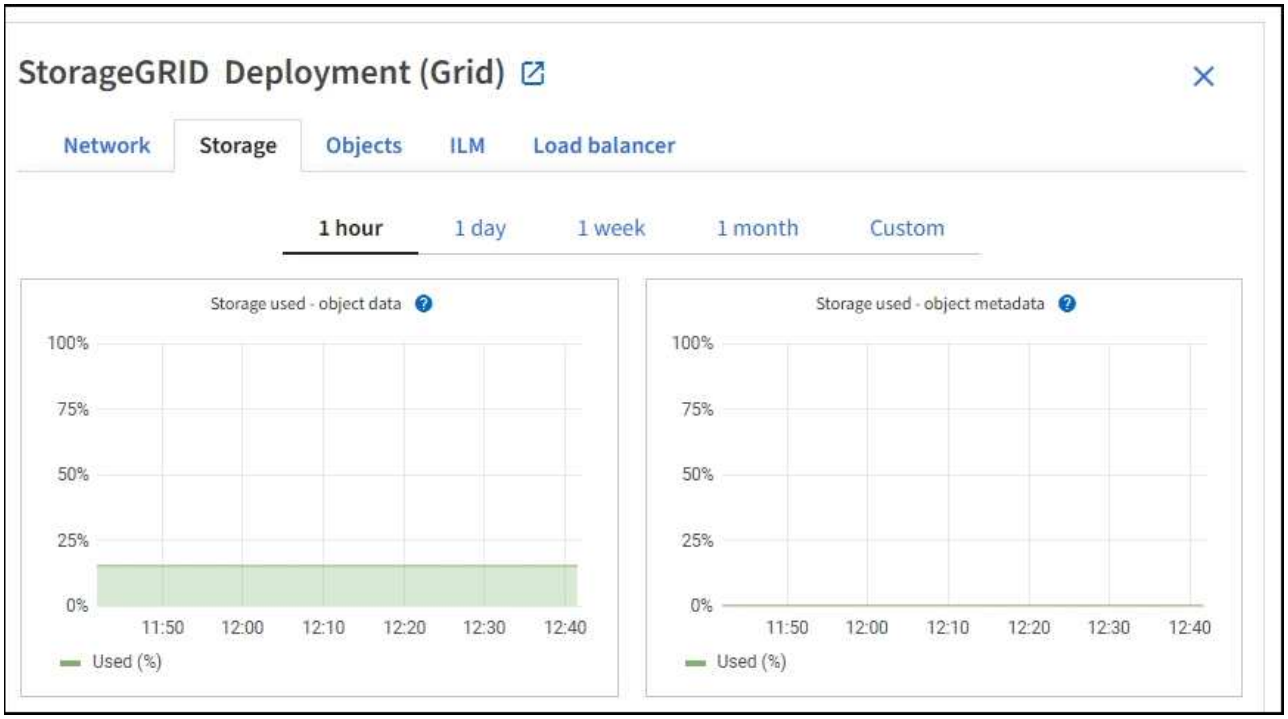


c. Position your cursor over the **Storage used - object data** and the **Storage used - object metadata** charts to see how much object storage and object metadata storage is available for the entire grid, and how much has been used over time.

> 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.

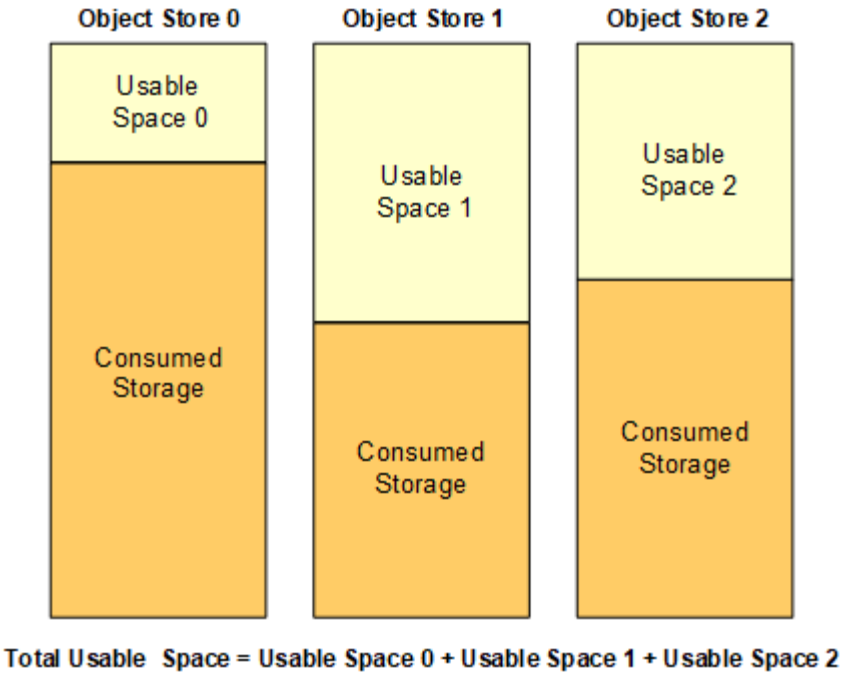


**Steps**

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.

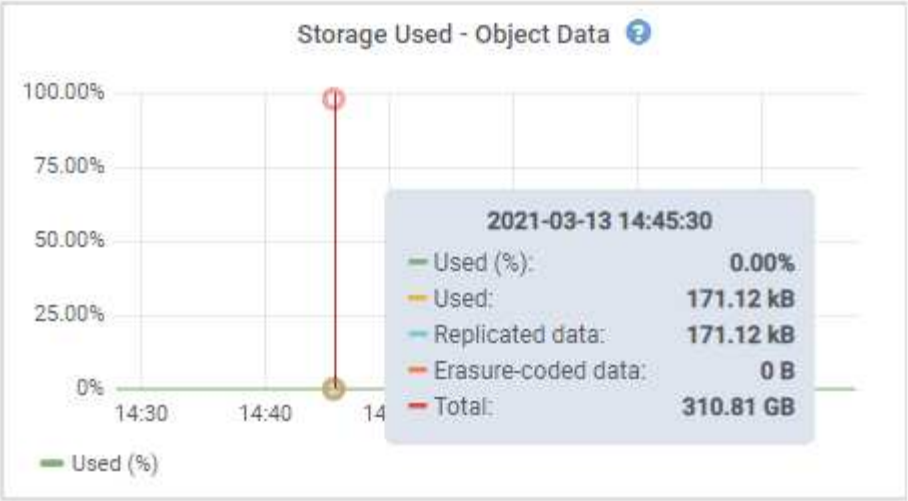


3. Review the Available values in the Volumes and Object stores tables, below the graphs.

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.

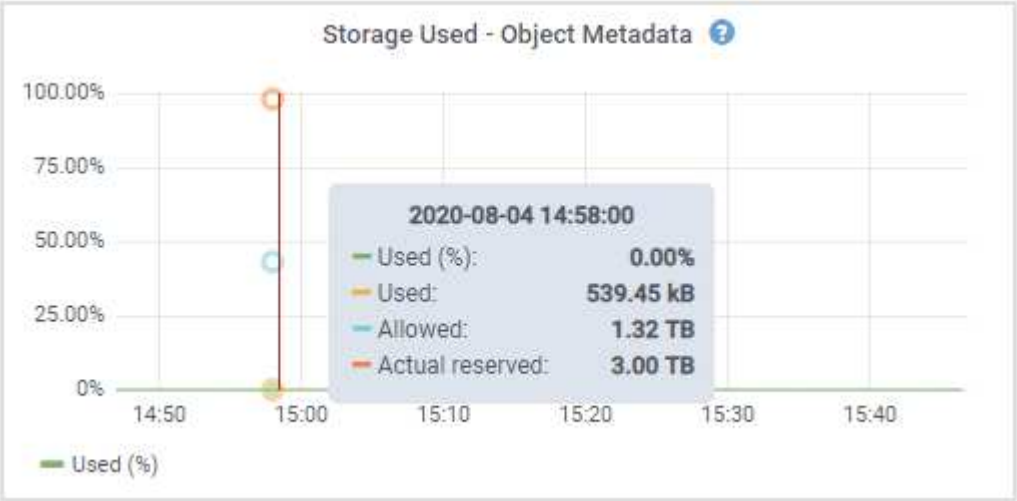


**Used (%)**

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 your 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.

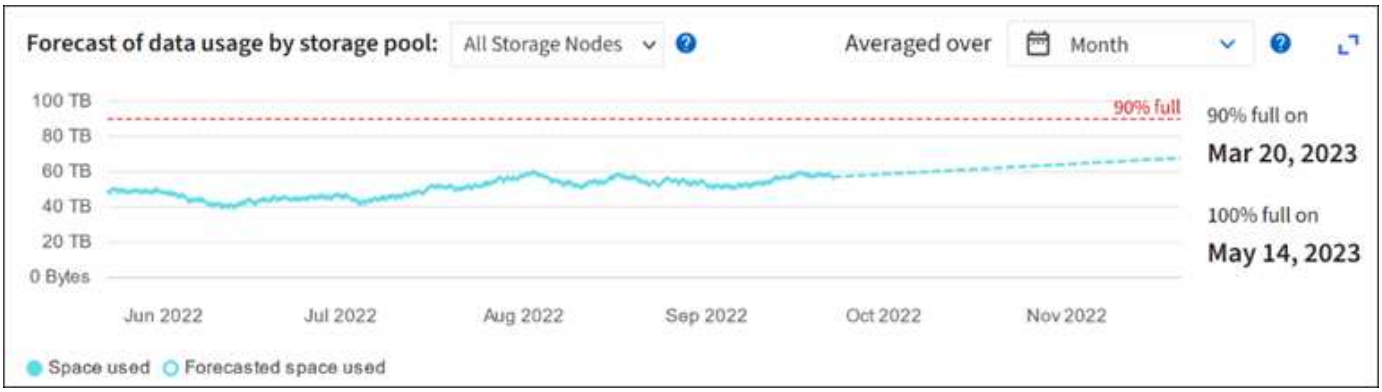


## Monitor information lifecycle management

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 required.

### About this task

The StorageGRID system manages objects by applying the active ILM policy. The ILM policy 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 can monitor whether StorageGRID is keeping up with client actions by charting the Awaiting - Client attribute.

To chart this attribute:

1. Sign in to the Grid Manager.
2. From the dashboard, locate the **Awaiting - Client** entry on the Information Lifecycle Management (ILM) tab.
3. Click the chart icon .

The example chart shows a situation where the number of objects awaiting ILM evaluation temporarily increased in an unsustainable manner, then eventually decreased. Such a trend indicates that ILM was temporarily not fulfilled in near real time.

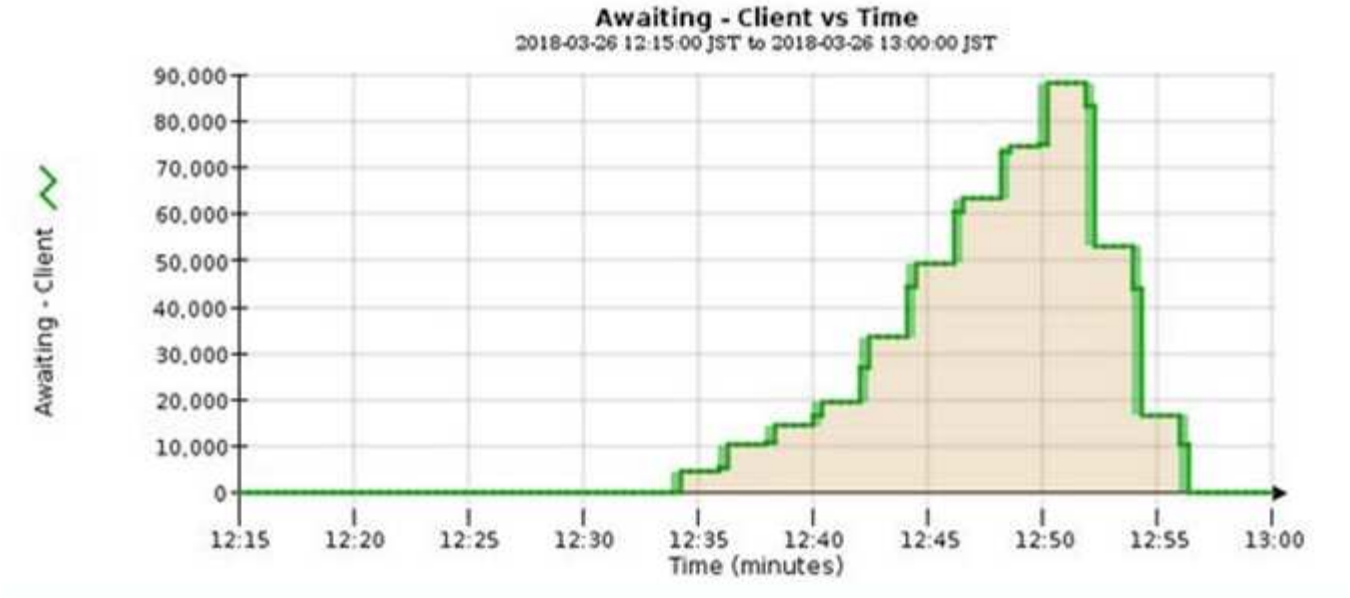


Temporary spikes in the chart of Awaiting - Client are to be expected. But if the value shown on the chart continues to increase and never declines, the grid requires more resources to operate efficiently: either more Storage Nodes, or, if the ILM policy places objects in remote locations, more network bandwidth.

You can further investigate ILM queues using the **NODES** page.

**Steps**

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.

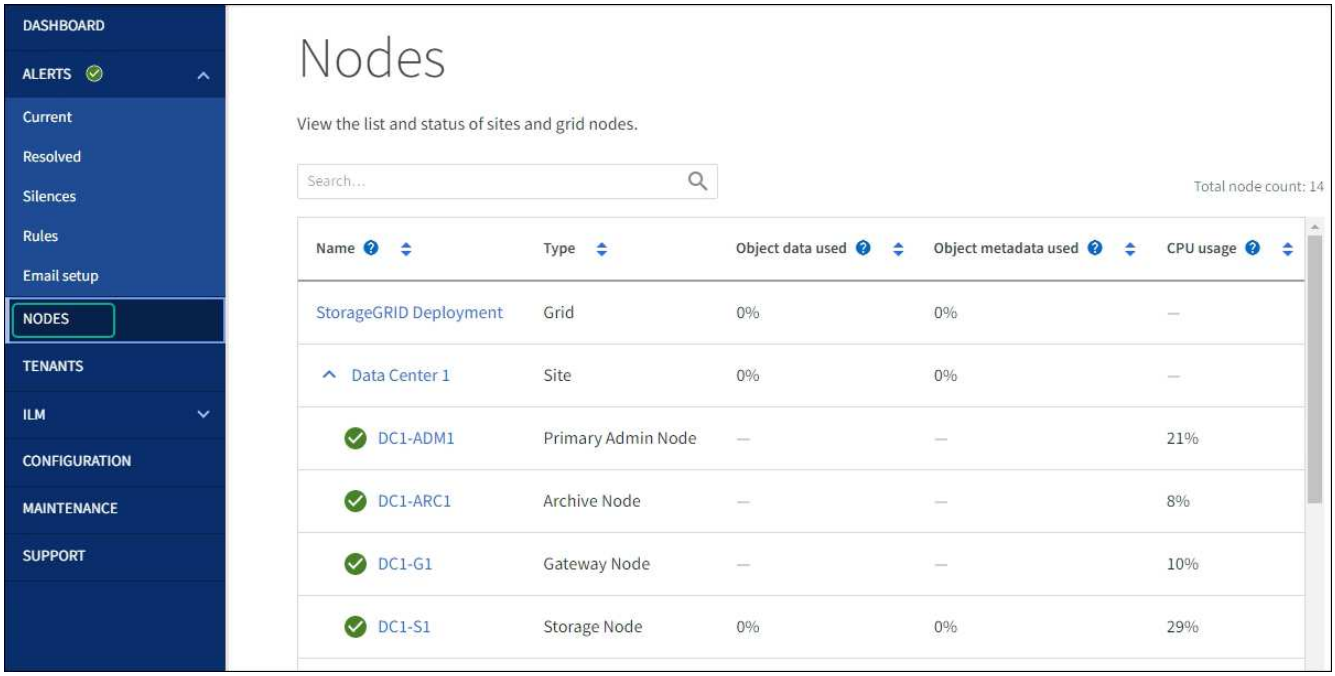


2. Select the grid name, a specific data center site, or a grid node, and then select the **Network** tab.

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.

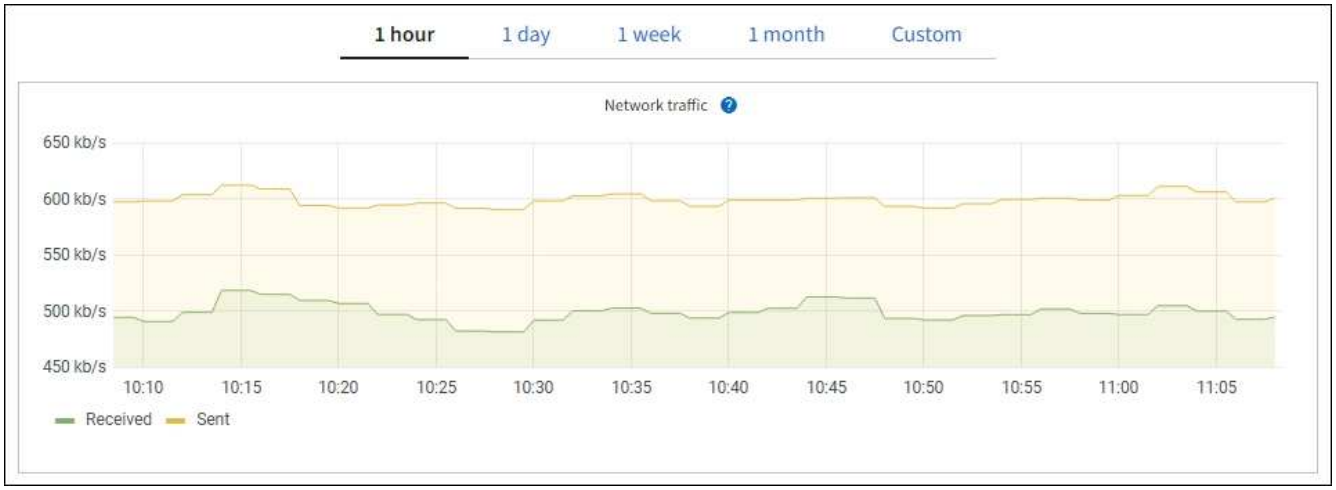


a. If you selected a grid node, scroll down to review the **Network Interfaces** section of the page.

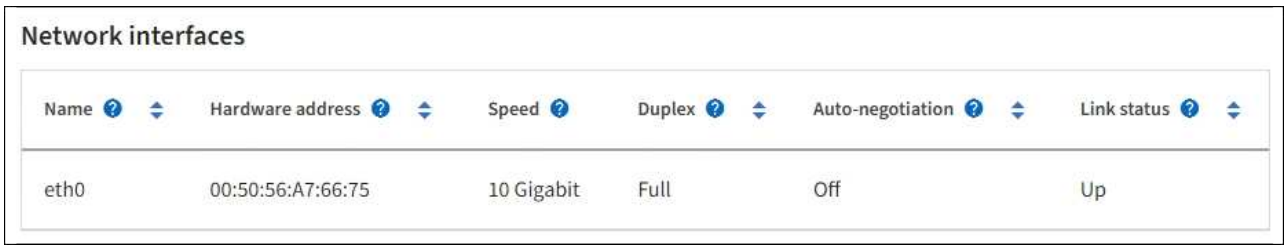

**Network interfaces**

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

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.

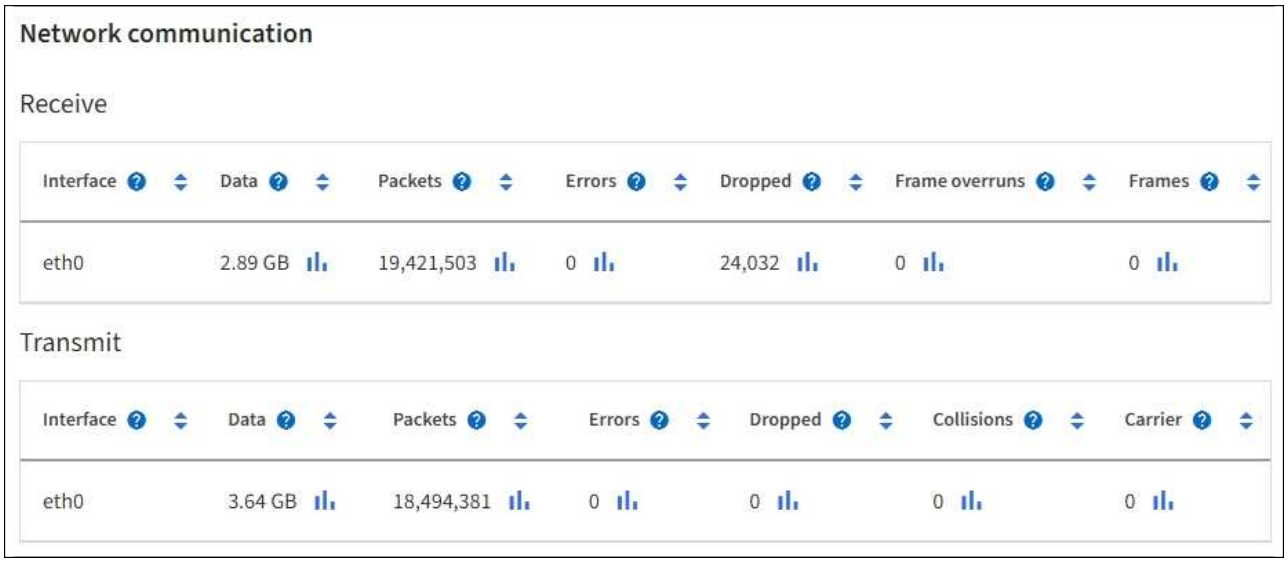

**Network communication**

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.

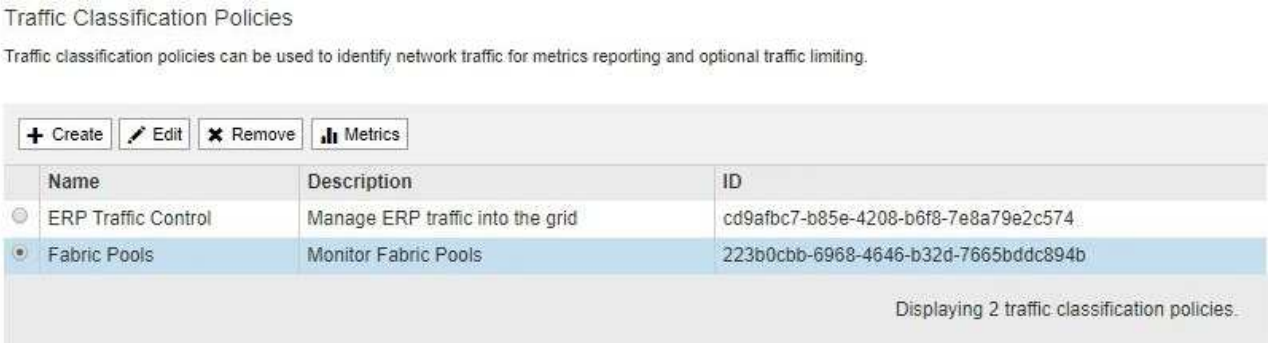


b. To view graphs that show the networking metrics associated with a policy, select the radio button to the left of the policy, and then click **Metrics**.

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.

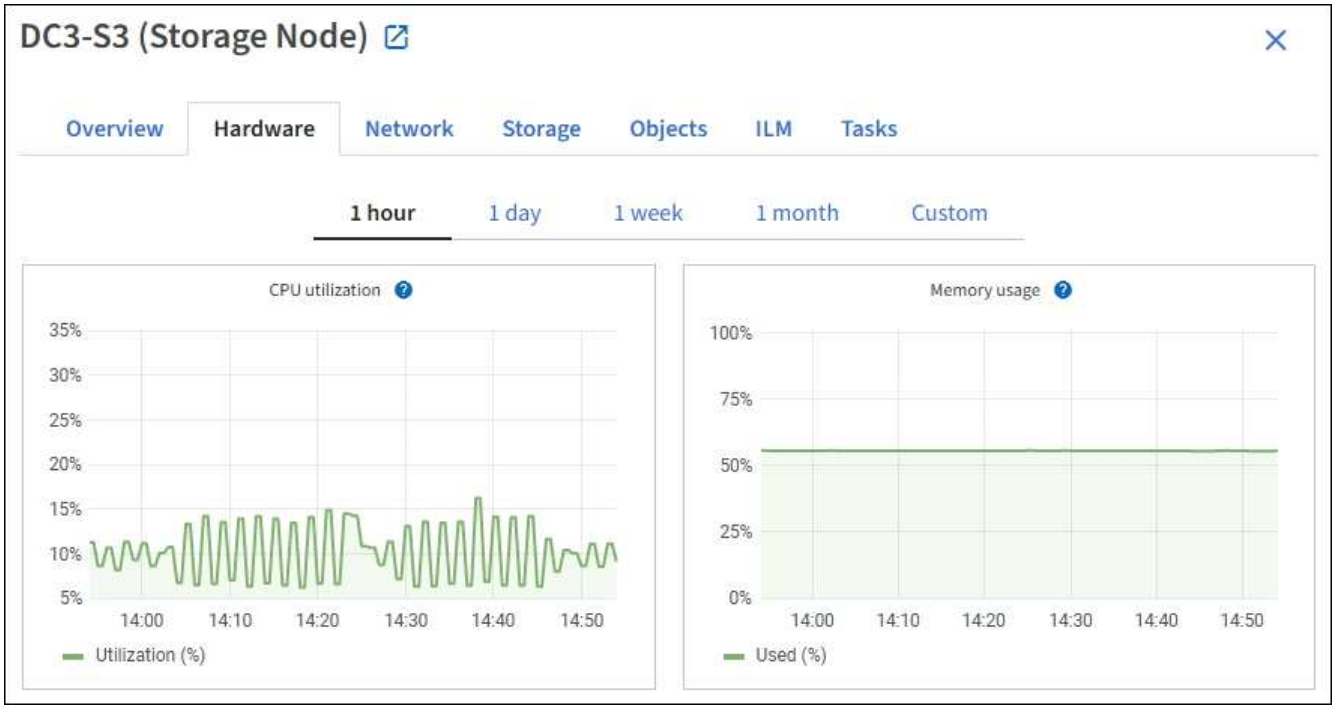

3. To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.
4. If the node is hosted on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

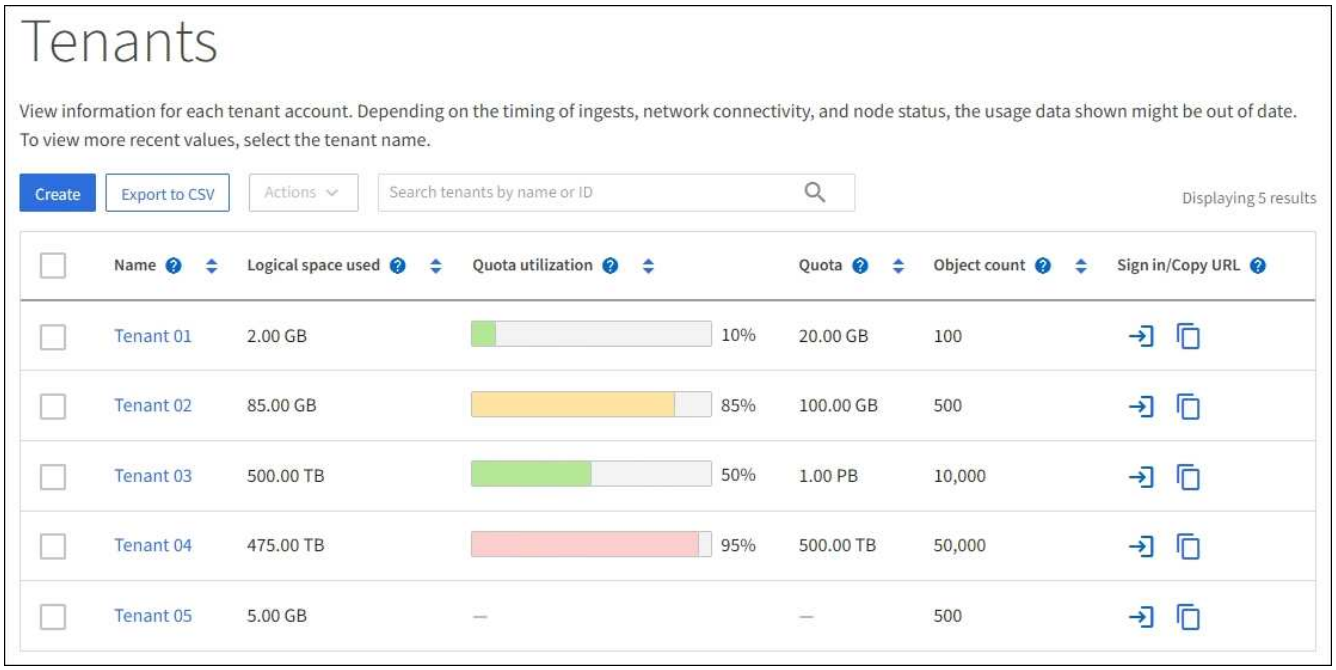


3. Optionally, sign in to a tenant account by selecting the sign-in link in the **Sign in/Copy URL** column.
4. Optionally, copy the URL for a tenant's sign-in page by selecting the copy URL link in the **Sign in/Copy URL** column.
5. Optionally, select **Export to CSV** to view and export a `.csv` file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

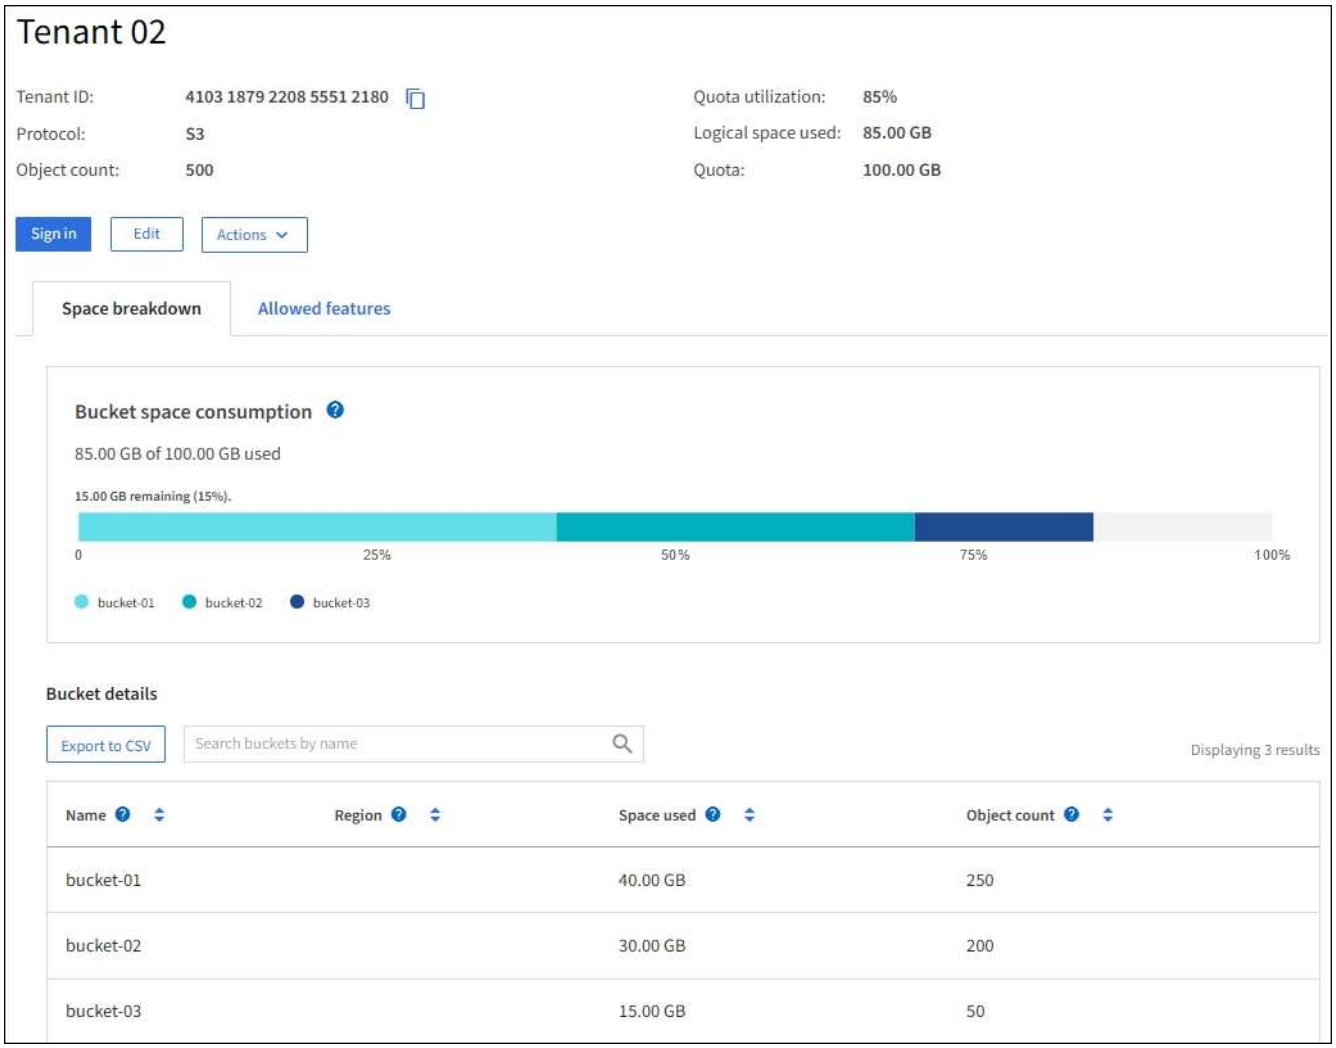


2. Review the tenant overview at the top of the page.

   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.

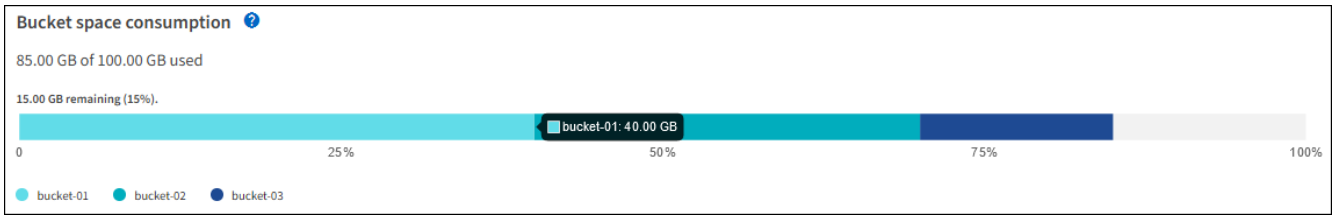

Quota utilization is based on internal estimates and might be exceeded in some cases. For example, StorageGRID checks the quota when a tenant starts uploading objects and rejects new ingests if the tenant has exceeded the quota. However, StorageGRID does not take into account the size of the current 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 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.

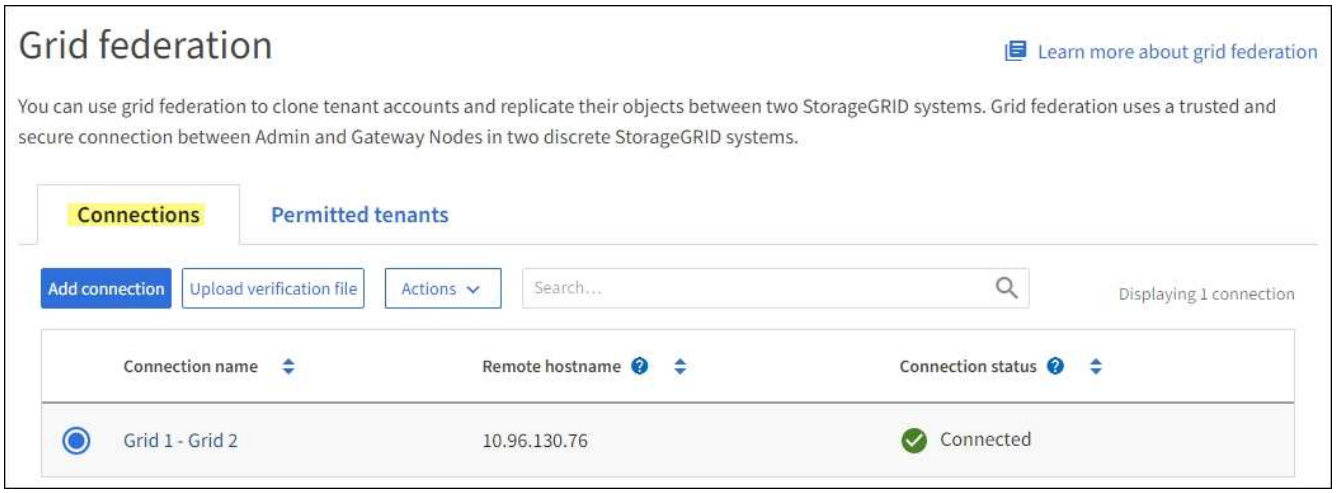


3. To see basic information for all tenant accounts on this grid that have the **Use grid federation connection** permission, select the **Permitted tenants** tab.

   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.

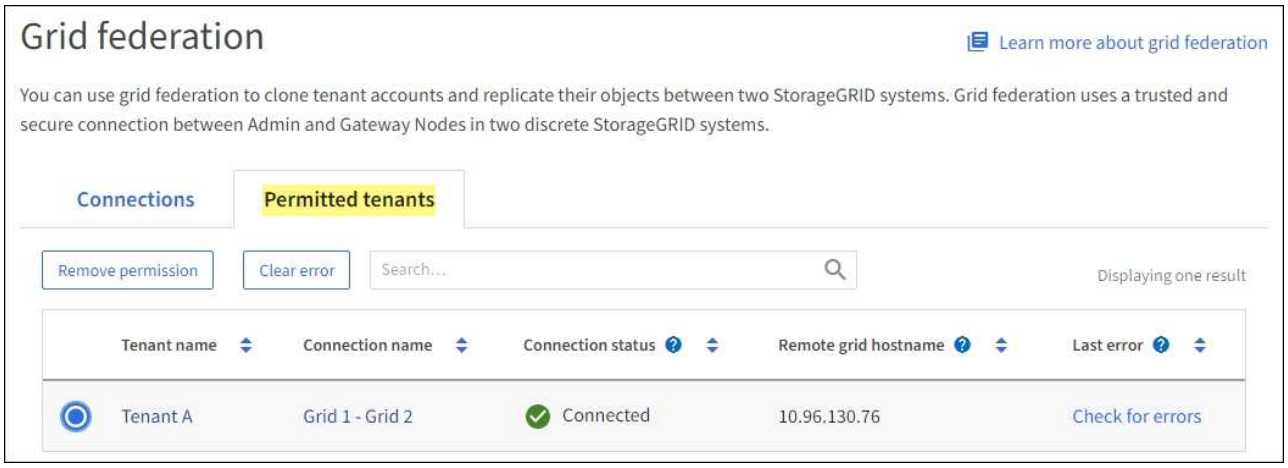


### View a specific connection

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.

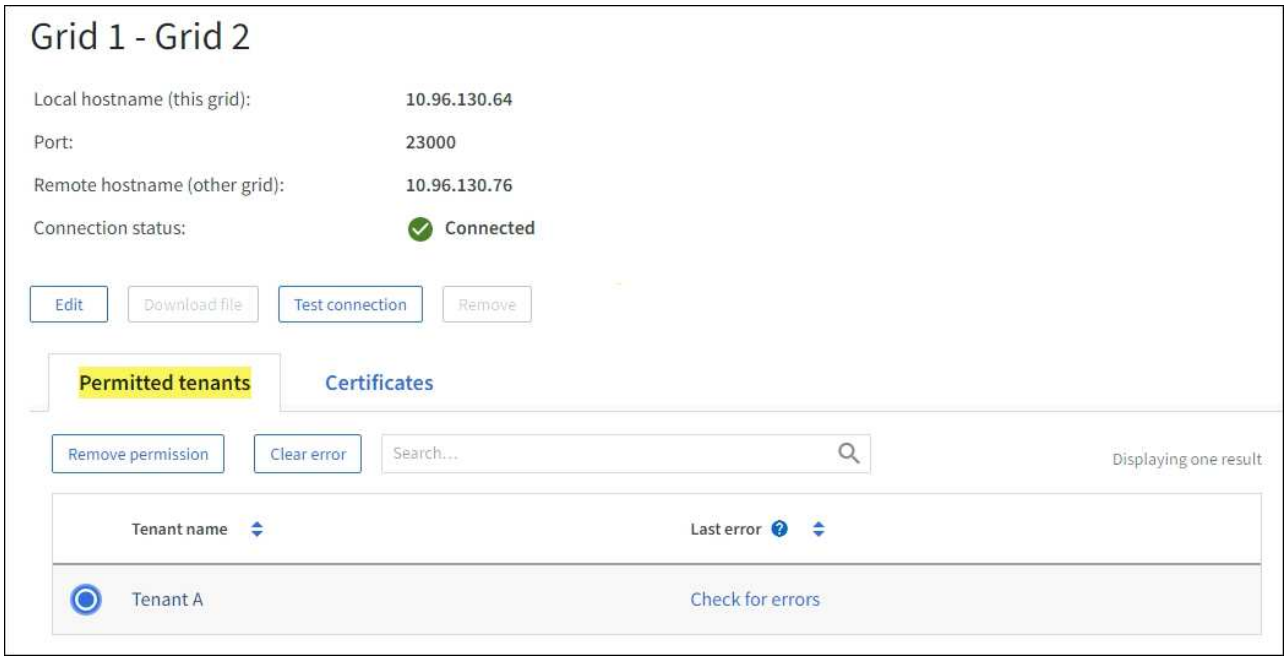


3. When viewing a specific connection, select the **Certificates** tab to view the system-generated server and client certificates for this connection.

   From this tab, you can:

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

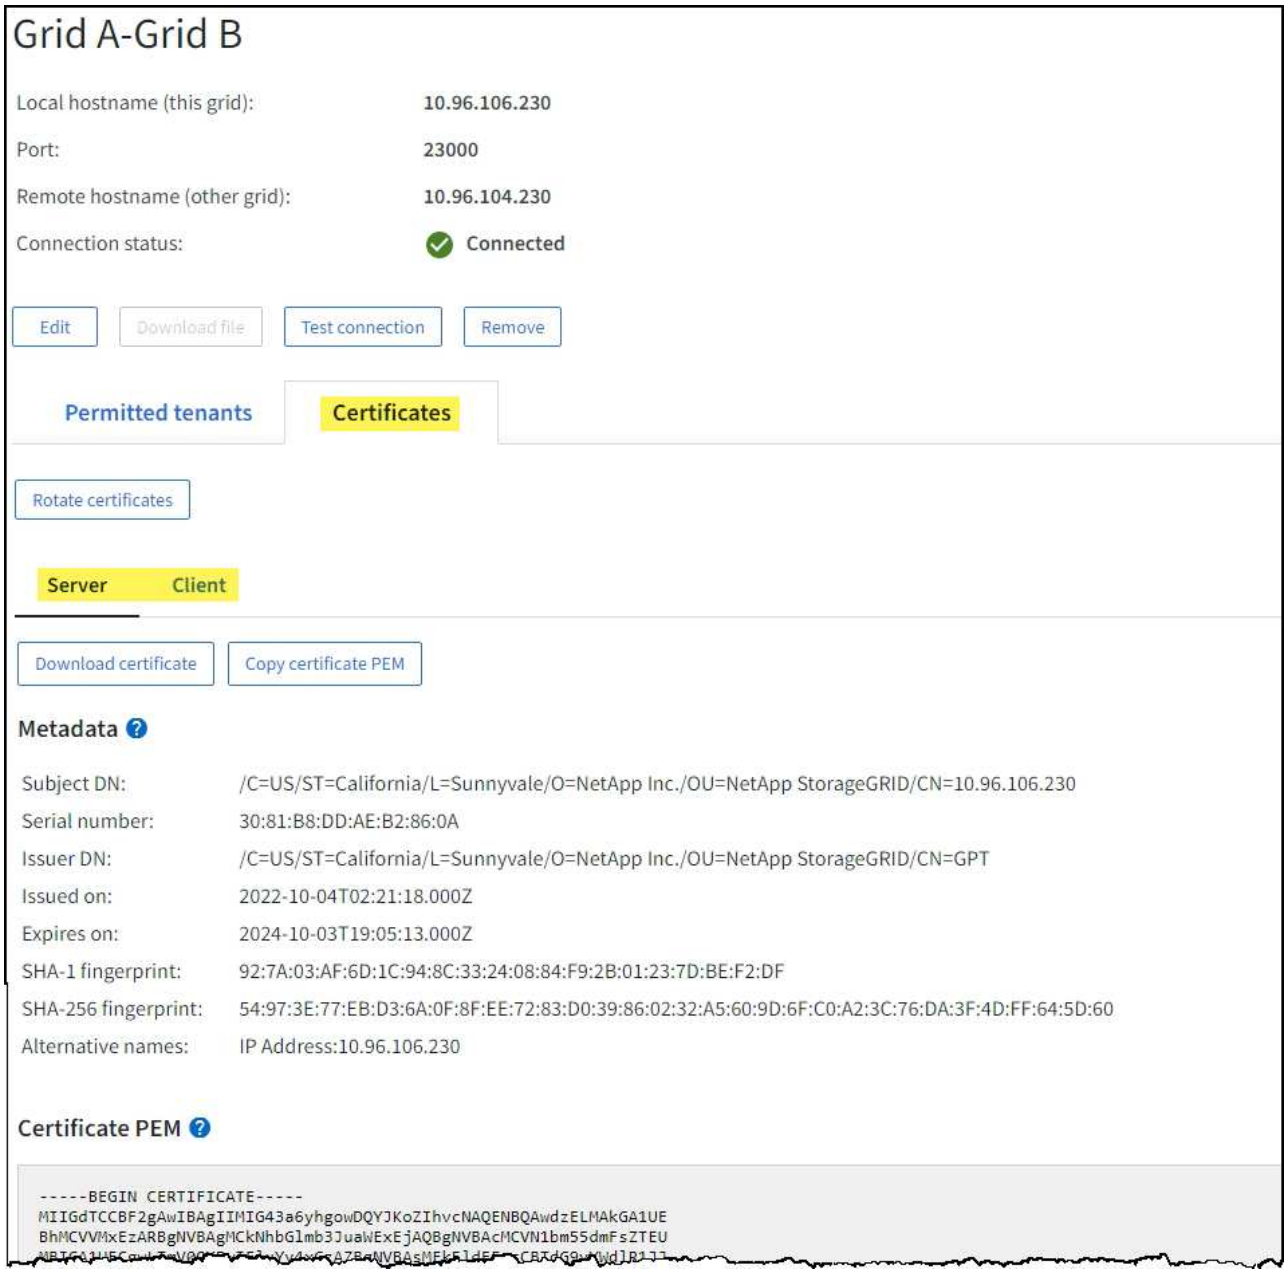


### Review cross-grid replication metrics

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.

## Apply hotfixes or upgrade software if necessary

If a hotfix or a new version of StorageGRID software is available, you should assess whether the update is appropriate for your system, and install it if required.

### About this task

StorageGRID hotfixes contain software changes that are made available outside of a feature or patch release. The same changes are included in a future release.

### Steps

1. Go to NetApp Downloads: StorageGRID.
2. Select the down arrow for the **Type/Select Version** field to see a list of the updates that are available to download:
   - **StorageGRID software versions**: 11.*x.y*
   - **StorageGRID hotfixes**: 11.*x.y.z*
3. Review the changes that are included in the update:
   a. Select the version from the menu, and select **Go**.
   b. Sign in using the username and password for your NetApp account.
   c. Read and accept the End User License Agreement.

      The downloads page for the version you selected appears.
4. Learn about the changes included in the software version or hotfix.
   - For a new software version, see What's new for the version you selected.
   - For a hotfix, download the README file for a summary of the changes included in the hotfix.
5. If you decide a software update is required, locate the instructions before proceeding.
   - For a new software version, carefully follow the instructions for upgrading software.
   - For a hotfix, see the StorageGRID hotfix procedure.

## 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.

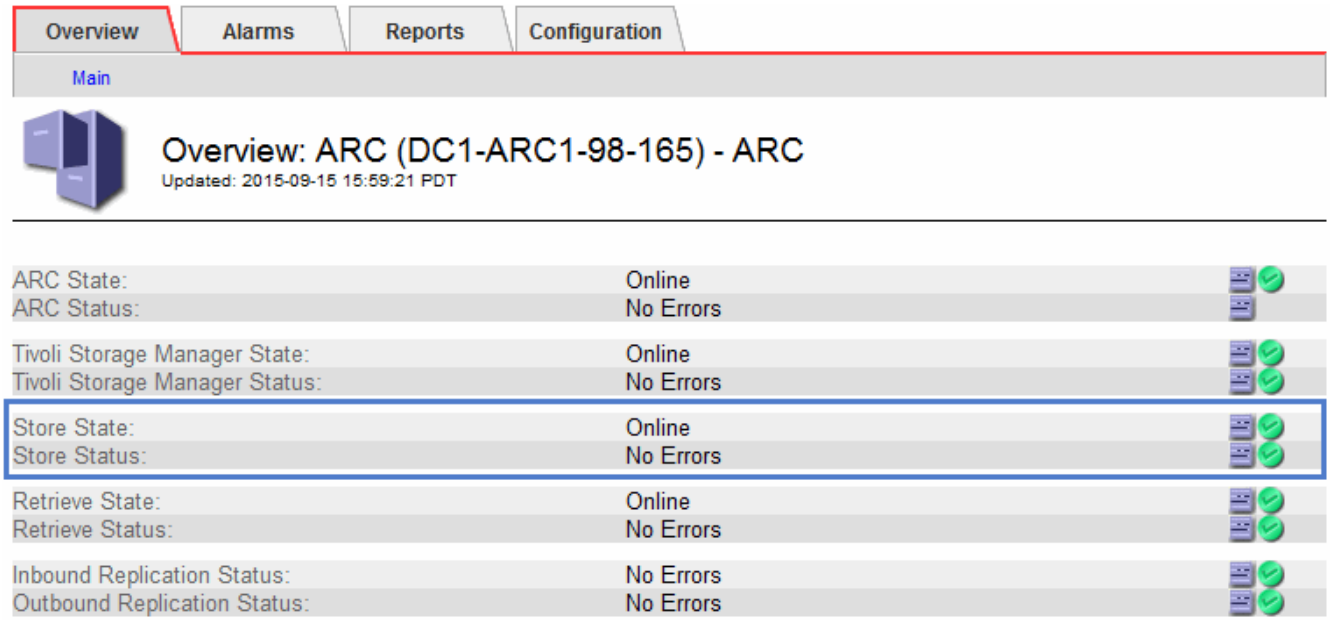


An offline Store component or one with errors might indicate that targeted archival storage system can no longer accept object data because it has reached capacity.

# 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.<br>• Select the alert on the **NODES** > **Overview** page.<br>• Select **ALERTS** > **Current**.<br><br>View current alerts | Select **SUPPORT** > **Alarms (legacy)** > **Current alarms**.<br><br>Manage 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.<br><br>View alert rules | Alarms are triggered when a StorageGRID attribute reaches a threshold value.<br><br>Manage alarms (legacy system) |

| | Alerts | 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.<br><br>As required, additional information is provided in the StorageGRID documentation.<br><br>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.<br><br>Alarms reference (legacy system) |
| Where can I see a list of alerts or alarms that have been resolved? | Select **ALERTS** > **Resolved**.<br><br>View current and resolved alerts | Select **SUPPORT** > **Alarms (legacy)** > **Historical alarms**.<br><br>Manage alarms (legacy system) |
| Where do I manage the settings? | Select **ALERTS** > **Rules**.<br><br>Manage alerts | Select **SUPPORT**. Then, use the options in the **Alarms (legacy)** section of the menu.<br><br>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.<br>• You must have the Manage alerts permission to manage silences, alert notifications, and alert rules.<br><br>Administer StorageGRID | • Anyone who can sign in to the Grid Manager can view legacy alarms.<br>• You must have the Acknowledge alarms permission to acknowledge alarms.<br>• You must have both the Grid topology page configuration and Other grid configuration permissions to manage global alarms and email notifications.<br><br>Administer StorageGRID |
| How do I manage email notifications? | Select **ALERTS** > **Email setup**.<br><br>**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.<br><br>Set up email notifications for alerts | Select **SUPPORT** > **Alarms (legacy)** > **Legacy email setup**.<br><br>Manage alarms (legacy system) |

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

| | Alerts | Alarms (legacy system) |
|---|---|---|
| How do I customize the conditions and triggers? | 1. Select **ALERTS** > **Rules**.<br>2. Select a default rule to edit, or select **Create custom rule**.<br><br>Edit alert rules<br><br>Create custom alert rules | 1. Select **SUPPORT** > **Alarms (legacy)** > **Global alarms**.<br>2. Create a Global Custom alarm to override a Default alarm or to monitor an attribute that does not have a Default alarm.<br><br>Manage alarms (legacy system) |
| How do I disable an individual alert or alarm? | 1. Select **ALERTS** > **Rules**.<br>2. Select the rule, and select **Edit rule**.<br>3. Clear the **Enabled** checkbox.<br><br>Disable alert rules | 1. Select **SUPPORT** > **Alarms (legacy)** > **Global alarms**.<br>2. Select the rule, and select the Edit icon.<br>3. Clear the **Enabled** checkbox.<br><br>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 11.7

- Review the video: Video: Using metrics to create custom alerts in StorageGRID 11.7

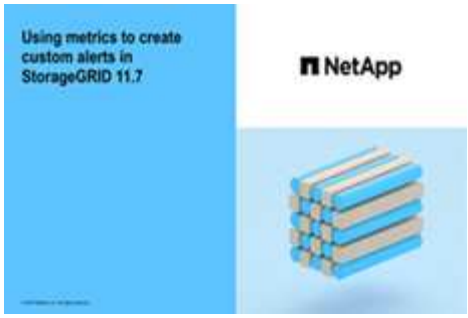


- See the Alerts reference.

### 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 11.7

#### 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**<br>The battery in the appliance's storage controller has expired. | storagegrid_appliance_component_failure{type="REC_EXPIRED_BATTERY"}<br>*Major* > 0 | Default | Enabled |
| **Appliance battery failed**<br>The battery in the appliance's storage controller has failed. | storagegrid_appliance_component_failure{type="REC_FAILED_BATTERY"}<br>*Major* > 0 | Default | Enabled |
| **Appliance battery has insufficient learned capacity**<br>The battery in the appliance's storage controller has insufficient learned capacity. | storagegrid_appliance_component_failure{type="REC_BATTERY_WARN"}<br>*Major* > 0 | Default | Enabled |
| **Appliance battery near expiration**<br>The battery in the appliance's storage controller is nearing expiration. | storagegrid_appliance_component_failure{type="REC_BATTERY_NEAR_EXPIRATION"}<br>*Major* > 0 | Default | Enabled |
| **Appliance battery removed**<br>The battery in the appliance's storage controller is missing. | storagegrid_appliance_component_failure{type="REC_REMOVED_BATTERY"}<br>*Major* > 0 | Default | Enabled |
| **Appliance battery too hot**<br>The battery in the appliance's storage controller is overheated. | storagegrid_appliance_component_failure{type="REC_BATTERY_OVERTEMP"}<br>*Major* > 0 | Default | Enabled |
| **Appliance cache backup device failed**<br>A persistent cache backup device has failed. | storagegrid_appliance_component_failure{type="REC_CACHE_BACKUP_DEVICE_FAILED"}<br>*Major* > 0 | Default | Enabled |
| **Appliance cache backup device insufficient capacity**<br>There is insufficient cache backup device capacity. | storagegrid_appliance_component_failure{type="REC_CACHE_BACKUP_DEVICE_INSUFFICIENT_CAPACITY"}<br>*Major* > 0 | Default | Enabled |
| **Appliance cache backup device write-protected**<br>A cache backup device is write-protected. | storagegrid_appliance_component_failure{type="REC_CACHE_BACKUP_DEVICE_WRITE_PROTECTED"}<br>*Major* > 0 | Default | Enabled |
| **Appliance cache memory size mismatch**<br>The two controllers in the appliance have different cache sizes. | storagegrid_appliance_component_failure{type="REC_CACHE_MEM_SIZE_MISMATCH"}<br>*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.<br>• **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.<br>• **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.<br>• **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:<br>• **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.<br>• **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.<br>• **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 11.7

#### 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.

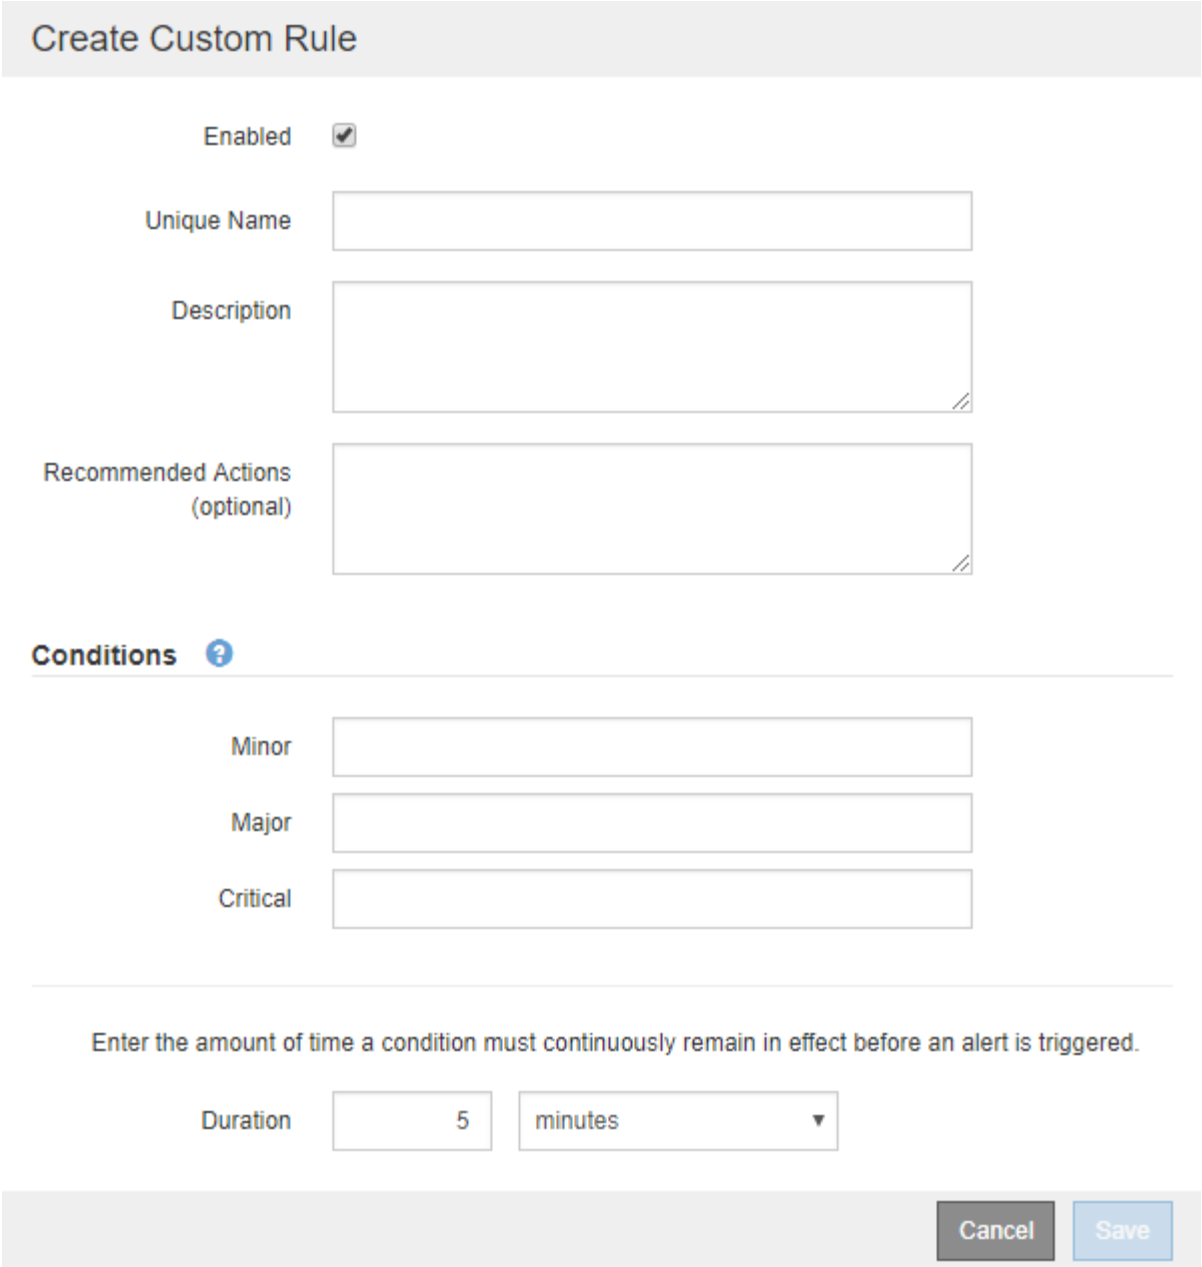


3. Select or clear the **Enabled** checkbox to determine if this alert rule is currently enabled.

   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.

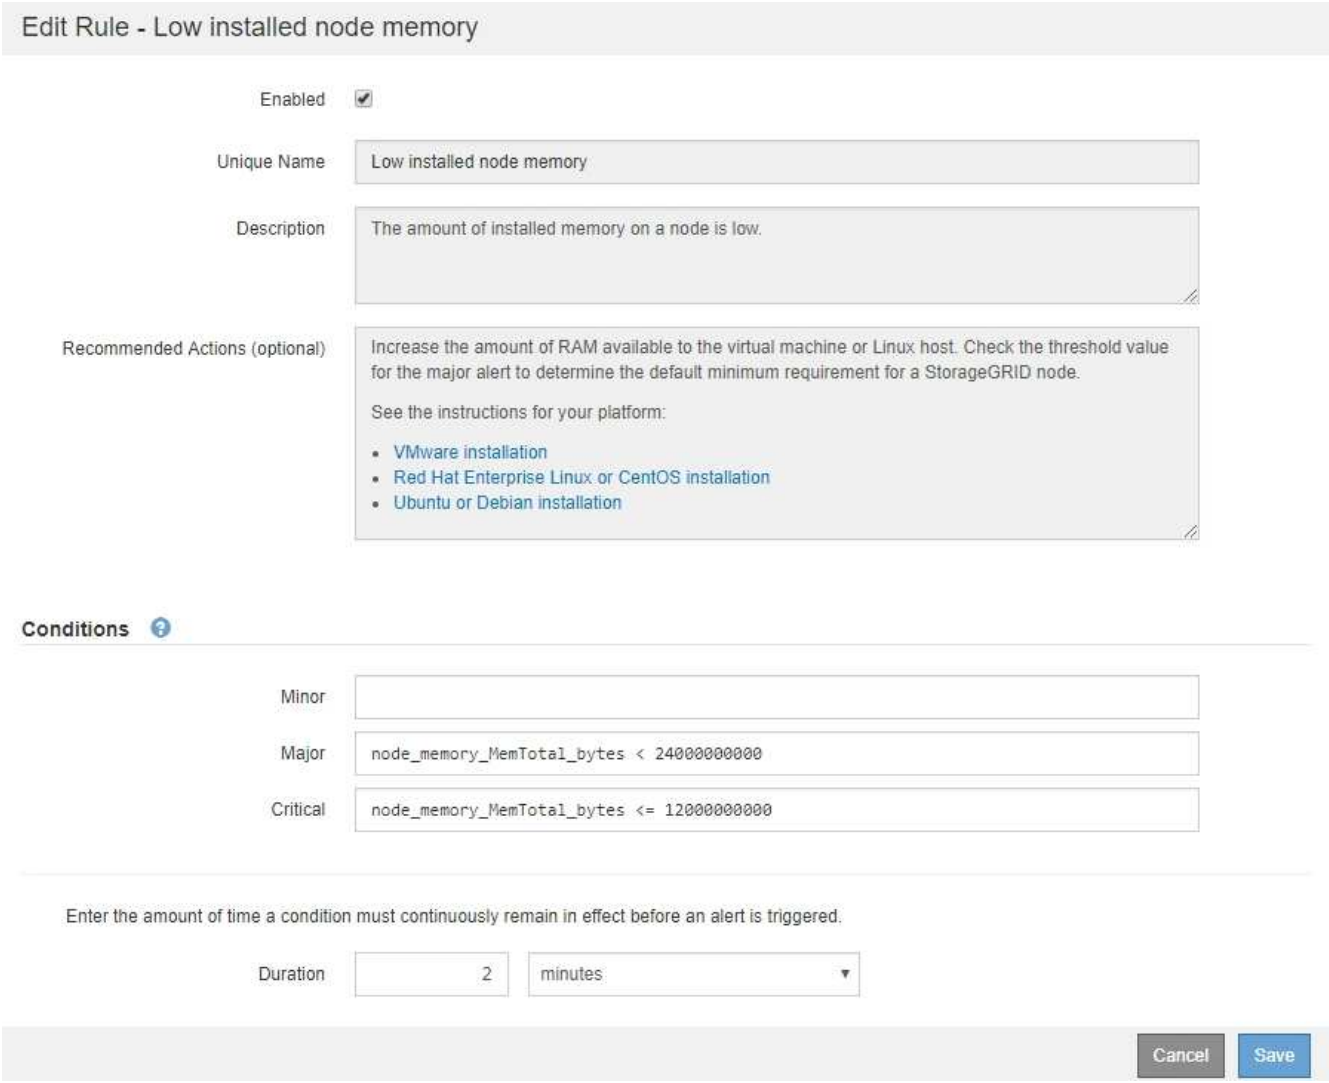


4. Select or clear the **Enabled** checkbox to determine if this alert rule is currently enabled.

   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.

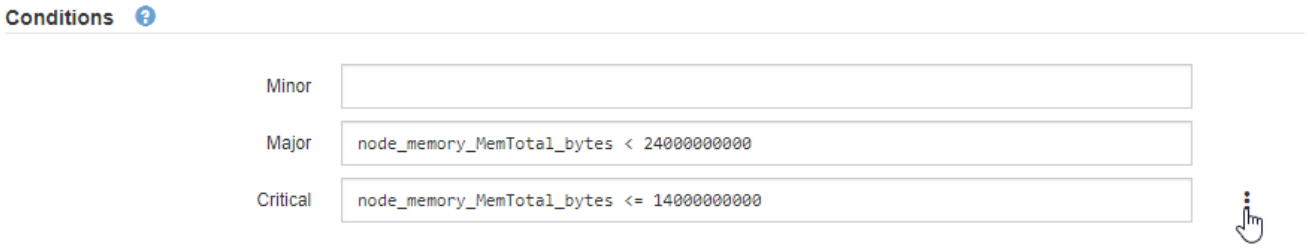


> If you update the conditions for a current alert, your changes might not be implemented until the previous condition is resolved. The next time one of the conditions for the rule is met, the alert will reflect the updated values.

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 messages, 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 messages. 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 messages, 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.

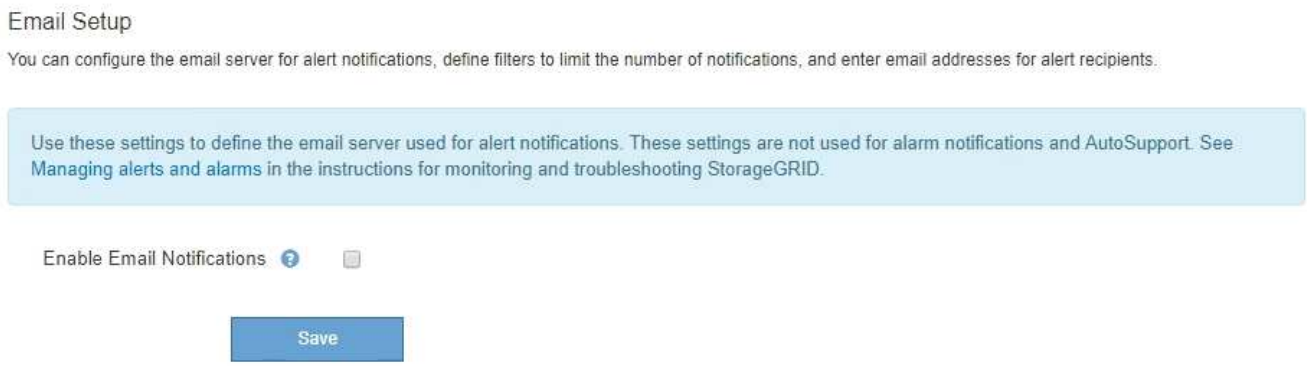


2. Select the **Enable Email Notifications** checkbox to indicate that you want notification emails to be sent when alerts reach configured thresholds.

   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.

**Transport Layer Security (TLS)**

Require TLS ☑

CA Certificate

```
-----BEGIN CERTIFICATE-----
1234567890abcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890
-----END CERTIFICATE-----
```

Browse

Send Client Certificate ☑

Client Certificate

```
-----BEGIN CERTIFICATE-----
1234567890abcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890
-----END CERTIFICATE-----
```

Browse

Private Key

```
-----BEGIN PRIVATE KEY-----
1234567890abcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890
-----BEGIN PRIVATE KEY-----
```

Browse

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 messages 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:

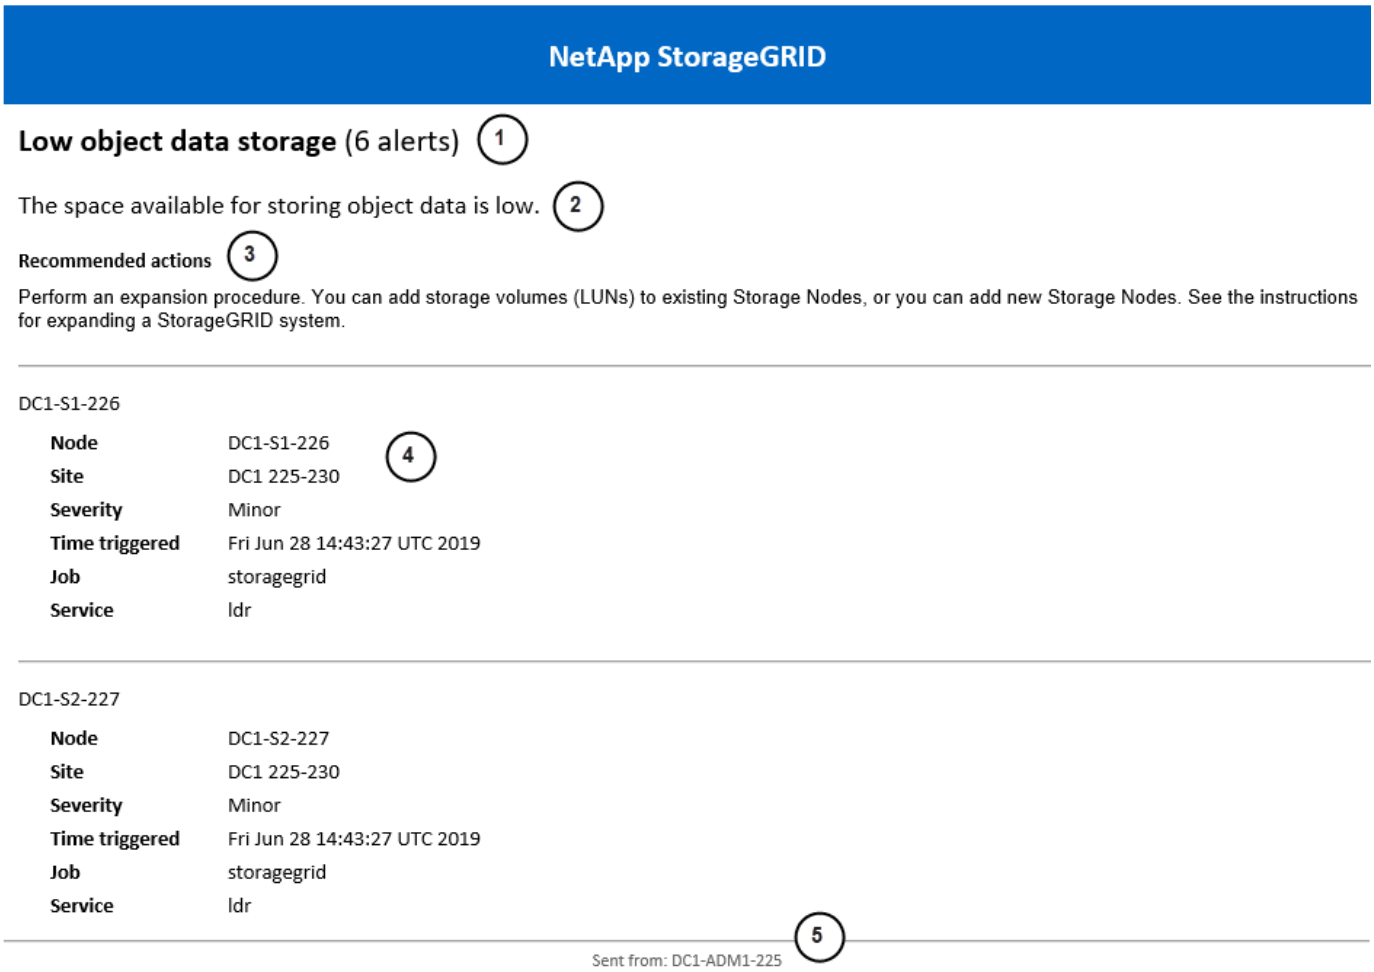


| Callout | Description |
|---|---|
| 1 | The name of the alert, followed by the number of active instances of this alert. |
| 2 | The description of the alert. |
| 3 | Any recommended actions for the alert. |
| 4 | Details about each active instance of the alert, including the node and site affected, 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.<br>• 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.<br>2. Alert A is triggered on node 2 at 08:01. No notification is sent.<br>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.<br>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.<br>2. Alert A is triggered for node 2. A second notification is sent.<br>3. Alert A is resolved for node 2, but it remains active for node 1.<br>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.<br>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.

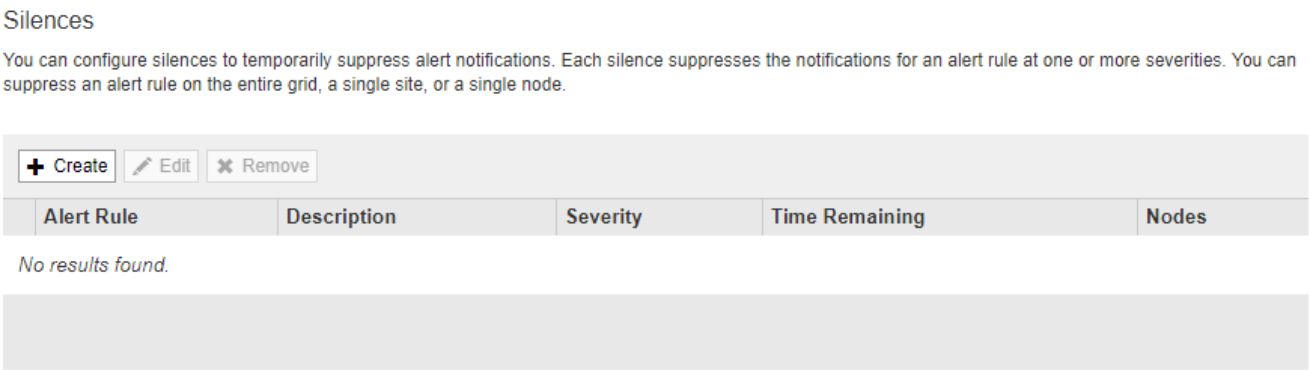


2. Select **Create**.

   The Create Silence dialog box appears.

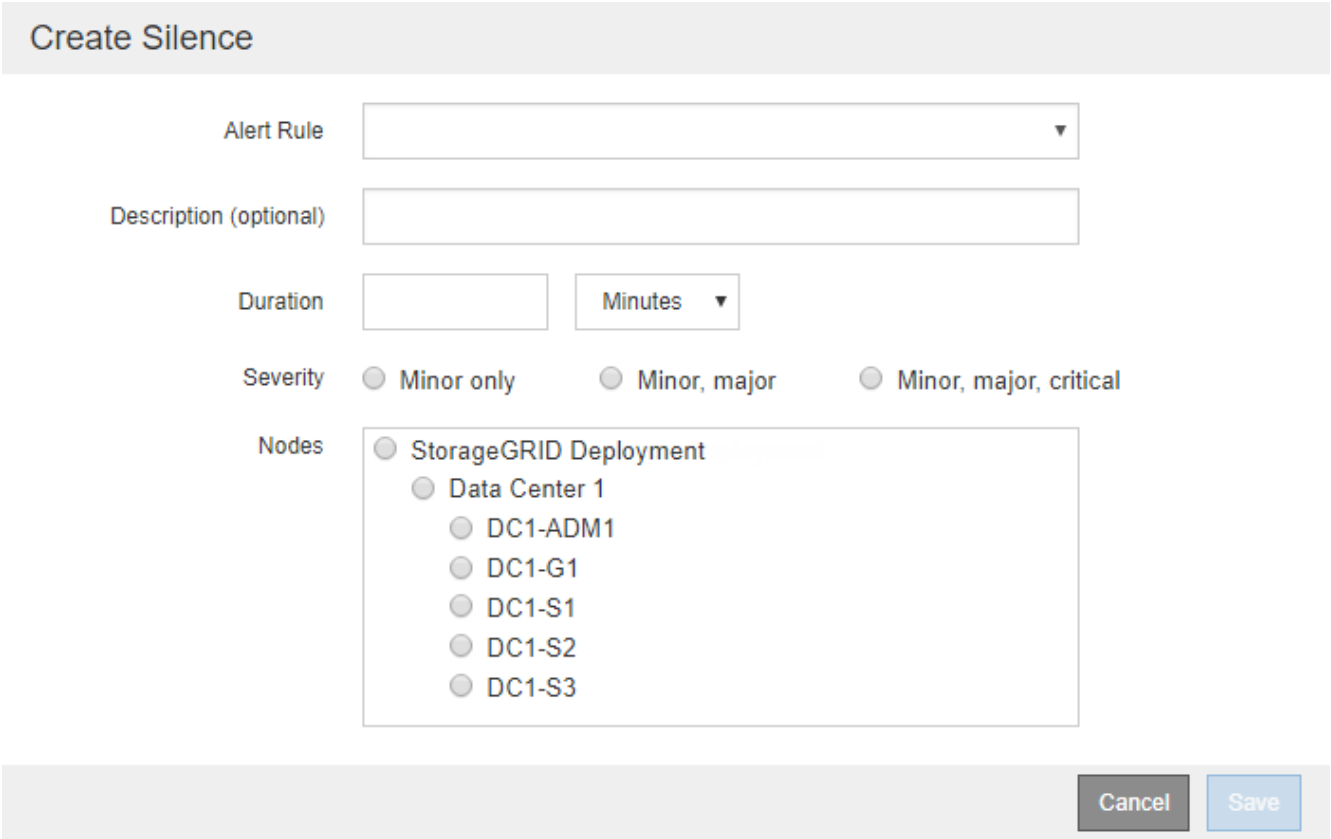


3. Select or enter the following information:

| 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.<br><br>**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).<br><br>**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.<br><br>**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**.<br>b. From the table, select the radio button for the silence you want to edit.<br>c. Select **Edit**.<br>d. Change the description, the amount of time remaining, the selected severities, or the affected node.<br>e. Select **Save**. |
| Remove a silence | a. Select **ALERTS** > **Silences**.<br>b. From the table, select the radio button for the silence you want to remove.<br>c. Select **Remove**.<br>d. Select **OK** to confirm you want to remove this silence.<br><br>**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 fault detected | A problem was detected with a direct-attached storage (DAS) drive in the appliance. |
| 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. |
| 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. |

| Alert name | Description |
| --- | --- |
| 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. |
| 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. |
| Storage device inaccessible | A storage device can’t be accessed. |

### Audit and syslog alerts

| Alert name | Description |
| --- | --- |
| Audit logs are being added to the in-memory queue | Node can’t send logs to the local syslog server and the in-memory queue is filling up. |

| Alert name | Description |
| --- | --- |
| External syslog server forwarding error | Node can’t 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 can’t 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 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. |
| 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. |

| Alert name | Description |
| --- | --- |
| Improved read availability disabled | When improved read availability is disabled, GET and HEAD requests might fail when Storage Nodes are unavailable. |

### 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 |
| --- | --- |
| 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 can’t 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 |

| Alert name | Description |
| --- | --- |
| 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. |
| 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. |

| Alert name | Description |
| --- | --- |
| 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 StorageGRID system data on the /var/local mount point is low. |
| Low tmp directory free space | The space available in the /tmp directory is low. |

**Node or node network alerts**

| Alert name | Description |
| --- | --- |
| Firewall configuration failure | Failed to apply firewall configuration. |
| 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. |
| 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. |

| Alert name | Description |
| --- | --- |
| 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 can’t transition to the desired state because of an internal error or volume related issue |
| Unable to communicate with node | One or more services are unresponsive, or the node can’t 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. |

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

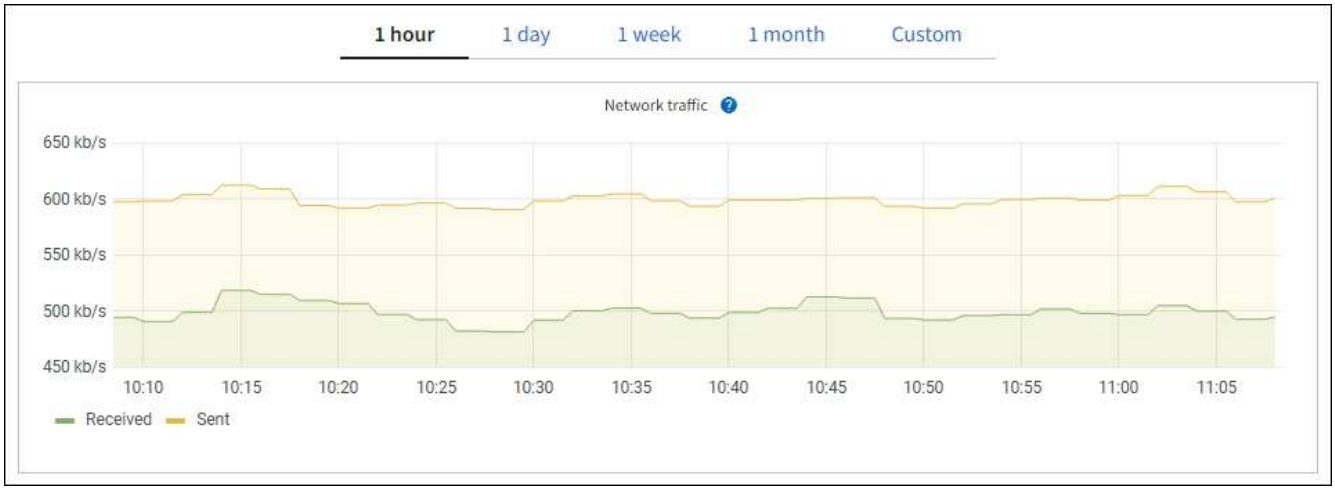


- **Alerts**: Alerts are triggered at specific severity levels when alert rule conditions that use Prometheus metrics evaluate as true.
- **Grid Management API**: You can use Prometheus metrics in custom alert rules or with external automation tools to monitor your StorageGRID system. A complete list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

| | Global Custom alarm threshold (enabled) | Default alarm threshold (enabled) |
|---|---|---|
| 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).

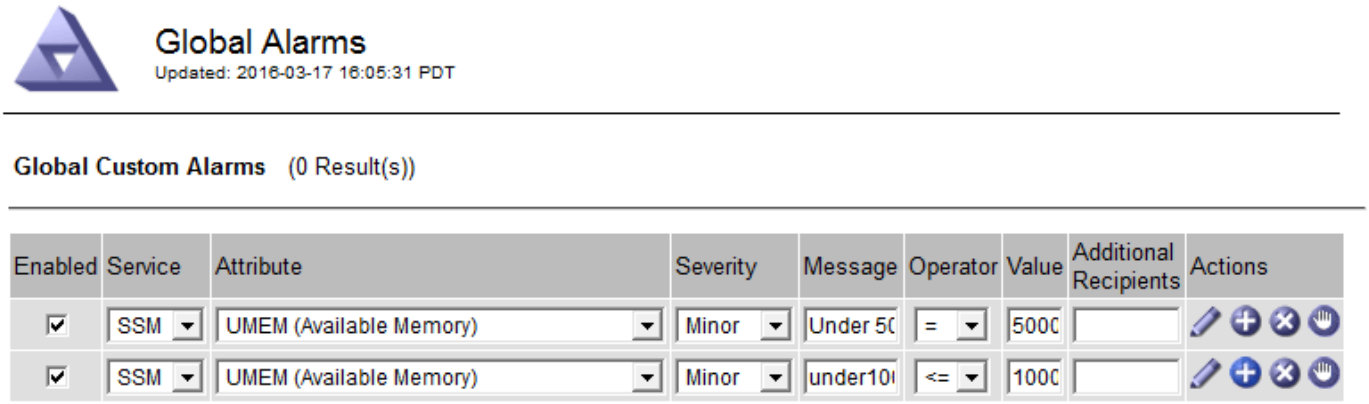


If the order is reversed, when UMEM drops to 100MB, the first alarm (<=100000000) is triggered, but not the one below it (= 50000000).

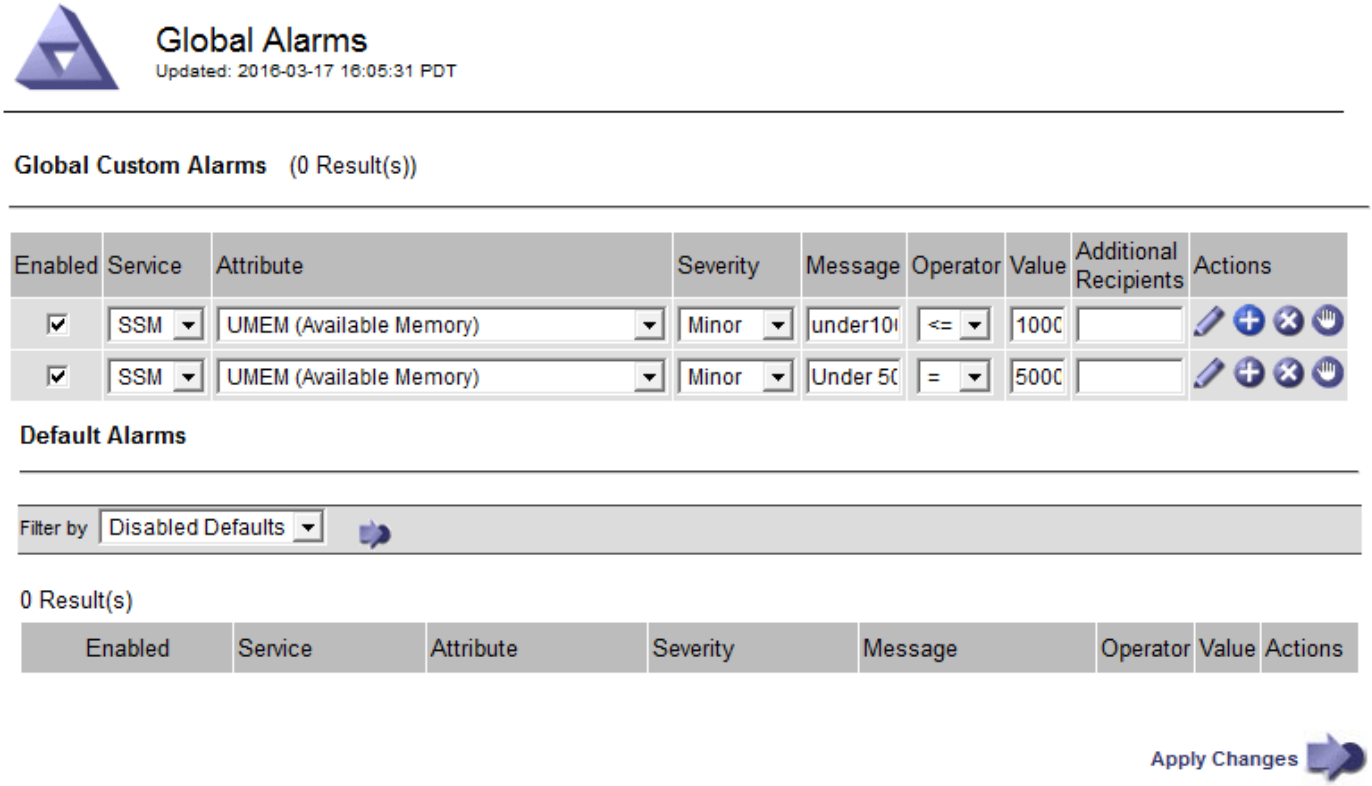


### Notifications

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**).

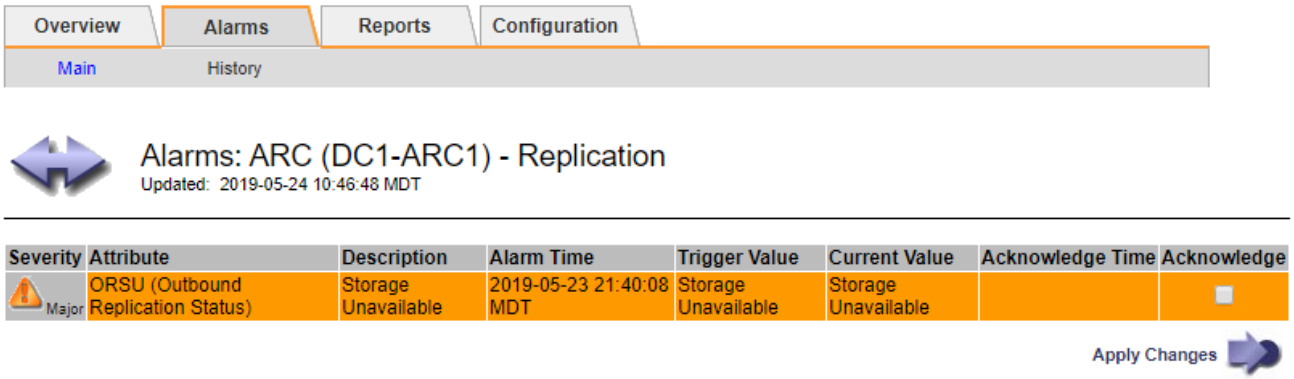


3. Select the **Acknowledge** checkbox for the alarm, and click **Apply Changes**.

   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.

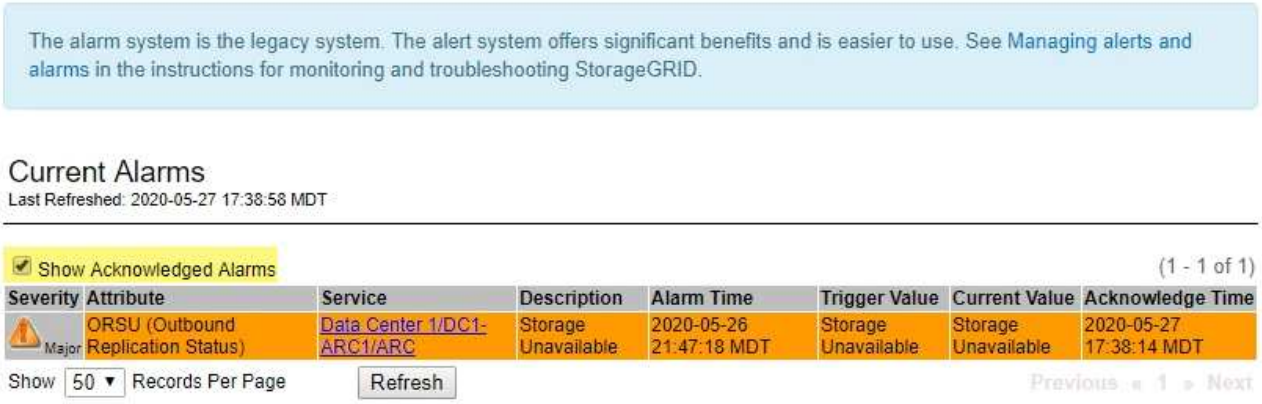


#### View Default alarms (legacy system)

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 must 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.

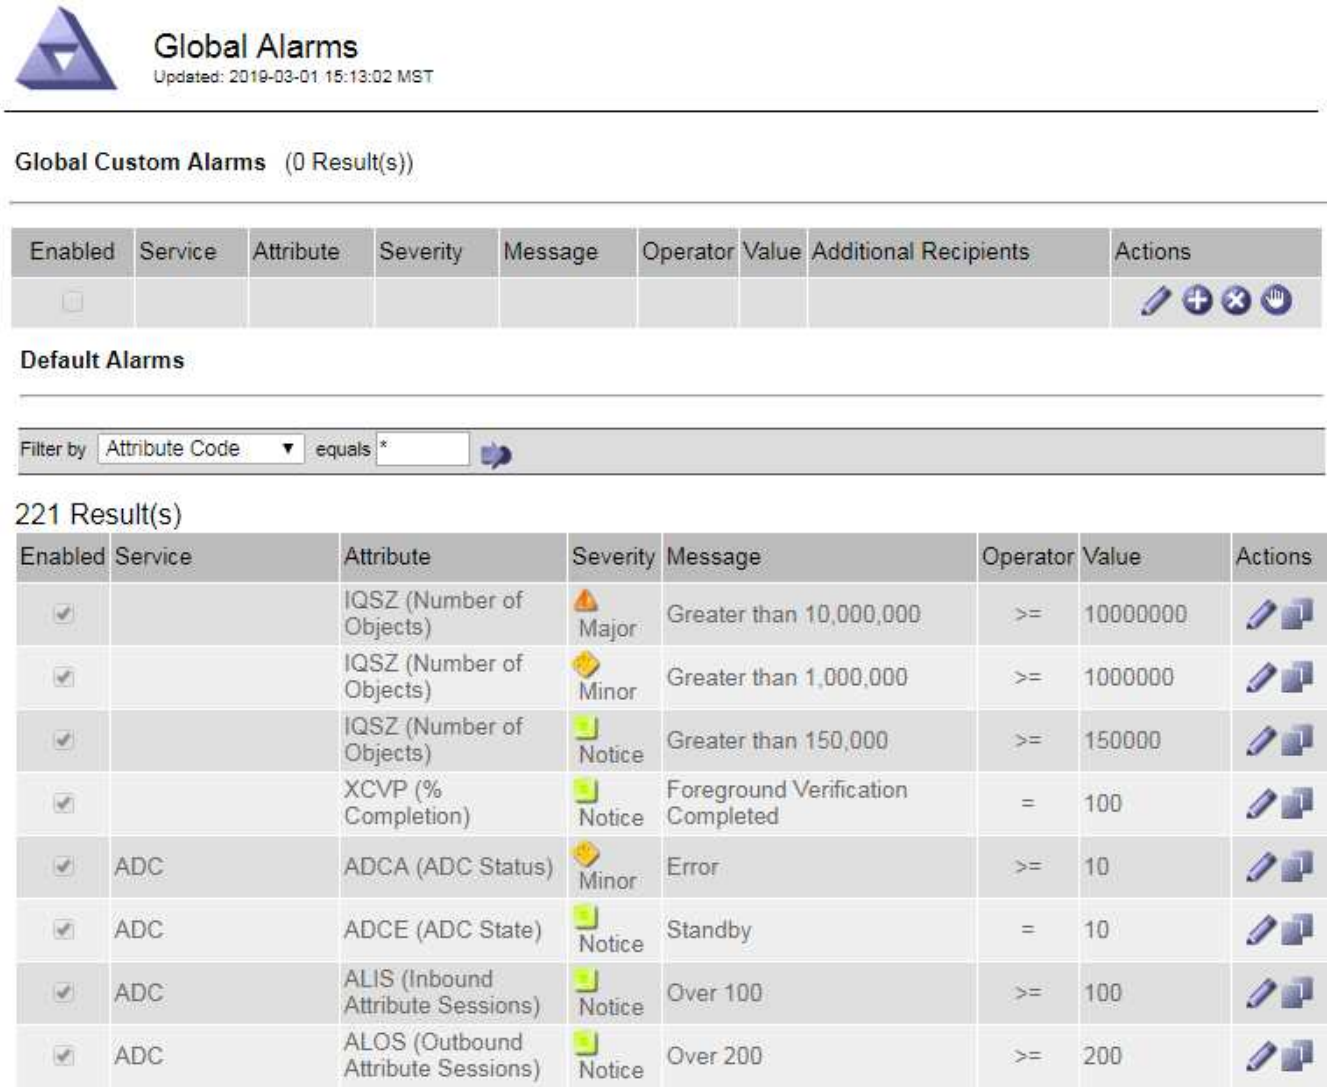


#### Review historical alarms and alarm frequency (legacy system)

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 must 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 must 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**.

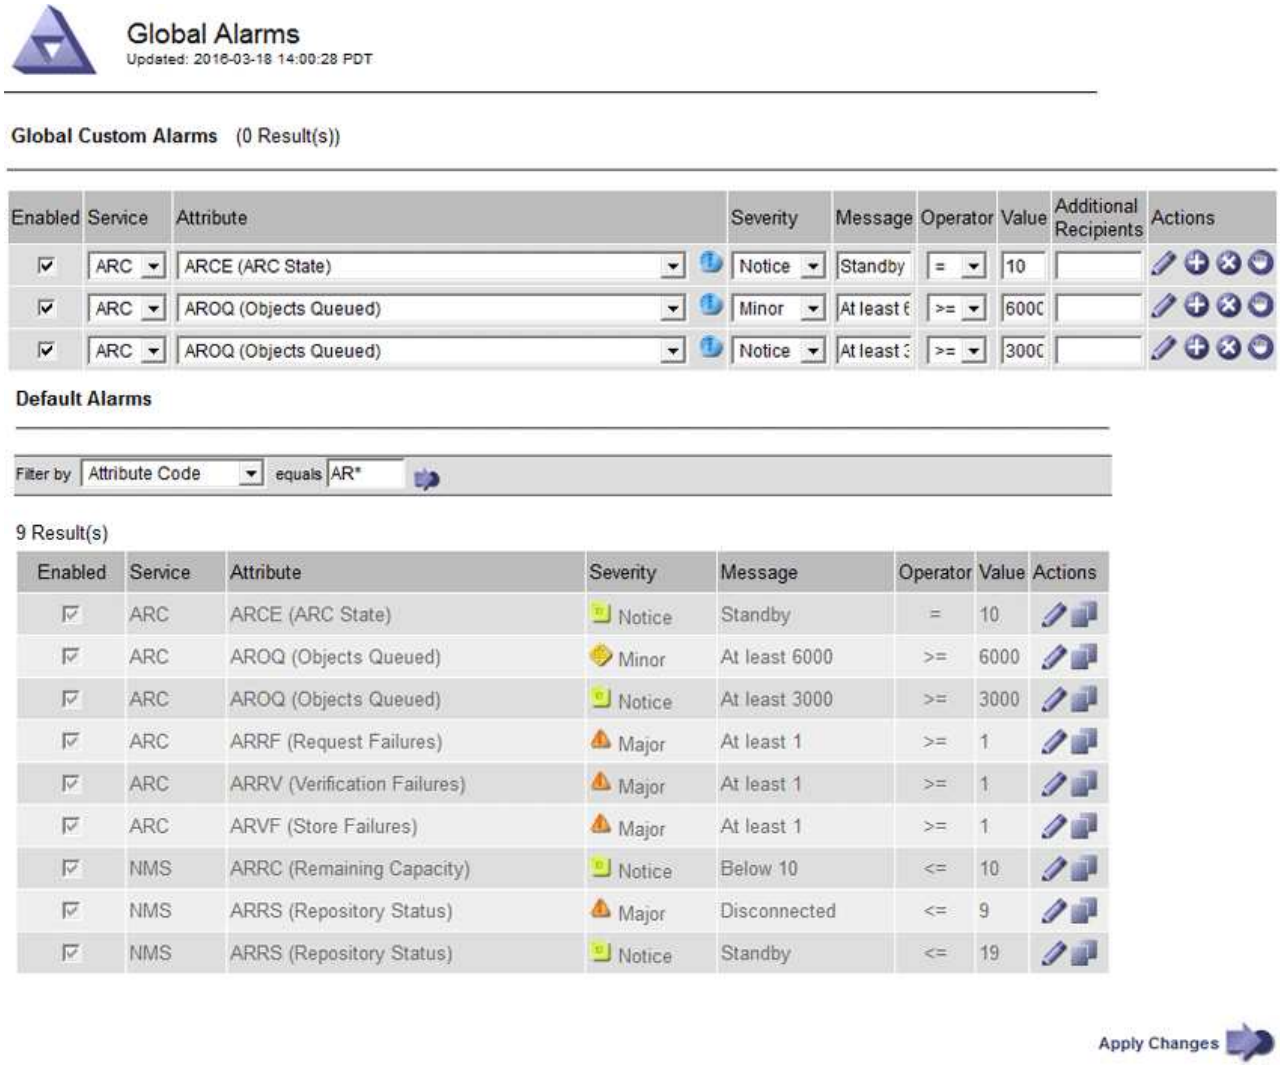


- To modify a Default alarm, search for the Default alarm.
  i. Under Filter by, select either **Attribute Code** or **Attribute Name**.
  ii. Type a search string.

     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:<br>• = equals<br>• > greater than<br>• < less than<br>• >= greater than or equal to<br>• <= less than or equal to<br>• ≠ 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.<br>**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<br>+ Insert a row<br>+ Delete a row<br>+ Drag a row up or down<br>+ 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 must 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.

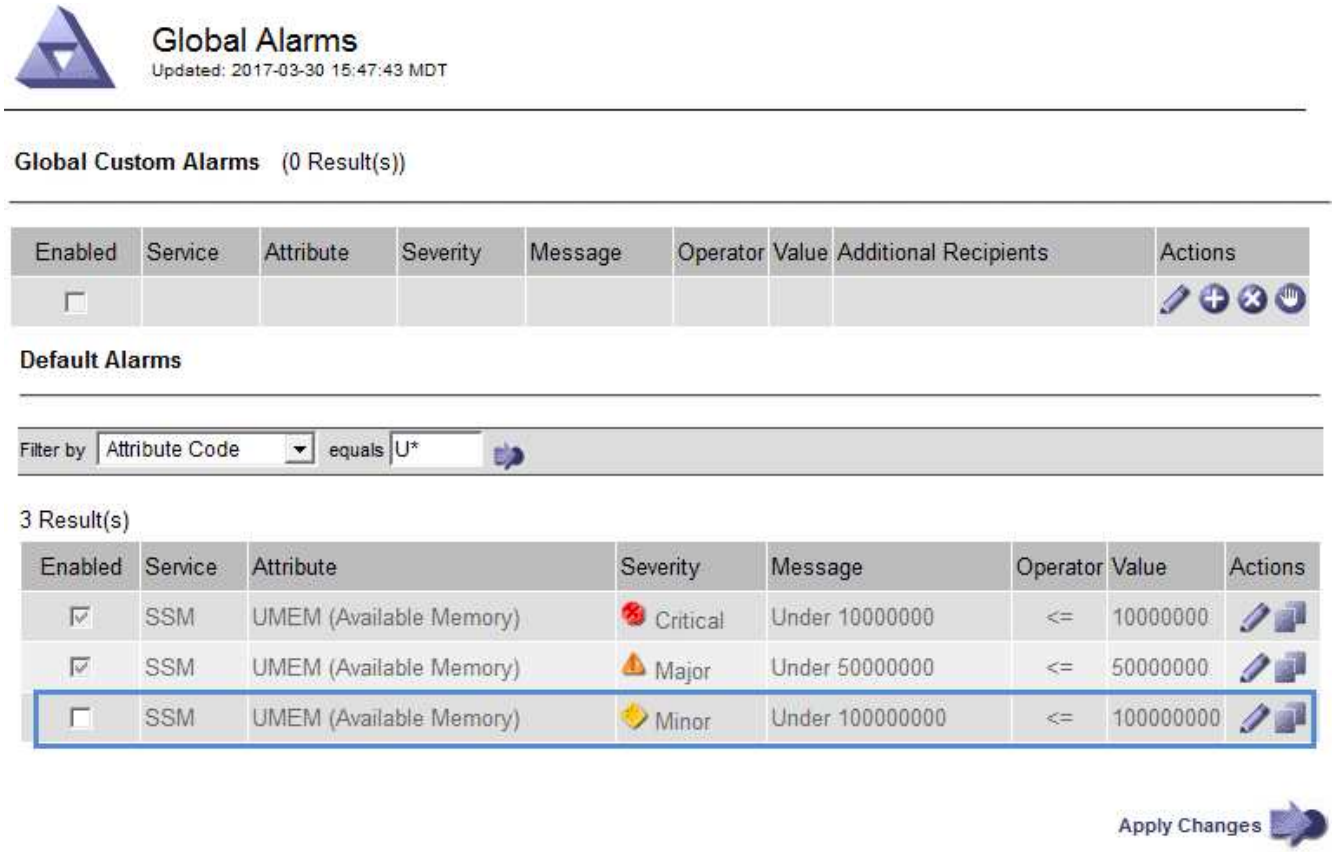


The **Enabled** checkbox for the selected alarm becomes active.

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 must 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.

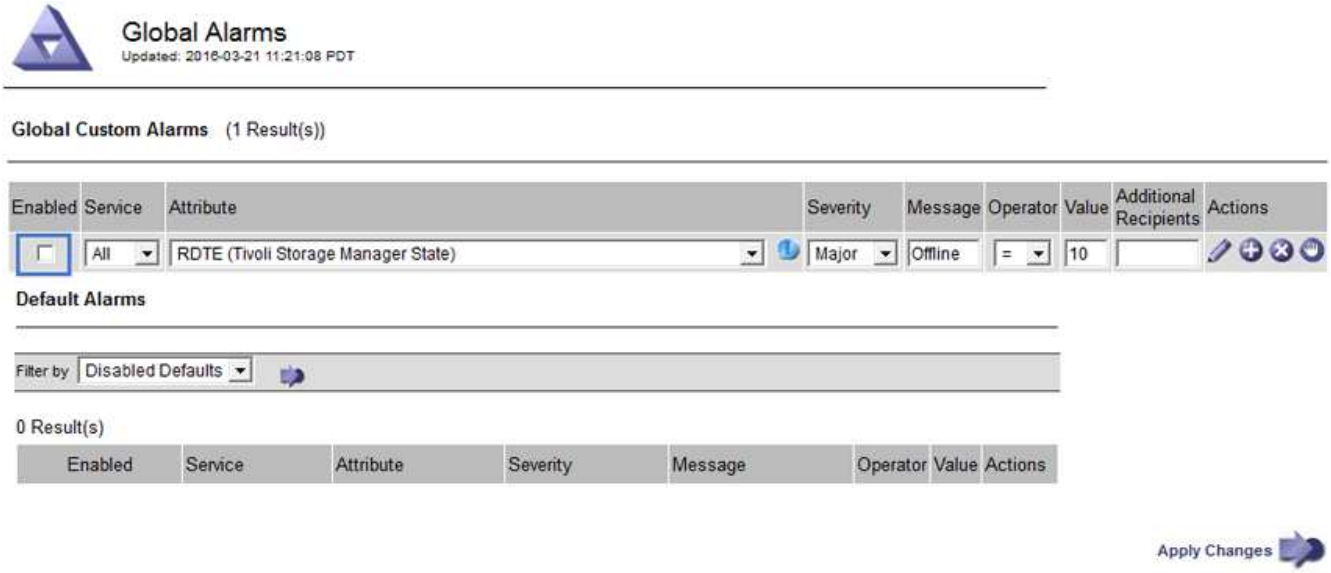


4. Click **Apply Changes**.

   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 must 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 messages, 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 messages.

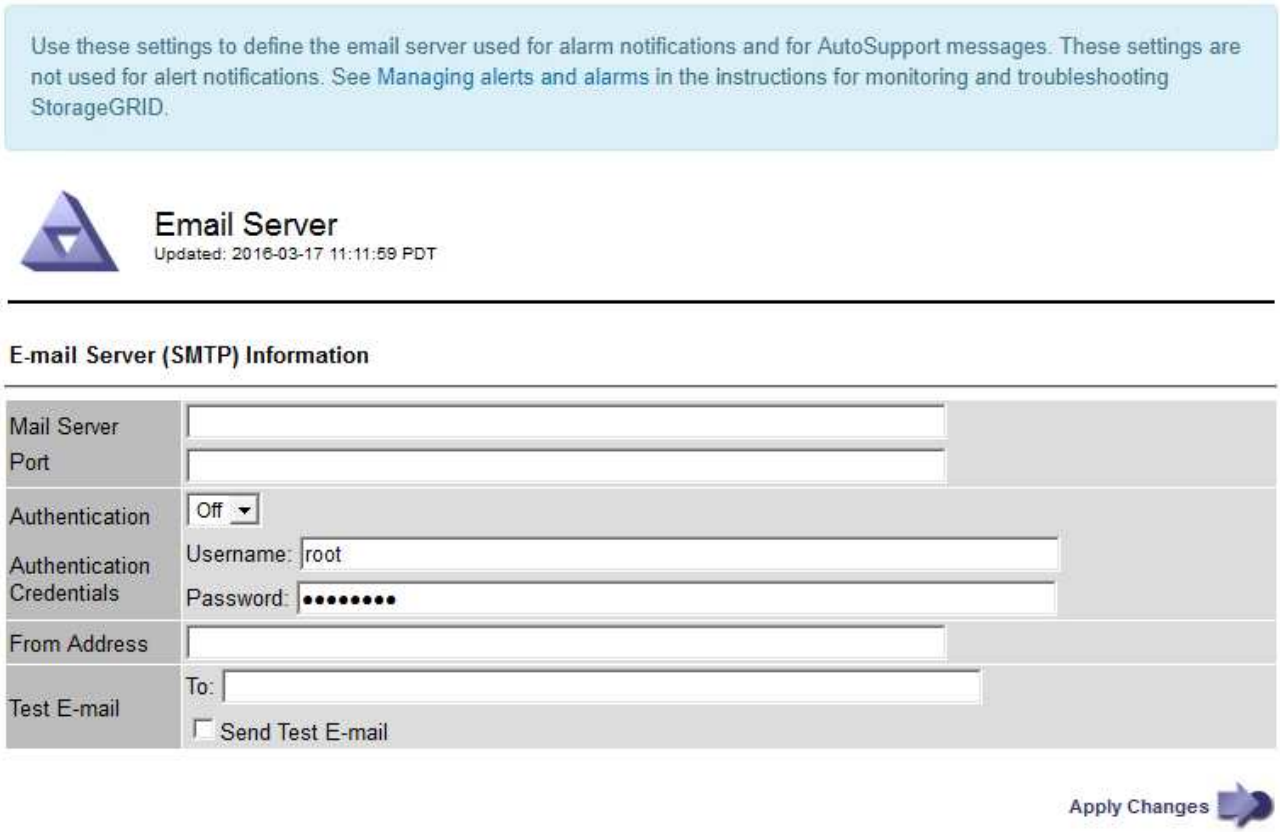


3. Add the following SMTP mail server settings:

| 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 must 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).

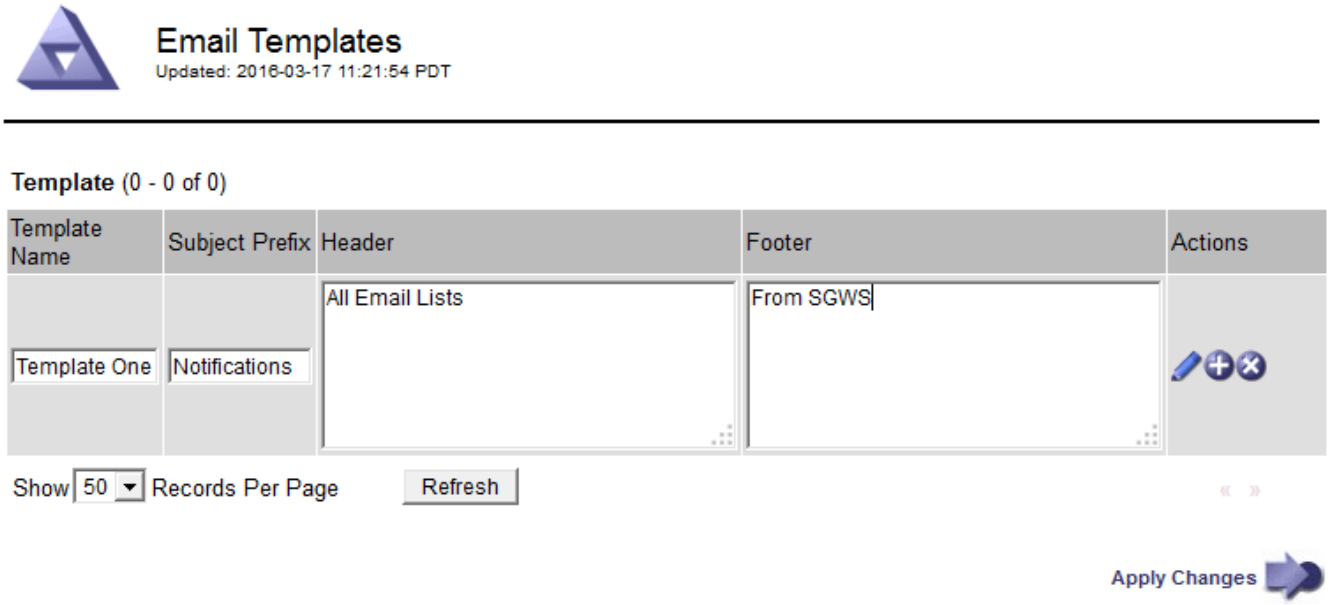


4. In the new row add the following:

| 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 must 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).

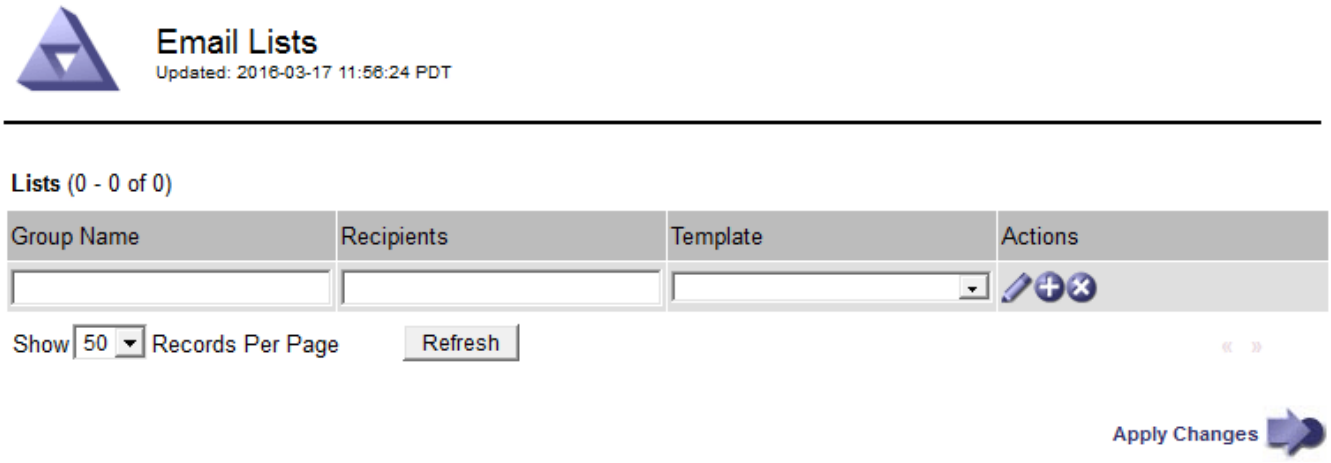


4. In the new row, add the following:

| Item | Description |
|---|---|
| Group Name | Unique name used to identify the mailing list. Mailing list names can't be duplicated.<br><br>**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.<br><br>**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 must 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 must 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**).

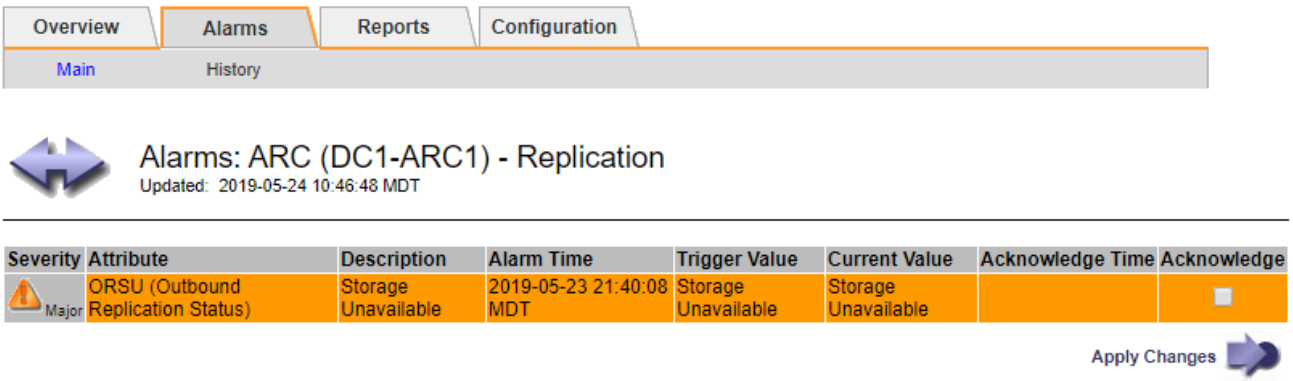


4. If you want to clear the count of current alarms, you can optionally do the following:
   - Acknowledge the alarm. An acknowledged alarm is no longer included in the count of legacy alarms unless it is triggered at the next severity level or it is resolved and occurs again.
   - Disable a particular 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.<br><br>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.<br><br>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).<br><br>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.<br><br>If a major alarm is triggered due to an authentication failure, verify the credentials associated with destination bucket and update values, if necessary.<br><br>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.<br><br>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.<br><br>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).<br><br>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.<br><br>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**.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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:<br><br>• Notice: More than 100,000 messages<br>• Minor: At least 500,000 messages<br>• Major: At least 2,000,000 messages<br>• Critical: At least 5,000,000 messages<br><br>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.<br><br>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).<br><br>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.<br><br>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.<br><br>If the value of Target Error, check the external archival storage system for errors.<br><br>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:<br><br>1. Plug in or replace the cable that does not display any indicator lights.<br>2. Wait one to five minutes.<br><br>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.<br><br>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.<br><br>If the value of ARC State does not transition from Standby to Online, check the status of the ARC components.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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 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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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**.<br><br>If the problem persists, contact technical support. |
| CASA | Data Store Status | DDS | An alarm is raised if the Cassandra metadata store becomes unavailable.<br><br>Check the status of Cassandra:<br><br>1. At the Storage Node, log in as admin and `su` to root using the password listed in the Passwords.txt file.<br>2. Enter: `service cassandra status`<br>3. If Cassandra is not running, restart it: `service cassandra restart`<br><br>This alarm might also indicate that the metadata store (Cassandra database) for a Storage Node requires rebuilding.<br><br>See information about troubleshooting the Services: Status - Cassandra (SVST) alarm in Troubleshoot metadata issues.<br><br>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).<br><br>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 your grid.<br><br>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.<br><br>**Note**: Contact technical support if you are unable to add Storage Nodes.<br><br>After new Storage Nodes are added, the system automatically rebalances object metadata across all Storage Nodes, and the alarm clears.<br><br>Also see information about troubleshooting the Low metadata storage alert in Troubleshoot metadata issues.<br><br>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.<br><br>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).<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>If the Storage Node is a StorageGRID appliance, the StorageGRID system indicates that the controller needs attention.<br><br>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.<br><br>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.<br><br>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:<br><br>• HTTP has been taken offline manually.<br>• The Auto-Start HTTP attribute has been disabled.<br>• The LDR service is shutting down.<br><br>The Auto-Start HTTP attribute is enabled by default. If this setting is changed, HTTP could remain offline after a restart.<br><br>If necessary, wait for the LDR service to restart.<br><br>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.<br><br>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.<br><br>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.<br><br>Check for multiple alarms. An increase in average latency times can be indicated by an excessive number of triggered alarms.<br><br>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.<br><br>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.<br><br>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.<br><br>Investigate LOST (LOST Object) alarms immediately. If the problem persists, contact technical support.<br><br>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.<br><br>1. From the Grid Manager, select **CONFIGURATION** > **Security** > **Certificates**.<br>2. On the **Global** tab, select **Management interface certificate**.<br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>Check the network router, switch, and adapter configurations.<br><br>Restart the server.<br><br>If the problem persists, contact technical support. |

| Code | Name | Service | Recommended action |
|---|---|---|---|
| NRER | Receive Errors | SSM | The following can be causes of NRER alarms:<br>• Forward error correction (FEC) mismatch<br>• Switch port and NIC MTU mismatch<br>• High link error rates<br>• NIC ring buffer overrun<br>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.<br>Restore connectivity to an ADC service as soon as possible.<br>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.<br>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.<br><br>Adjusting the alarm threshold allows you to proactively manage when additional storage needs to be allocated.<br><br>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.<br><br>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.<br><br>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.<br><br>Check the status of Power Supply A or B to determine which power supply is operating abnormally.<br><br>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.<br><br>1. Select **SUPPORT** > **Tools** > **Grid topology**.<br>2. Select **site** > **Storage Node** > **LDR** > **Verification** > **Configuration** > **Main**.<br>3. Select **Delete Quarantined Objects**.<br>4. Click **Apply Changes**.<br><br>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**.<br><br>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.<br><br>Figures under 80% are normal. Over 90% is considered a problem.<br><br>If memory usage is high for a single service, monitor the situation and investigate.<br><br>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.<br><br>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.<br><br>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).<br><br>If the value of Tivoli Storage Manager State is Offline, check Tivoli Storage Manager Status and resolve any problems.<br><br>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).<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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”).<br><br>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**.<br><br>If the value of Status for a grid task being stopped is Error, retry ending the grid task.<br><br>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.<br><br>1. Select **CONFIGURATION** > **Security** > **Certificates**.<br>2. On the **Global** tab, select **S3 and Swift API certificate**.<br>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.<br><br>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.<br><br>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.<br><br>If necessary, replace the component.<br><br>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:<br><br>• problems with the volume being mounted<br>• file system errors |
| SLSA | CPU Load Average | SSM | The higher the value the busier the system.<br><br>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**.<br><br>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.<br><br>When an issue is resolved, reset the counter to clear the alarm. Select **NODES** > ***site*** > ***grid node*** > **Events** > **Reset event counts**.<br><br>To reset event counts, you must have the Grid topology page configuration permission.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>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.<br><br>Additional storage should be added. It is not impacting end user functionality, but the alarm persists until additional storage is added.<br><br>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.<br><br>If the value of Storage Status is Error, contact technical support.<br><br>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.<br><br>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:<br><br>• The service has been manually stopped (`/etc/init.d/<service\> stop`).<br>• There is an issue with the MySQL database and Server Manager shuts down the MI service.<br>• A grid node has been added, but not started.<br>• During installation, a grid node has not yet connected to the Admin Node.<br><br>If a service is listed as Not Running, restart the service (`/etc/init.d/<service\> restart`).<br><br>This alarm might also indicate that the metadata store (Cassandra database) for a Storage Node requires rebuilding.<br><br>If the problem persists, contact technical support.<br><br>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.<br><br>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**.<br><br>Check the operating system for any signs of block-device or file system errors.<br><br>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.<br><br>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. |

| Archive location | Description |
| --- | --- |
| `install` | `Gdu-server` and install logs. |
| `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.<br><br>**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 |

| File name | Notes | Found on |
| --- | --- | --- |
| `/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.<br>• Captures events related to the operation of the NMS service, for example, alarm processing, email notifications, and configuration changes.<br>• Contains XML bundle updates resulting from configuration changes made in the system.<br>• Contains error messages related to the attribute downsampling done once a day.<br>• 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.<br><br>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 |
| `/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 |

| File name | Notes | Found on |
| --- | --- | --- |
| `/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 |
| `/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<br><br>Only certain Admin and Gateway Nodes contain this log. See the S3 Select requirements and limitations for Admin and Gateway Nodes. |

| File name | Notes | Found on |
| --- | --- | --- |
| `/var/local/log/ldr.errlog` | | Storage Nodes |
| `/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.<br><br>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`<br><br>`/var/local/log/mysql-slow.log` | Log files generated by MySQL.<br><br>`mysql.err` captures database errors and events such as startups and shutdowns.<br><br>`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`<br><br>`/var/lib/ntp/var/log/ntpstats/` | `/var/local/log/ntp.log` contains the log file for NTP error messages.<br><br>`/var/lib/ntp/var/log/ntpstats/` directory contains NTP timing statistics.<br><br>`loopstats` records loop filter statistics information.<br><br>`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 messages and log destinations

Audit messages and logs record system activities and security events, and are essential tools for monitoring and troubleshooting. You can adjust audit levels to increase or decrease the type and number of audit messages recorded. Optionally, you can define any HTTP request headers you want to include in client read and write audit messages. You can also configure an external syslog server and change the destination of audit information.

For more information about audit messages, see Review audit logs.

**Before you begin**

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

- You have Maintenance or Root access permissions.

**About this task**

All StorageGRID nodes generate audit messages and logs to track system activity and events. By default, audit information is sent to the audit log on Admin Nodes. You can adjust audit levels to increase or decrease the type and number of audit messages recorded in the audit log. Optionally, you can configure audit information to be stored temporarily on the originating nodes for manual collection.

> If you have a large grid, use multiple types of S3 applications, or want to retain all audit data, configure an external syslog server and save audit information remotely. Using an external server minimizes the performance impact of audit message logging without reducing the completeness of of audit data. See Considerations for external syslog server for details.

## Change audit message levels in the audit log

You can set a different audit level for each of the following categories of messages in the audit log:

| Audit category | Description |
|---|---|
| System | By default, this level is set to Normal. See System audit messages. |
| Storage | By default, this level is set to Error. See Object storage audit messages. |
| Management | By default, this level is set to Normal. See Management audit message. |
| Client reads | By default, this level is set to Normal. See Client read audit messages. |
| Client writes | By default, this level is set to Normal. See Client write audit messages. |
| ILM operations | By default, this level is set to Normal. See ILM operations audit messages. |

> These defaults apply if you initially installed StorageGRID using version 10.3 or later. If you have upgraded from 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). |

| Audit level | Description |
| --- | --- |
| 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. Optionally, under **Audit protocol headers**, define any 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.

   > Audit protocol headers apply to S3 and Swift requests only.

4. 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**.

5. Select **Save**

   A green banner displays indicating your configuration has been saved successfully.

## Use an external syslog server

You can configure an external syslog server if you want to save audit information remotely.

- If you want to save audit information to an external syslog server, go to Configure an external syslog server.
- If you aren't using an external syslog server, go to Select audit information destinations.

## Select audit information destinations

You can specify where audit logs, security event logs, and application logs are sent.

> Some destination are available only if you are using an external syslog server. See Configure an external syslog server to configure an external syslog server.

> For more information about StorageGRID software logs, see StorageGRID software logs.

1. On the Audit and syslog server page, select the destination for audit information from the listed options:

| Option | Description |
| --- | --- |
| Default (Admin nodes/local nodes) | Audit messages are sent to the audit log (`audit.log`) on the Admin Node, and security event logs and application logs are stored on the nodes where they were generated (also referred to as "the local node"). |
| External syslog server | Audit information is sent to an 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. |
| Admin Node and external syslog server | Audit messages are sent to the audit log (`audit.log`) on the Admin Node, 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. |
| Local nodes only | No audit information is sent to an Admin Node or remote syslog server. Audit information is saved only on the nodes that generated it.<br><br>**Note**: StorageGRID periodically removes these 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. |

Audit information generated on every local node is stored in `/var/local/log/localaudit.log`

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 appears notifying you that your audit configuration has been saved.

   New logs are sent to the destinations you selected. Existing logs remain in their current location.

**Related information**

Considerations for external syslog server

Administer StorageGRID

Troubleshoot the external syslog server

# Use an external syslog server

## Considerations for external syslog server

Use the following guidelines to estimate the size of the external syslog server you need.

### What is 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 configure the destinations of your audit information so you can reduce network traffic on your Admin Nodes and manage the information more efficiently. 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

### 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. |

| Code | Field | Description |
| --- | --- | --- |
| 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 will be receive (and typically store) application data at a rate of 0.6 MB per second.

For more information about configuring audit message levels and an external syslog server, see the following:

- Configure an external syslog server
- Configure audit messages and log destinations

## Configure an external syslog server

If you want to save audit logs, application logs, and security event logs to a location outside of your grid, use this procedure to configure an external syslog server.

### Before you begin

- You are signed in to the Grid Manager using a supported web browser.
- You have Maintenance or Root access permissions.
- You have a syslog server with the capacity to receive and store the log files. For more information, see Considerations for external syslog server.
- You have the correct server and client certifications if you plan to use TLS or RELP/TLS.

### About this task

If you want to send audit information to an external syslog server, you must configure the external server first.

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 8192 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 20GB of local logs of audit records (localaudit.log) are maintained on each node.

> If the configuration options available in this procedure aren't flexible enough to meet your requirements, additional configuration options can be applied using the private API `audit-destinations` endpoints. For example, it is possible to use different syslog servers for different groups of nodes.

## Configure the external server

### Access the wizard

To start, access the Configure external syslog server wizard.

**Steps**

1. Select **CONFIGURATION** > **Monitoring** > **Audit and syslog server**.
2. From the Audit and syslog server page, select **Configure external syslog server**. If you have previously configured an external syslog server, select **Edit external syslog server**.

   The Configure external syslog server wizard appears.

### Enter syslog info

You must provide the information StorageGRID needs to access the external syslog server.

**Steps**

1. 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.
2. Enter the destination port on the external syslog server (must be an integer between 1 and 65535). The default port is 514.
3. 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.

4. Select **Continue**.
5. If you selected **TLS** or **RELP/TLS**, upload 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. The file you upload here might be a CA bundle.
   - **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.

a. Select **Browse** for the certificate or key you want to use.
b. Select the certificate file 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.

6. 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`
2. Use the drop-down menus to select the severity and facility (type of message) for the category of audit information you want to send.

   If you select **Passthrough** for severity and facility, the information sent to the remote syslog server will receive the same severity and facility as it did when logged locally onto the node. Setting facility and severity can help you aggregate the logs in customizable ways for easier analysis.

   > For more information about StorageGRID software logs, see StorageGRID software logs.

   a. For **Severity**, select **Passthrough** if you want each message sent to the external syslog to have the same severity value as it does in the local syslog.

      For audit logs, if you select **Passthrough**, the severity is 'info.'

For security events, if you select **Passthrough**, the severity values are generated by the Linux distribution on the nodes.

For application logs, if you select **Passthrough**, 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.'

b. If you don't want to use the passthrough value, select a severity value between 0 and 7.

The selected value will be applied to all messages of this type. Information about different severities will be lost when you choose to override severity with a fixed value.

| Severity | Description |
|---|---|
| 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 |

c. For **Facility**, select **Passthrough** if you want each message sent to the external syslog to have the same facility value as it does in the local syslog.

For audit logs, if you select **Passthrough**, the facility sent to the external syslog server is 'local7.'

For security events, if you select **Passthrough**, the facility values are generated by the linux distribution on the nodes.

For application logs, if you select **Passthrough**, the application logs sent to the external syslog server have the following facility values:

| Application log | Passthrough value |
|---|---|
| bycast.log | user or daemon |
| bycast-err.log | user, daemon, local3, or local4 |
| jaeger.log | local2 |

| Application log | Passthrough value |
| --- | --- |
| nms.log | local3 |
| prometheus.log | local4 |
| raft.log | local5 |
| hagroups.log | local6 |

d. If you don't want to use the passthrough value, select the facility value between 0 and 23.

The selected value will be applied to all messages of this type. Information about different facilities will be lost when you choose to override facility with a fixed value.

| Facility | Description |
| --- | --- |
| 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) |
| 11 | FTP |
| 12 | NTP |
| 13 | logaudit (log audit) |
| 14 | logalert (log alert) |

| Facility | Description |
| --- | --- |
| 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 appears indicating your configuration has been saved successfully.

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 Troubleshooting the 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 appears notifying you that your syslog server configuration has been saved successfully.

   > Your 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 security event logs, application logs, and audit message logs are sent.

> For more information about StorageGRID software logs, see StorageGRID software logs.

#### Steps

1. On the Audit and syslog server page, select the destination for audit information from the listed options:

| Option | Description |
|---|---|
| Default (Admin nodes/local nodes) | Audit messages are sent to the audit log (`audit.log`) on the Admin Node, and security event logs and application logs are stored on the nodes where they were generated (also referred to as "the local node"). |
| External syslog server | Audit information is sent to an 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. |
| Admin Node and external syslog server | Audit messages are sent to the audit log (`audit.log`) on the Admin Node, 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. |
| Local nodes only | No audit information is sent to an Admin Node or remote syslog server. Audit information is saved only on the nodes that generated it.<br><br>**Note**: StorageGRID periodically removes these 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. |

> Audit information generated on every local node is stored in `/var/local/log/localaudit.log`

2. Select **Save**. Then, select **OK** to accept the change to the log destination.
3. If you selected either **External syslog server** or **Admin Nodes and external syslog server** as the destination for audit information, an additional warning appears. Review the warning text.

   > You must confirm that the external syslog server can receive test StorageGRID messages.

4. Select **OK** to confirm that you want to change the destination for audit information.

   A green banner appears notifying you that your audit configuration has been saved successfully.

   New logs are sent to the destinations you selected. Existing logs remain in their current location.

**Related information**

Audit message overview

Configure audit messages and log destinations

System audit messages

Object storage audit messages

Management audit message

Client read audit messages

Administer StorageGRID

# 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** 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 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 | 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 | 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.

**Related information**

- Alerts reference
- Alarms reference (legacy system)
- Silence alert notifications

## Configure the SNMP agent

You can configure the StorageGRID SNMP agent if you want 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 all three versions of the SNMP protocol. You can configure the agent for one or more versions.

**Steps**

1. Select **CONFIGURATION** > **Monitoring** > **SNMP agent**.

   The SNMP Agent page appears.

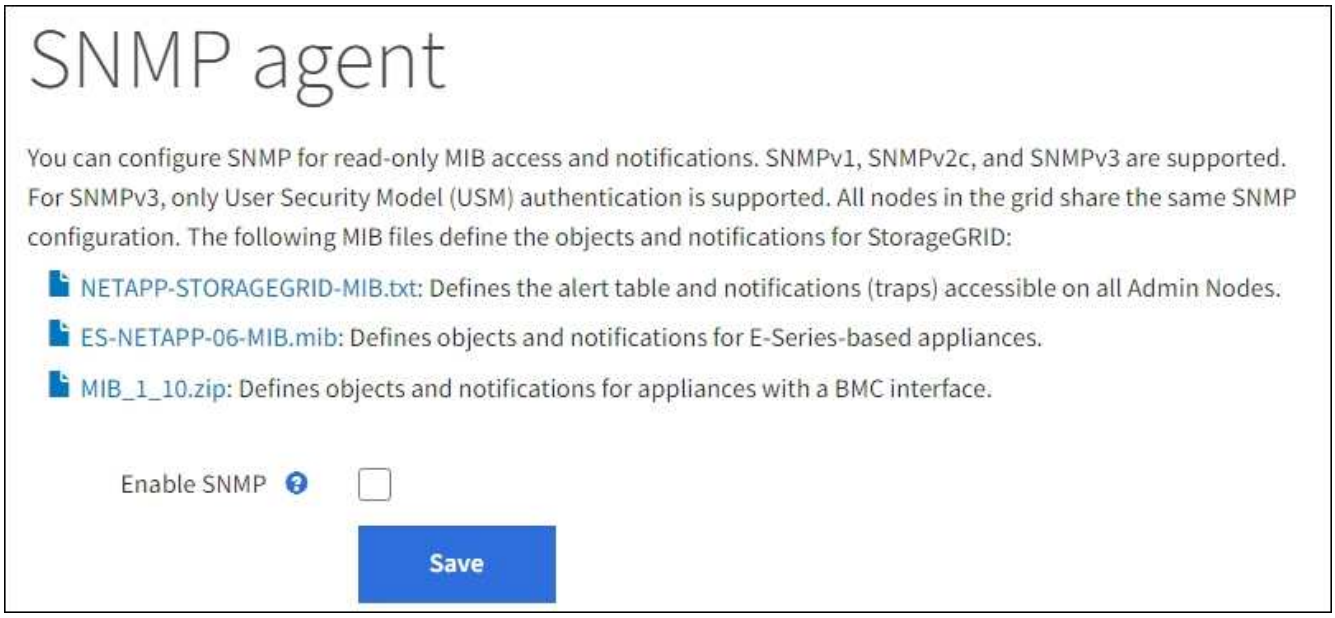


2. To enable the SNMP agent on all grid nodes, select the **Enable SNMP** checkbox.

   The fields for configuring an SNMP agent appear.

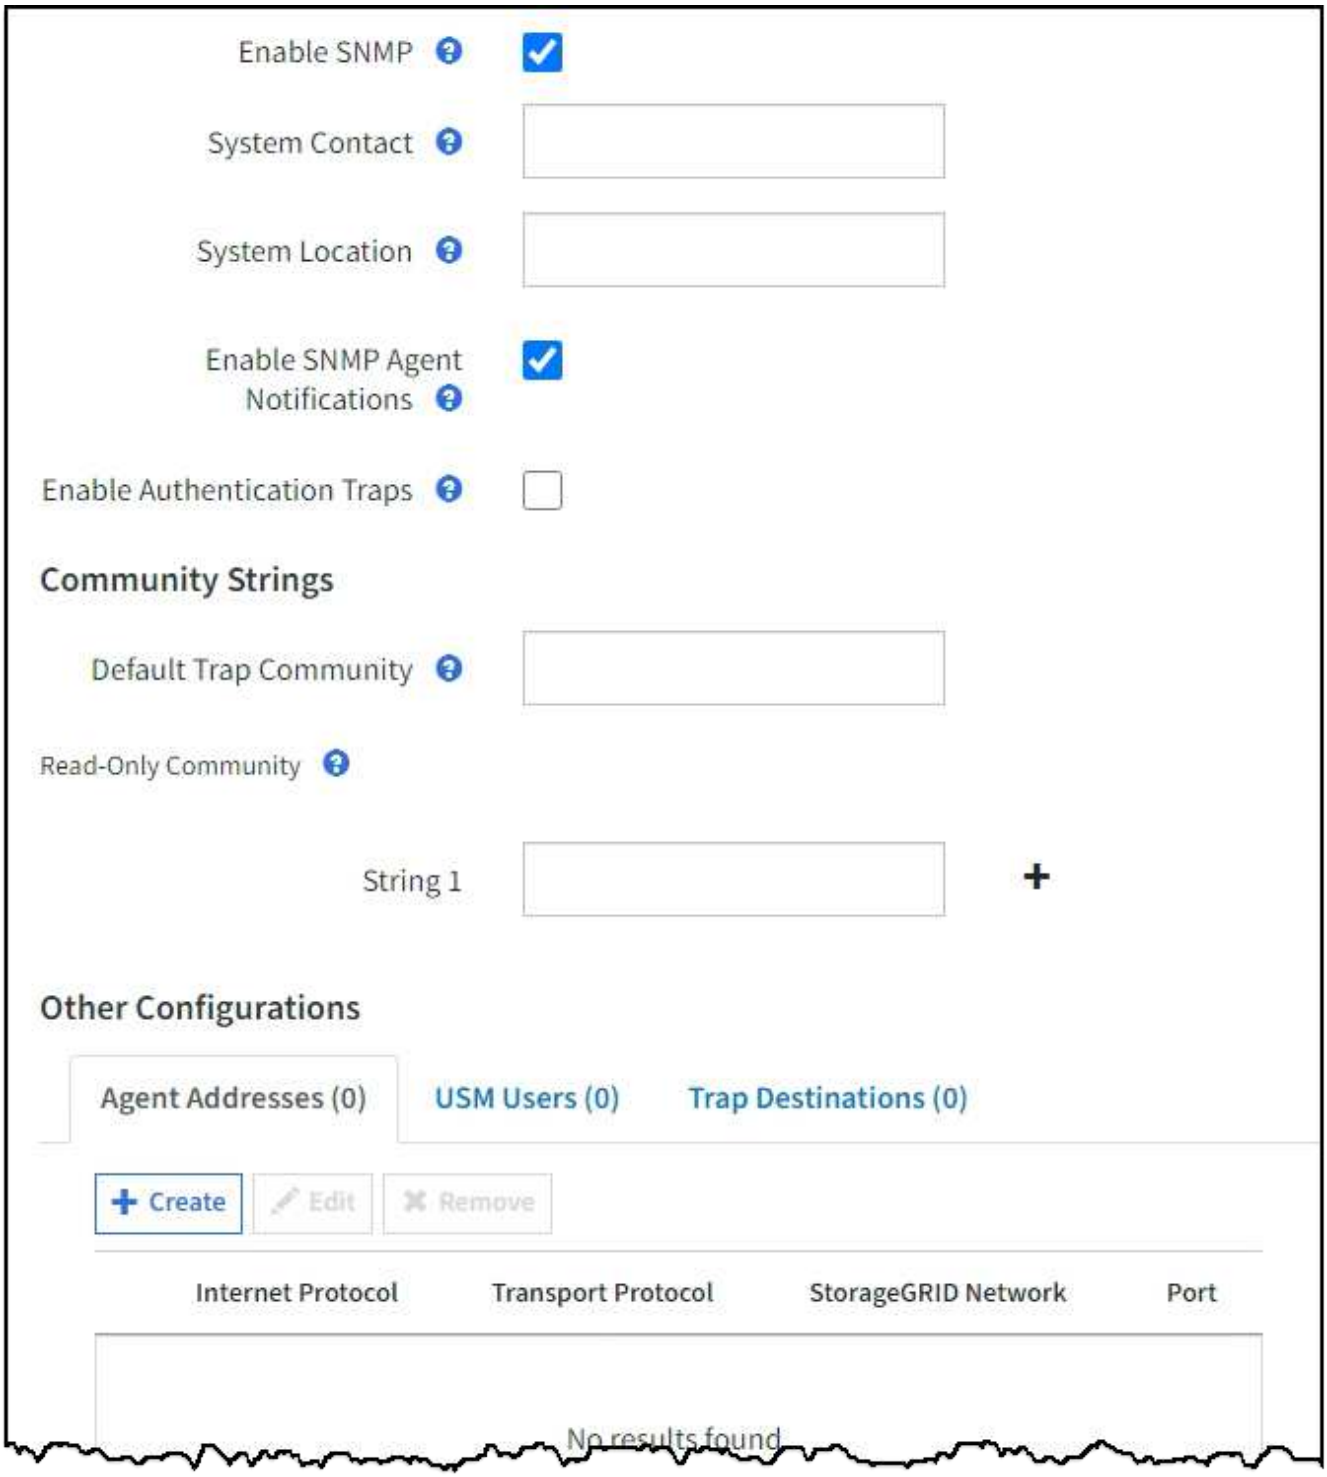


3. In the **System Contact** field, enter the value you want StorageGRID to provide in SNMP messages for sysContact.

   The System Contact typically is an email address. The value you provide applies to all nodes in the StorageGRID system. **System Contact** can be a maximum of 255 characters.

4. In the **System Location** field, enter the value you want StorageGRID to provide in SNMP messages for 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. The value you provide applies to all

nodes in the StorageGRID system. **System Location** can be a maximum of 255 characters.

5. Keep the **Enable SNMP Agent Notifications** checkbox selected if you want the StorageGRID SNMP agent to send trap and inform notifications.

   If this checkbox is cleared, the SNMP agent supports read-only MIB access, but it does not send any SNMP notifications.

6. Select the **Enable Authentication Traps** checkbox if you want the StorageGRID SNMP agent to send an authentication trap if it receives an improperly authenticated protocol message.
7. If you use SNMPv1 or SNMPv2c, complete the Community Strings section.

   The fields in this section are used for community-based authentication in SNMPv1 or SNMPv2c. These fields don't apply to SNMPv3.

   a. In the **Default Trap Community** field, optionally enter the default community string you want to use for 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.

   b. For **Read-Only Community**, enter one or more community strings to allow read-only MIB access on IPv4 and IPv6 agent addresses. Select the plus sign + to add multiple strings.

      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.

      Each community string can be a maximum of 32 characters and can't contain whitespace characters. Up to five strings are allowed.

      > To ensure the security of your StorageGRID system, don't use "public" as the community string. If you don't enter a community string, the SNMP agent uses the grid ID of your StorageGRID system as the community string.

8. Optionally, 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. Each agent address includes an internet protocol, a transport protocol, a StorageGRID network, and optionally a port.

   If you don't configure an agent address, the default listening address is UDP port 161 on all StorageGRID networks.

   a. Select **Create**.

      The Create Agent Address dialog box appears.

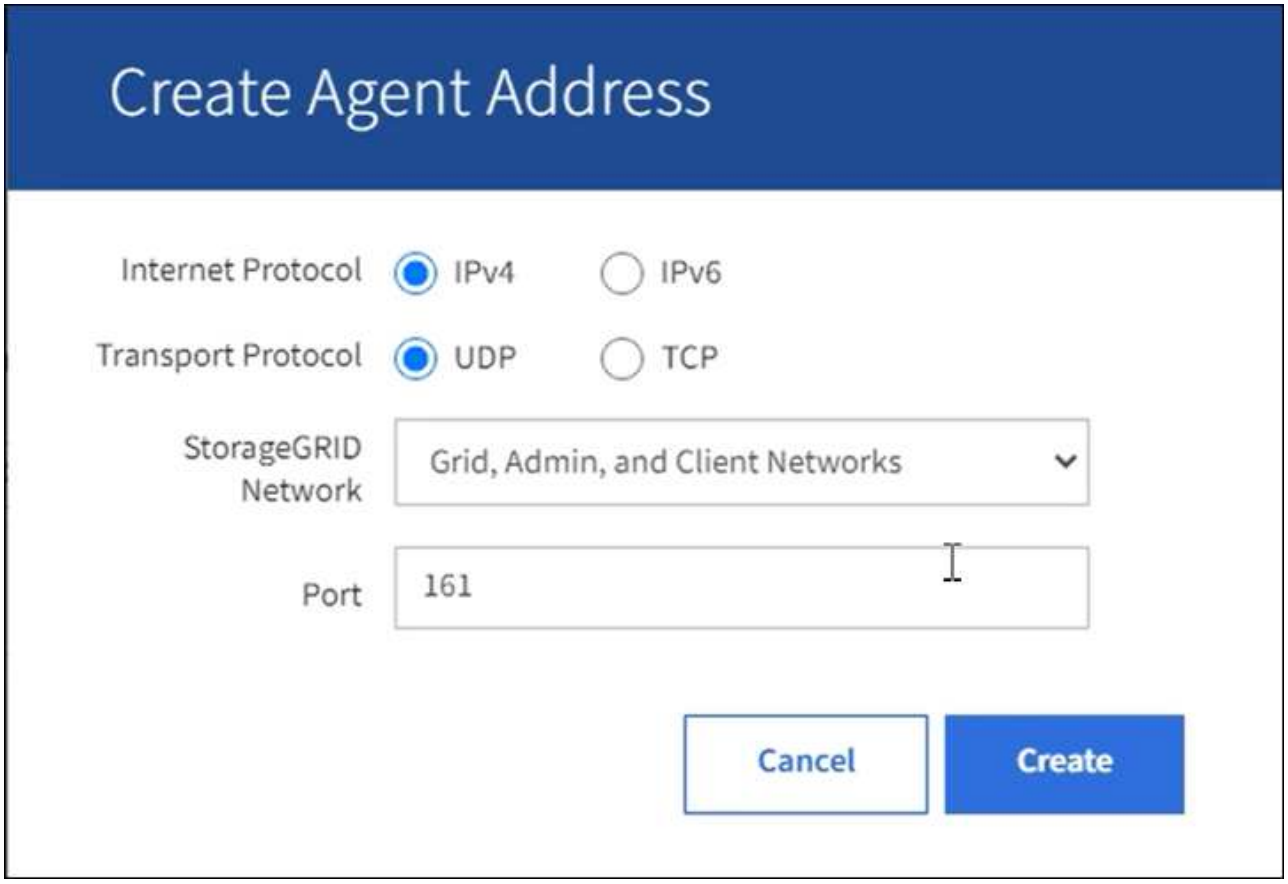


b. For **Internet Protocol**, select whether this address will use IPv4 or IPv6.

   By default, SNMP uses IPv4.

c. For **Transport Protocol**, select whether this address will use UDP or TCP.

   By default, SNMP uses UDP.

d. In the **StorageGRID Network** field, select which StorageGRID network the query will be received on.

   - Grid, Admin, and Client Networks: StorageGRID should listen for SNMP queries on all three networks.
   - Grid Network
   - Admin Network
   - Client Network

   > To ensure that client communications with StorageGRID remain secure, you should not create an agent address for the Client Network.

e. In the **Port** field, optionally enter 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.

   > 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.

f. Select **Create**.

The agent address is created and added to the table.

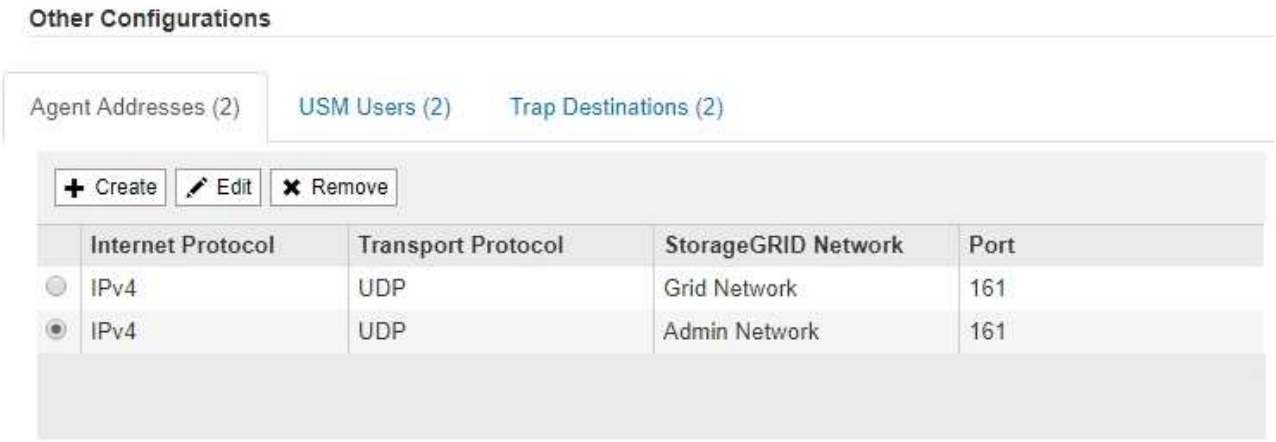


9. If you are using SNMPv3, select the USM Users tab in the Other Configurations section.

Use this tab to define the USM users who are authorized to query the MIB or to receive traps and informs.

> This step does not apply if you are only using SNMPv1 or SNMPv2c.

a. Select **Create**.

The Create USM User dialog box appears.

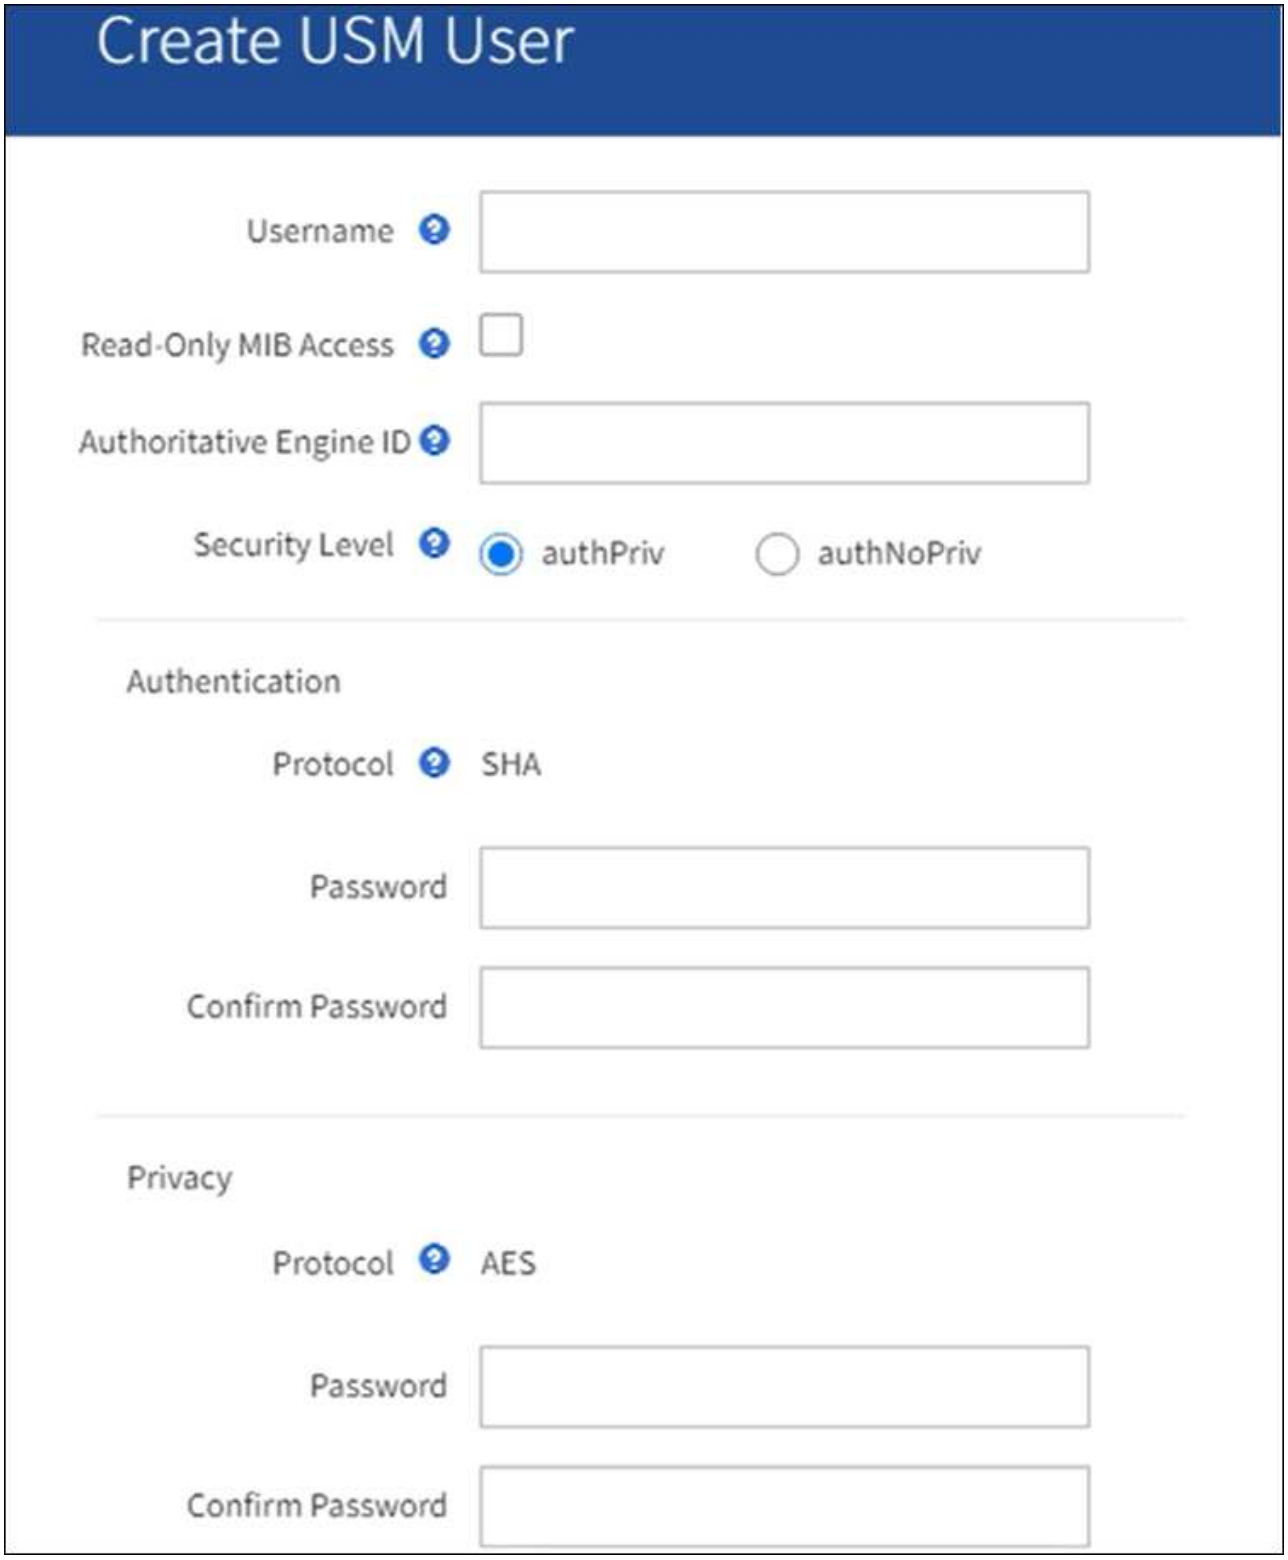


b. Enter a unique **Username** for this USM user.

   Usernames have a maximum of 32 characters and can't contain whitespace characters. The username can't be changed after the user is created.

c. Select the **Read-Only MIB Access** checkbox if this user should have read-only access to the MIB.

   If you select **Read-Only MIB Access**, the **Authoritative Engine ID** field is disabled.

> USM users who have read-only MIB access can't have engine IDs.

d. If this user will be used in an inform destination, enter the **Authoritative Engine ID** for this user.

> SNMPv3 inform destinations must have users with engine IDs. SNMPv3 trap destination can't have users with engine IDs.

The authoritative engine ID can be from 5 to 32 bytes in hexadecimal.

e. Select a 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.

f. Enter and confirm the password this user will use for authentication.

> The only authentication protocol supported is SHA (HMAC-SHA-96).

g. If you selected **authPriv**, enter and confirm the password this user will use for privacy.

> The only privacy protocol supported is AES.

h. Select **Create**.

The USM user is created and added to the table.

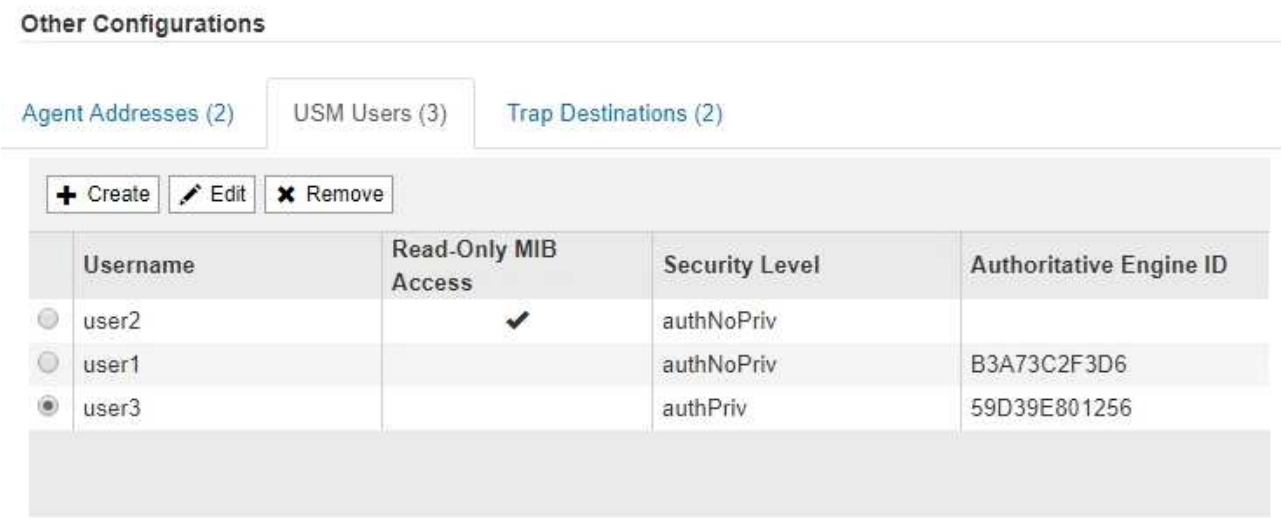


10. In the Other Configurations section, select the Trap Destinations tab.

The Trap Destinations tab allows you to define one or more destinations for StorageGRID trap or inform notifications. When you enable the SNMP agent and select **Save**, StorageGRID starts sending notifications to each defined destination. Notifications are sent when alerts are triggered. Standard notifications are also sent for the supported MIB-II entities (for example, ifDown and coldStart).

a. Select **Create**.

The Create Trap Destination dialog box appears.

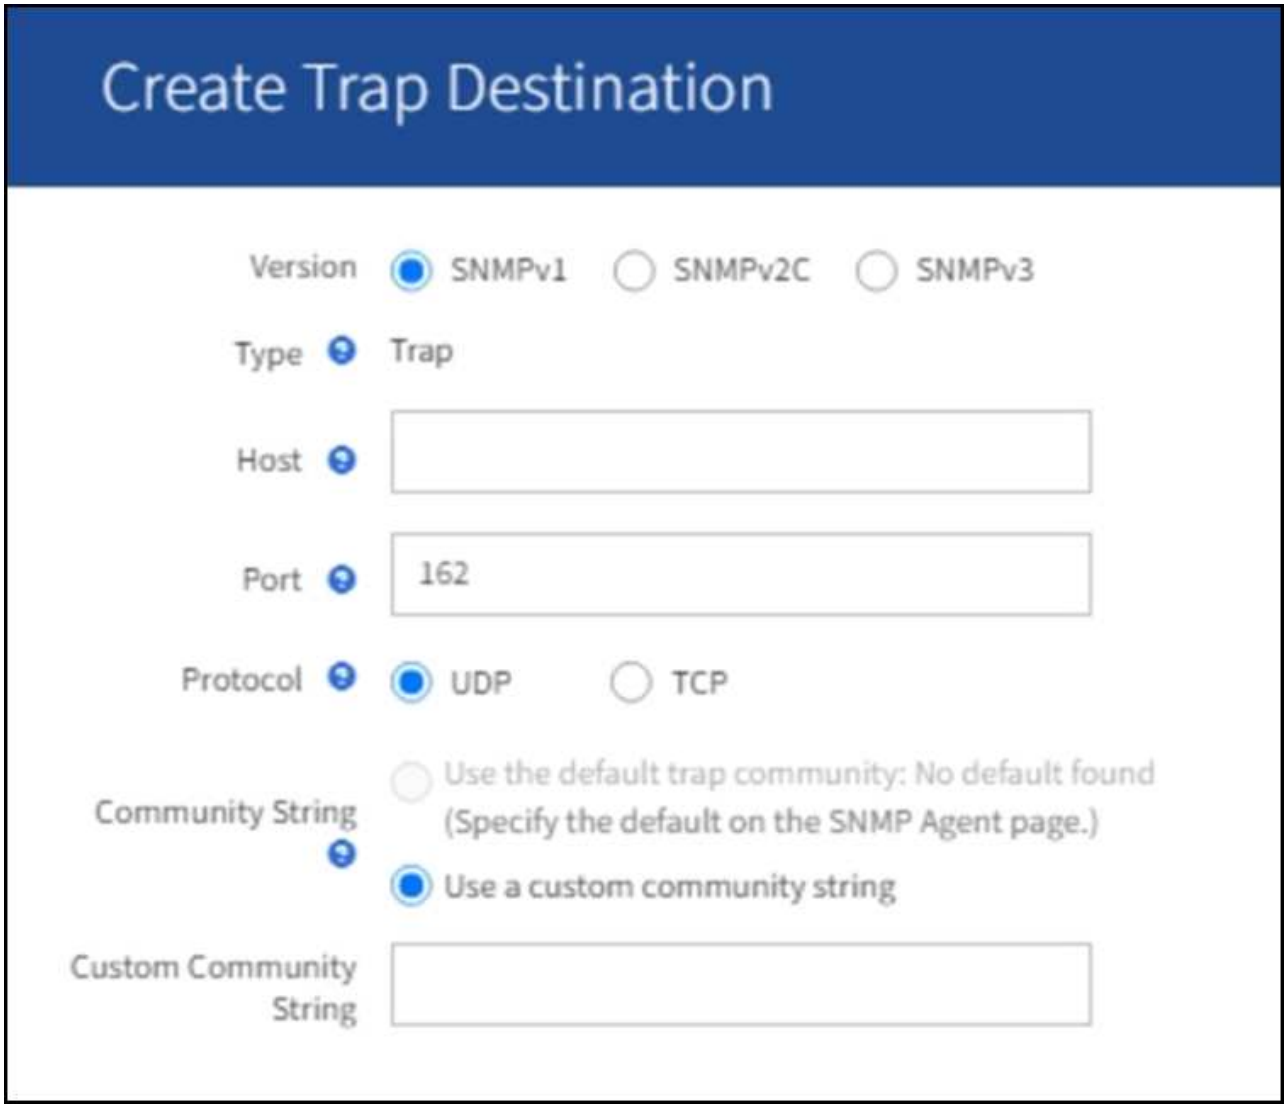


b. In the **Version** field, select which SNMP version will be used for this notification.

c. Complete the form, based on which version you selected

| Version | Specify this information |
|---|---|
| SNMPv1<br><br>(For SNMPv1, the SNMP agent can only send traps. Informs aren't supported.) | 1. In the **Host** field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap.<br>2. For **Port**, use the default (162), unless you must use another value. (162 is the standard port for SNMP traps.)<br>3. For **Protocol**, use the default (UDP). TCP is also supported. (UDP is the standard SNMP trap protocol.)<br>4. Use the default trap community, if one was specified on the SNMP Agent page, or enter a custom community string for this trap destination.<br><br>The custom community string can be a maximum of 32 characters and can't contain whitespace. |

| Version | Specify this information |
| --- | --- |
| SNMPv2c | 1. Select whether the destination will be used for traps or informs.<br>2. In the **Host** field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap.<br>3. For **Port**, use the default (162), unless you must use another value. (162 is the standard port for SNMP traps.)<br>4. For **Protocol**, use the default (UDP). TCP is also supported. (UDP is the standard SNMP trap protocol.)<br>5. Use the default trap community, if one was specified on the SNMP Agent page, or enter a custom community string for this trap destination.<br><br>The custom community string can be a maximum of 32 characters and can't contain whitespace. |
| SNMPv3 | 1. Select whether the destination will be used for traps or informs.<br>2. In the **Host** field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap.<br>3. For **Port**, use the default (162), unless you must use another value. (162 is the standard port for SNMP traps.)<br>4. For **Protocol**, use the default (UDP). TCP is also supported. (UDP is the standard SNMP trap protocol.)<br>5. Select the USM user that will be used for authentication.<br>◦ If you selected **Trap**, only USM users without authoritative engine IDs are shown.<br>◦ If you selected **Inform**, only USM users with authoritative engine IDs are shown. |

d. Select **Create**.

   The trap destination is created and added to the table.

11. When you have completed the SNMP agent configuration, select **Save**.

    The new SNMP agent configuration becomes active.

**Related information**

Silence alert notifications

## Update the SNMP agent

You might want to disable SNMP notifications, update community strings, or add or remove agent addresses, USM users, and trap destinations.

**Before you begin**

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

### About this task

Whenever you update the SNMP agent configuration, be aware that you must select **Save** at the bottom on the SNMP Agent page to commit any changes you have made on each tab.

### Steps

1. Select **CONFIGURATION** > **Monitoring** > **SNMP agent**.

   The SNMP Agent page appears.

2. If you want to disable the SNMP agent on all grid nodes, clear the **Enable SNMP** checkbox, and select **Save**.

   The SNMP agent is disabled for all grid nodes. If you later re-enable the agent, any previous SNMP configuration settings are retained.

3. Optionally, update the values you entered for **System Contact** and **System Location**.
4. Optionally, clear the **Enable SNMP Agent Notifications** checkbox if you no longer want the StorageGRID SNMP agent to send trap and inform notifications.

   When this checkbox is cleared, the SNMP agent supports read-only MIB access, but it does not send any SNMP notifications.

5. Optionally, clear the **Enable Authentication Traps** checkbox if you no longer want the StorageGRID SNMP agent to send an authentication trap when it receives an improperly authenticated protocol message.
6. If you use SNMPv1 or SNMPv2c, optionally update the Community Strings section.

   The fields in this section are used for community-based authentication in SNMPv1 or SNMPv2c. These fields don't apply to SNMPv3.

   > If you want to remove the default community string, you must first ensure that all trap destinations use a custom community string.

7. If you want to update agent addresses, select the Agent Addresses tab in the Other Configurations section.

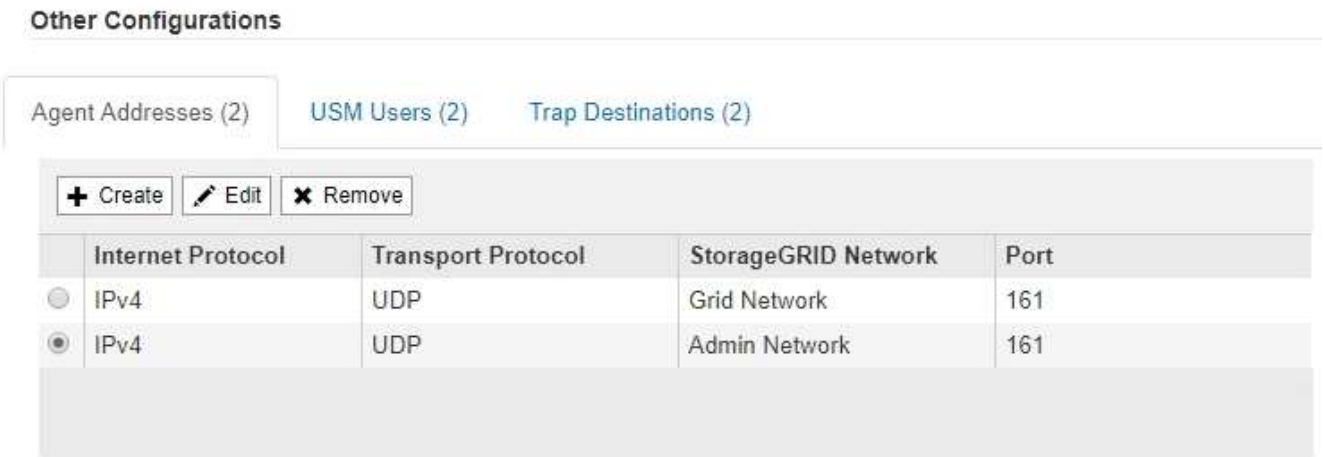

Use this tab to specify one or more “listening addresses.” These are the StorageGRID addresses on which the SNMP agent can receive queries. Each agent address includes an internet protocol, a transport protocol, a StorageGRID network, and a port.

a. To add an agent address, select **Create**. Then, refer to the step for agent addresses in the instructions for configuring the SNMP agent.
b. To edit an agent address, select the radio button for the address, and select **Edit**. Then, refer to the step for agent addresses in the instructions for configuring the SNMP agent.
c. To remove an agent address, select the radio button for the address, and select **Remove**. Then, select **OK** to confirm that you want to remove this address.
d. To commit your changes, select **Save** at the bottom of the SNMP Agent page.

8. If you want to update USM users, select the USM Users tab in the Other Configurations section.

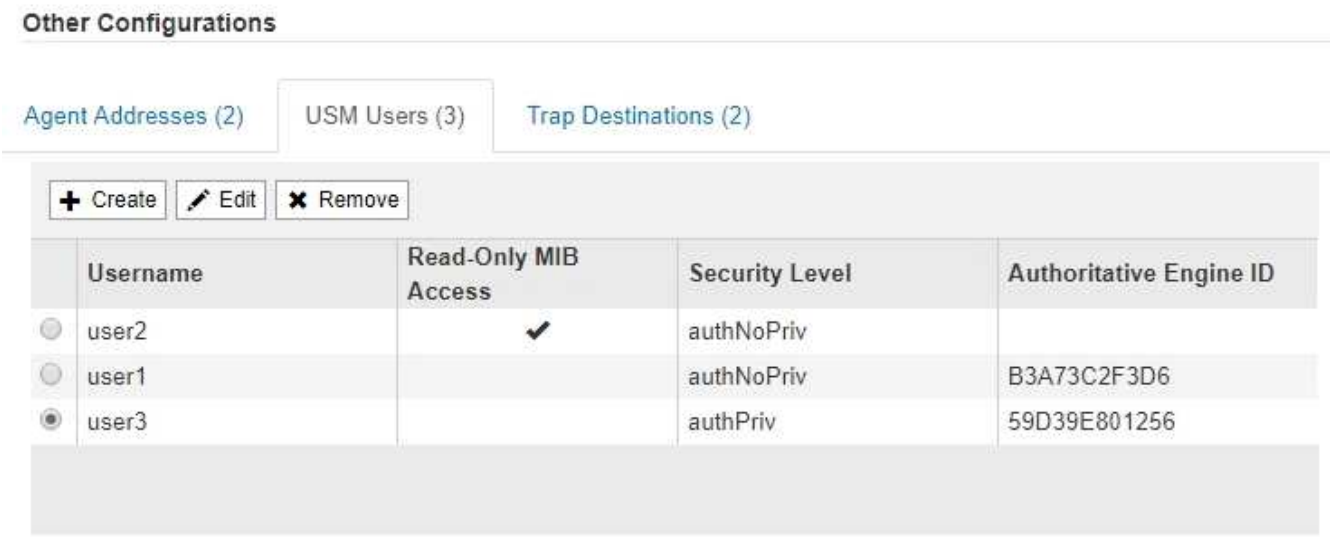


Use this tab to define the USM users who are authorized to query the MIB or to receive traps and informs.

a. To add a USM user, select **Create**. Then, refer to the step for USM users in the instructions for configuring the SNMP agent.
b. To edit a USM user, select the radio button for the user, and select **Edit**. Then, refer to the step for USM users in the instructions for configuring the SNMP agent.

   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, as described in step SNMP trap destination. Otherwise, a validation error occurs when you save the SNMP agent configuration.

c. To remove a USM user, select the radio button for the user, and select **Remove**. Then, select **OK** to confirm that you want to remove this user.

   > If the user you removed is currently selected for a trap destination, you must edit or remove the destination, as described in step SNMP trap destination. Otherwise, a validation error occurs when you save the SNMP agent configuration.

d. To commit your changes, select **Save** at the bottom of the SNMP Agent page.

9. If you want to update trap destinations, select the Trap Destinations tab in the Other Configurations section.

   The Trap Destinations tab allows you to define one or more destinations for StorageGRID trap or inform notifications. When you enable the SNMP agent and select **Save**, StorageGRID starts sending notifications to each defined destination. Notifications are sent when alerts and alarms are triggered. Standard notifications are also sent for the supported MIB-II entities (for example, ifDown and coldStart).

   a. To add a trap destination, select **Create**. Then, refer to the step for trap destinations in the instructions for configuring the SNMP agent.
   b. To edit a trap destination, select the radio button for the user, and select **Edit**. Then, refer to the step for trap destinations in the instructions for configuring the SNMP agent.
   c. To remove a trap destination, select the radio button for the destination, and select **Remove**. Then, select **OK** to confirm that you want to remove this destination.
   d. To commit your changes, select **Save** at the bottom of the SNMP Agent page.

10. 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

#### 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.
3. Optionally, you can access MIB files at the following location on any StorageGRID node: `/usr/share/snmp/mibs`
4. 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. Then 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. |
| `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` | 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.

### 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.

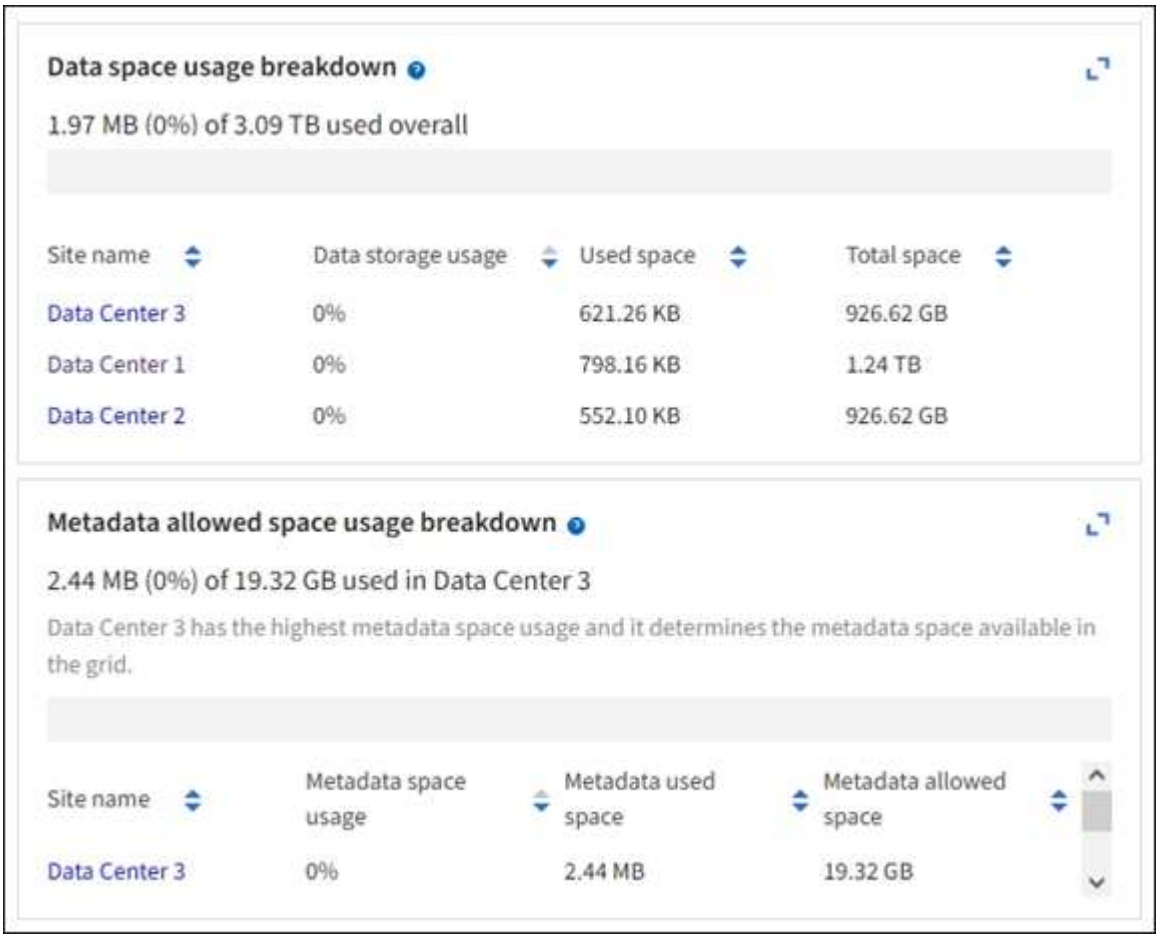


The Storage usage panel on the Tenant Manager dashboard displays the following:

- 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

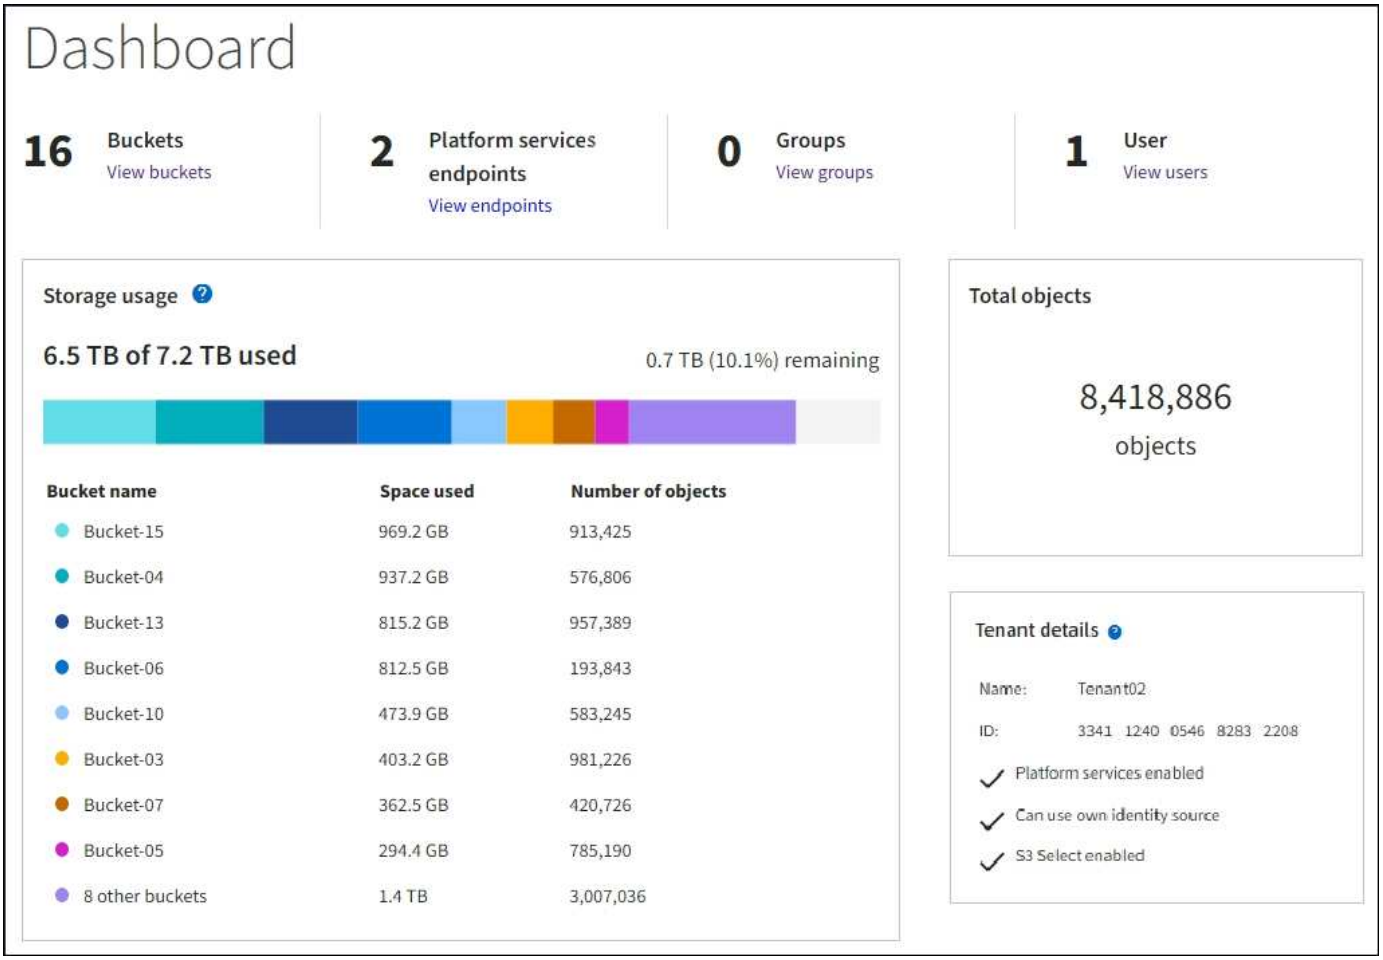


In addition, graphs that show how StorageGRID metrics and attributes change over time are available from the Nodes page and from the **SUPPORT** > **Tools** > **Grid topology** page.

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.

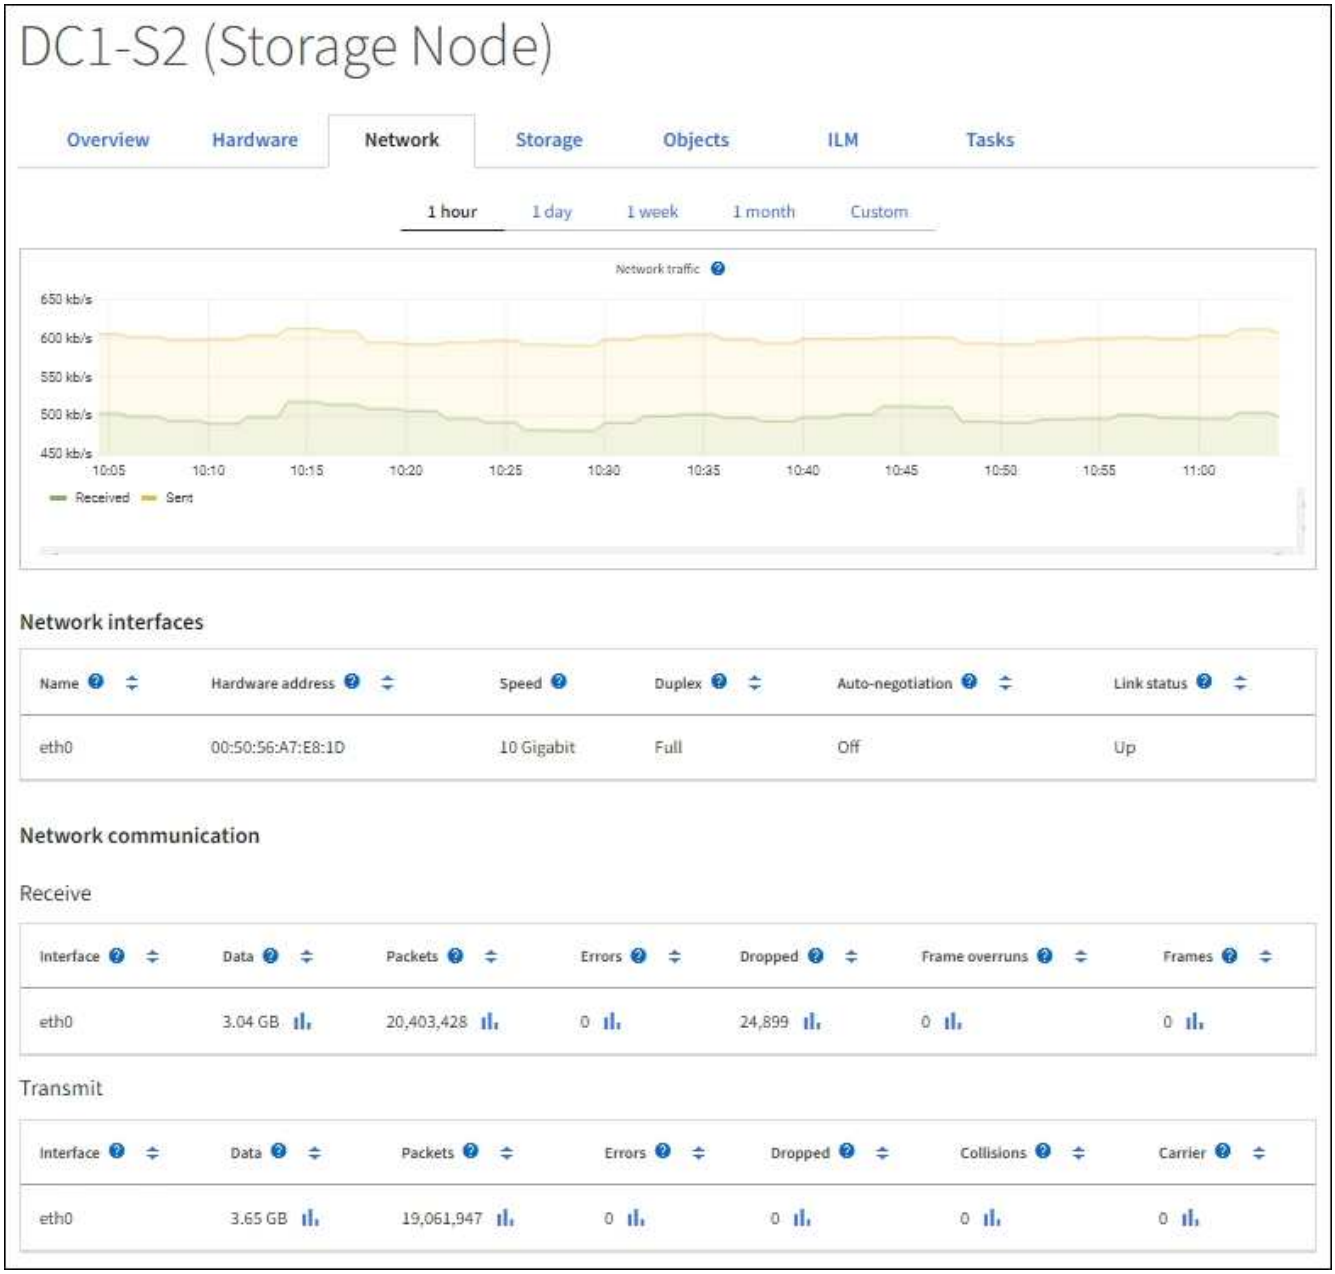


Grafana charts are also included on the pre-constructed dashboards available from the **SUPPORT** > **Tools** > **Metrics** page.

- **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.

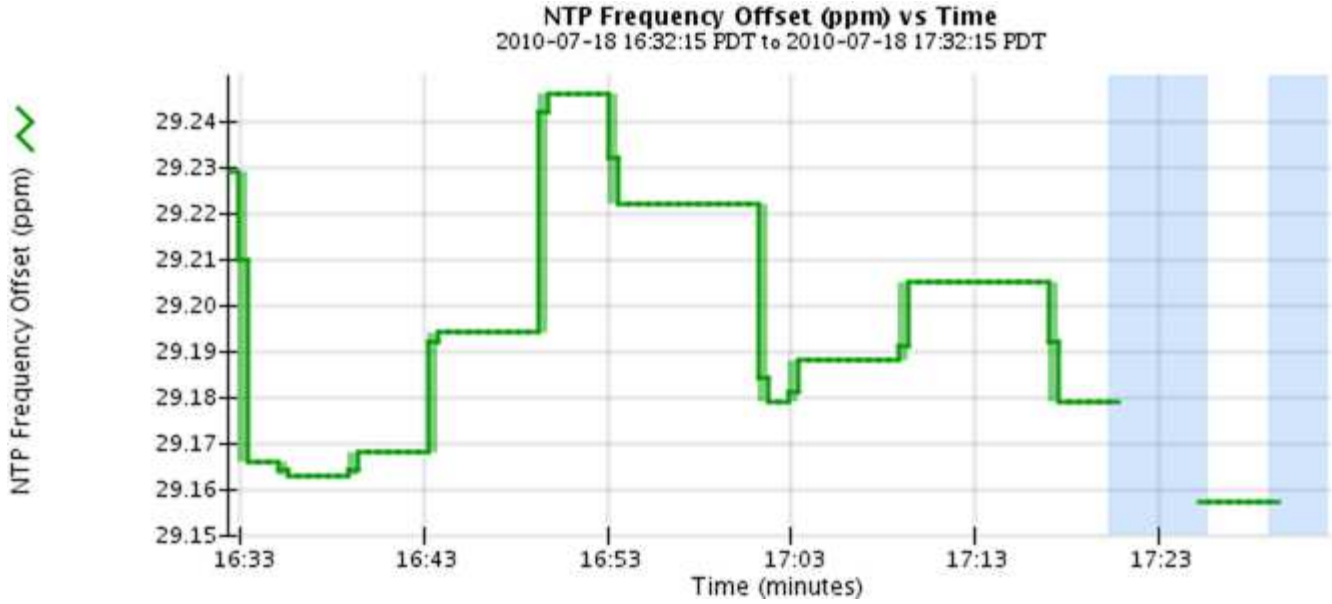


- **Area graphs**: Available from the Nodes page and from the **SUPPORT** > **Tools** > **Grid topology** page (select the chart icon after a data value), area graphs are used to plot volumetric attribute quantities, such as object counts or service load values. Area graphs are similar to line graphs, but include a light brown shading below the line. The changes in the value are plotted in regular data intervals (bins) over time.

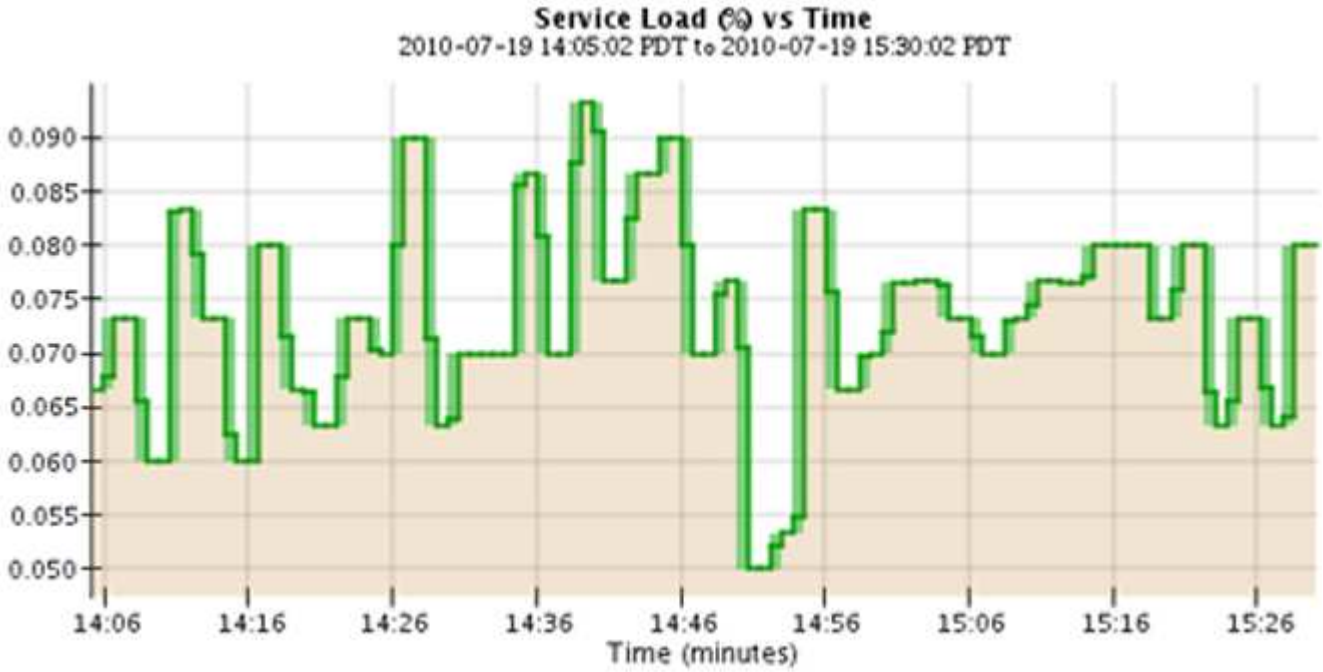


- Some graphs are denoted with a different type of chart icon and have a different format:

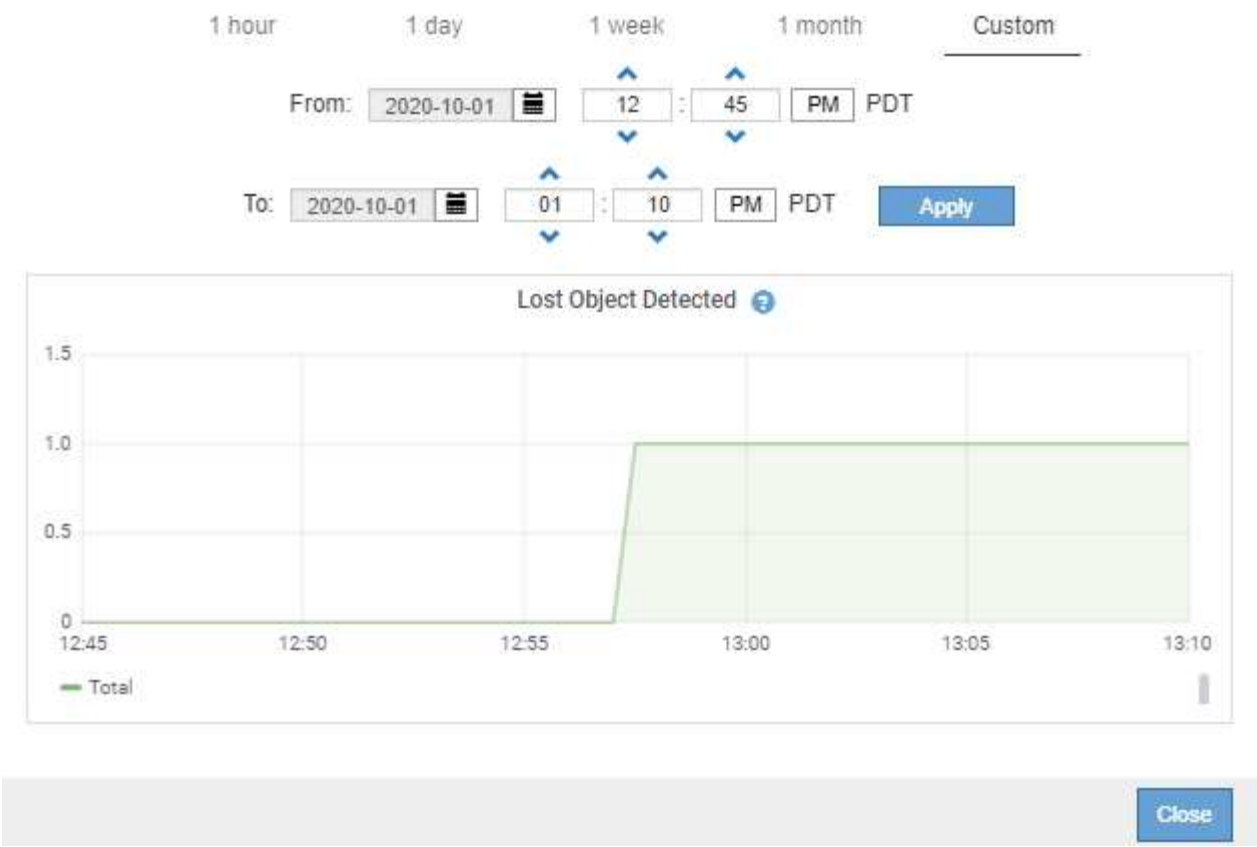


- **State graph**: Available from the **SUPPORT** > **Tools** > **Grid topology** page (select the chart icon after a data value), state graphs are used to plot attribute values that represent distinct states such as a service state that can be online, standby, or offline. State graphs are similar to line graphs, but the transition is discontinuous; that is, the value jumps from one state value to another.

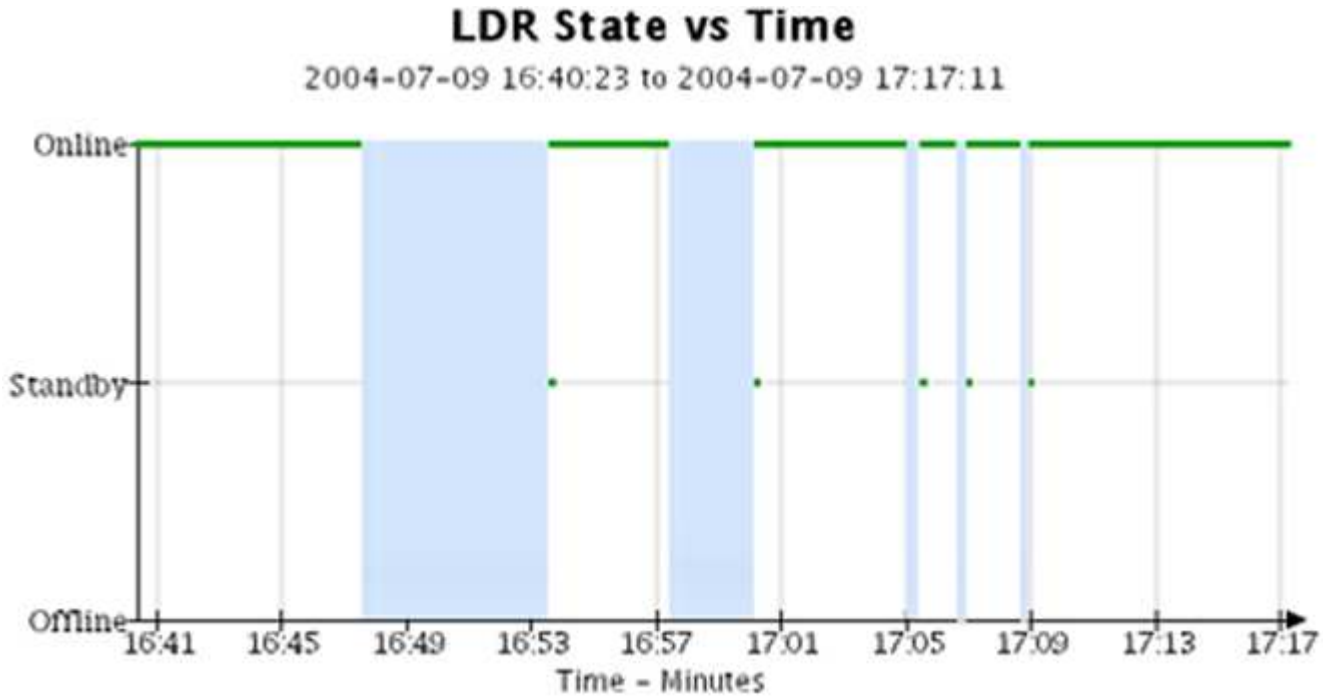


**Related information**

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.

| Sample | 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. |
| 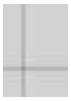 | Gray shading indicates that some or all of the attribute values at that time were not known because the service reporting the attributes was administratively down. |
| | A mixture of gray and blue shading indicates that some of the attribute values at the time were indeterminate (because the service 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.

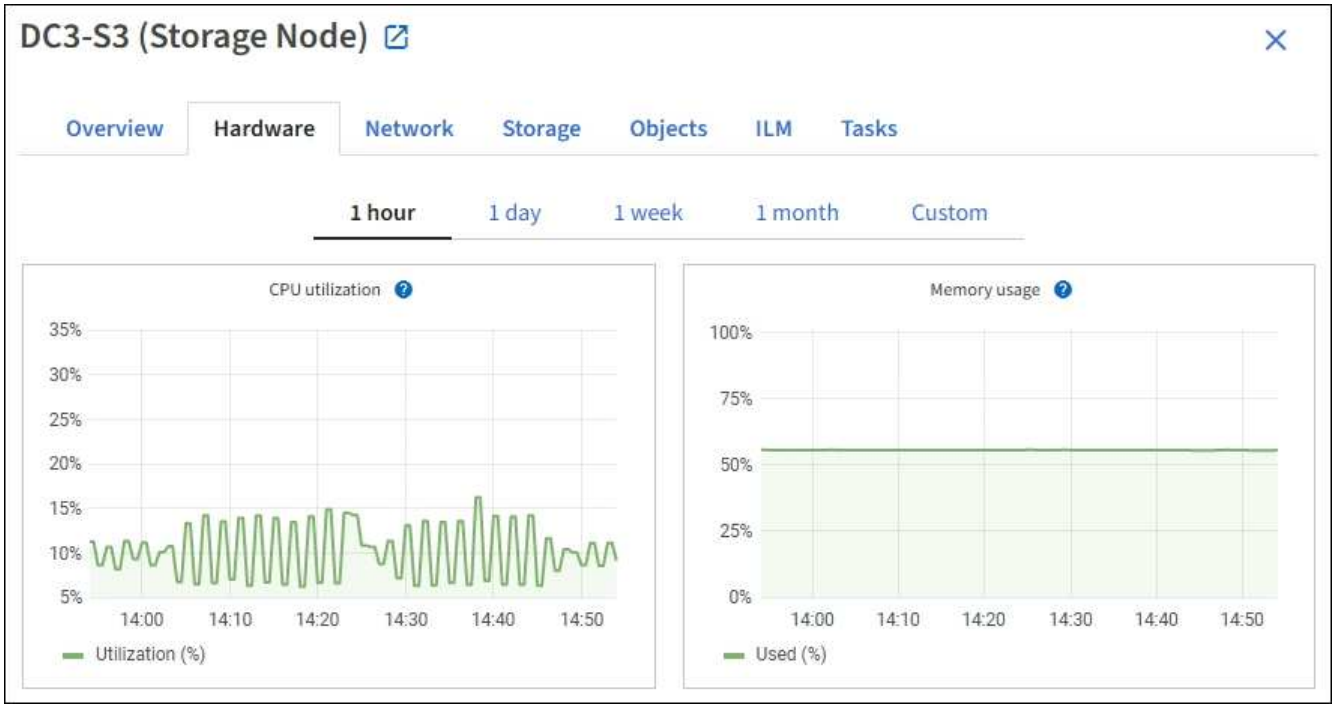


3. Optionally, position your cursor over the chart to see more detailed values for a particular point in time.

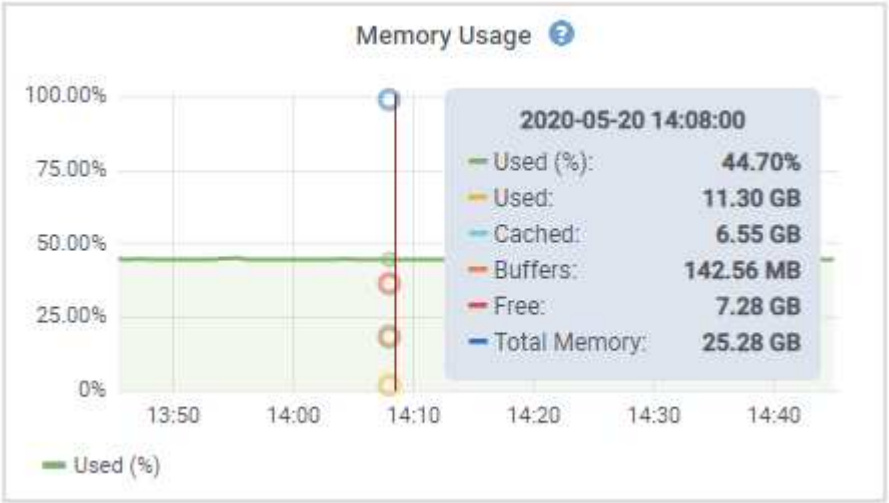


4. As required, you can often display a chart for a specific attribute or metric. From the table on the Nodes page, select the chart icon to the right of the attribute name.

> 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

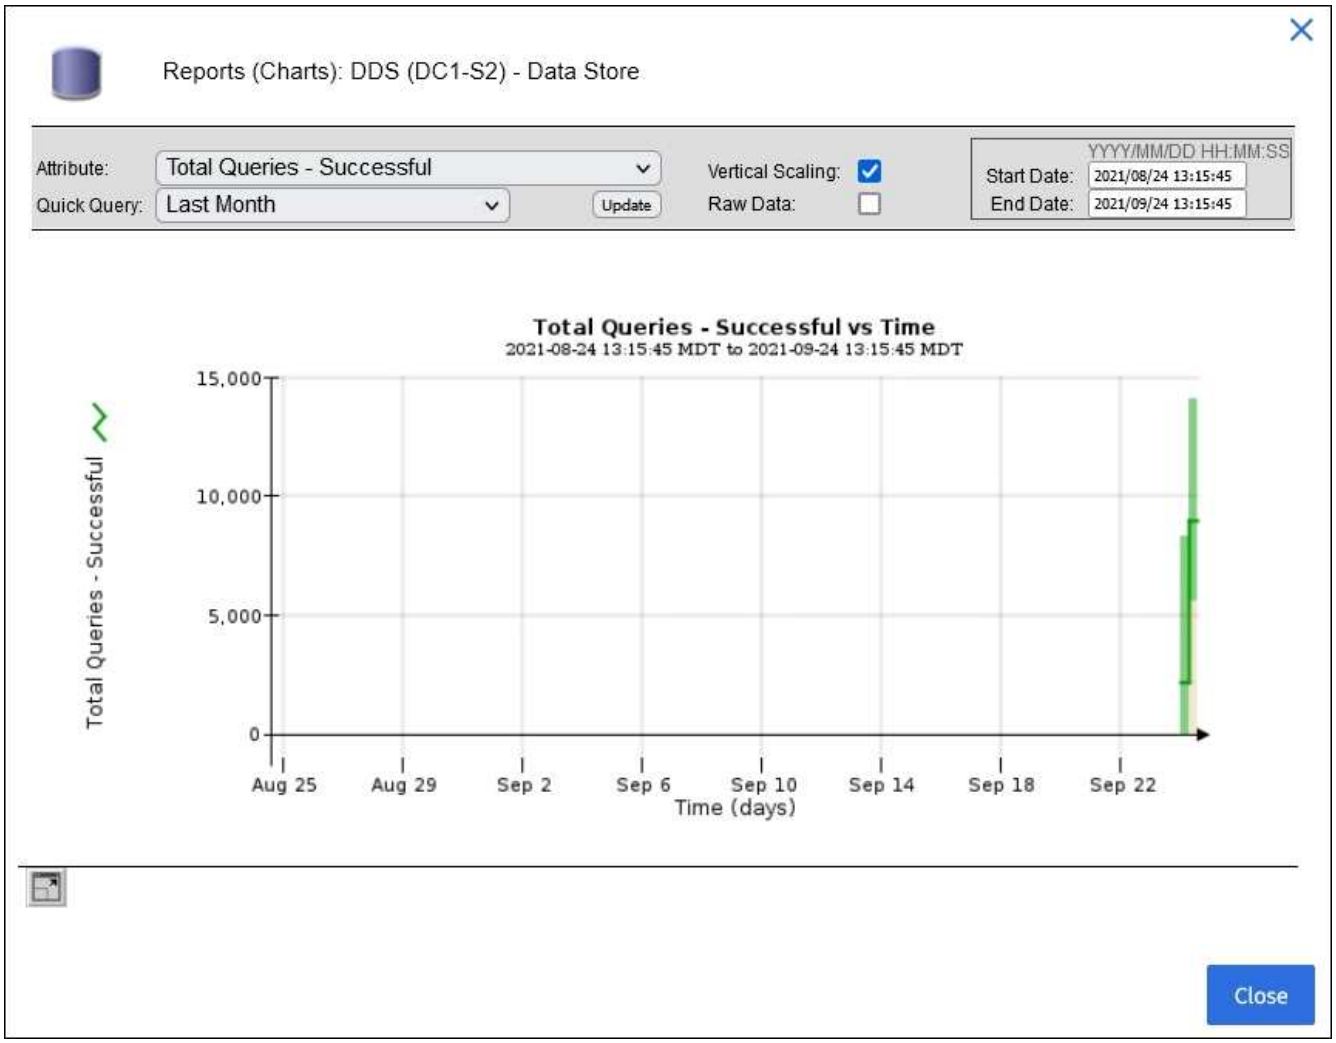


**Example 2**: From the Objects tab for a Storage Node, you can select the chart icon to see the Grafana graph of the count of lost objects detected over time.

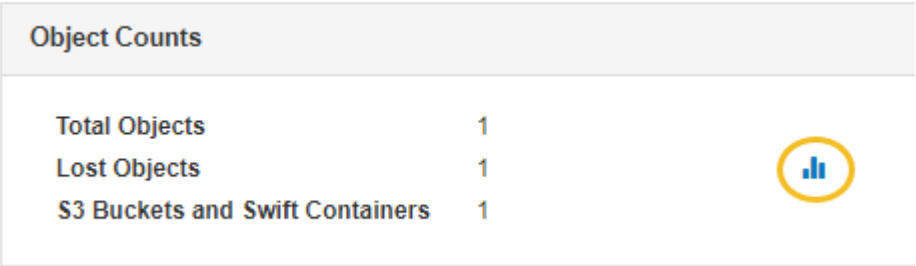

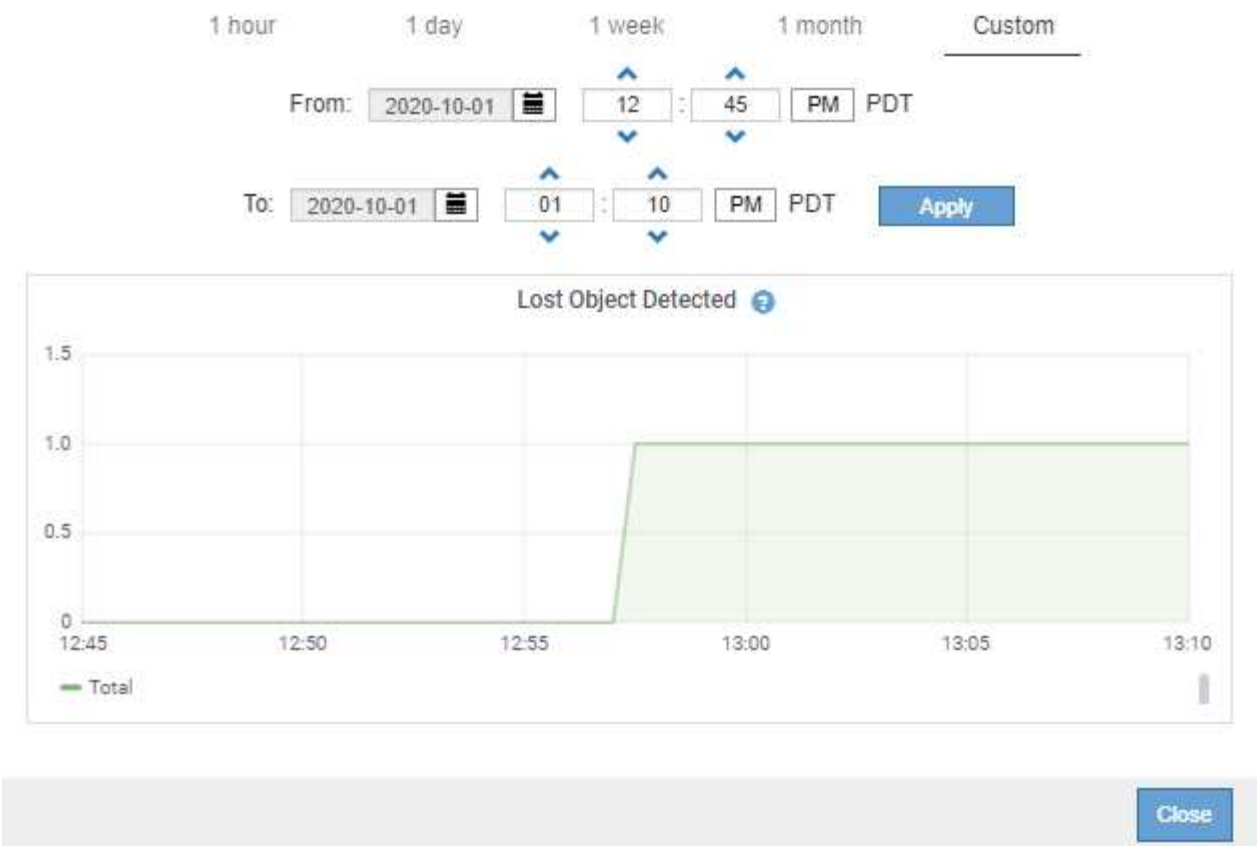


5. To display charts for attributes that aren't shown on the Node page, select **SUPPORT** > **Tools** > **Grid topology**.
6. Select ***grid node*** > ***component or service*** > **Overview** > **Main**.

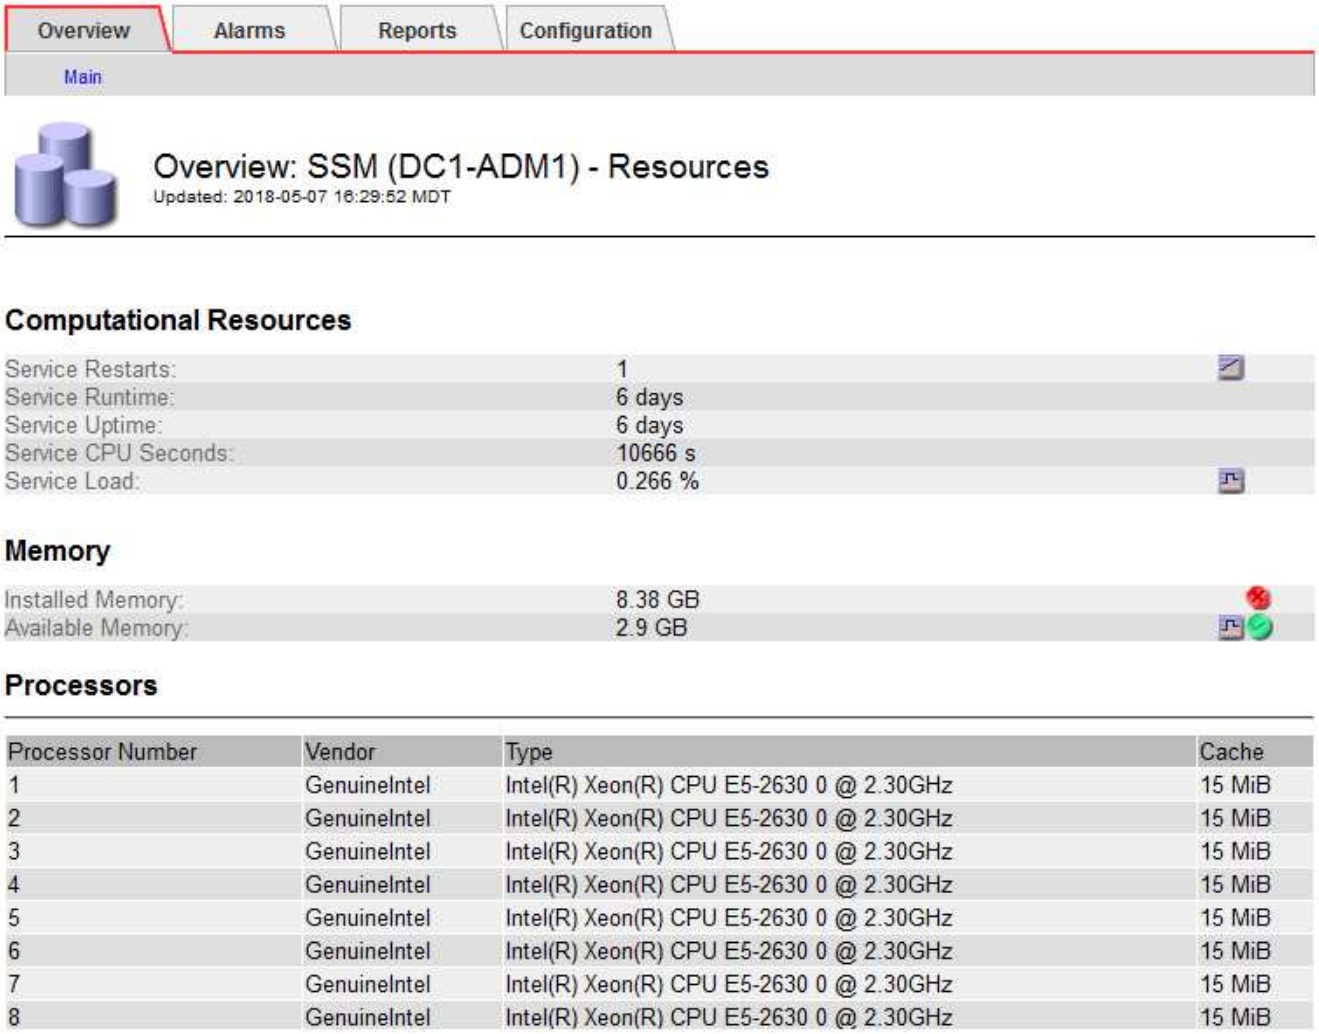


7. Select the chart icon next to the attribute.

   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 must 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 must 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*.

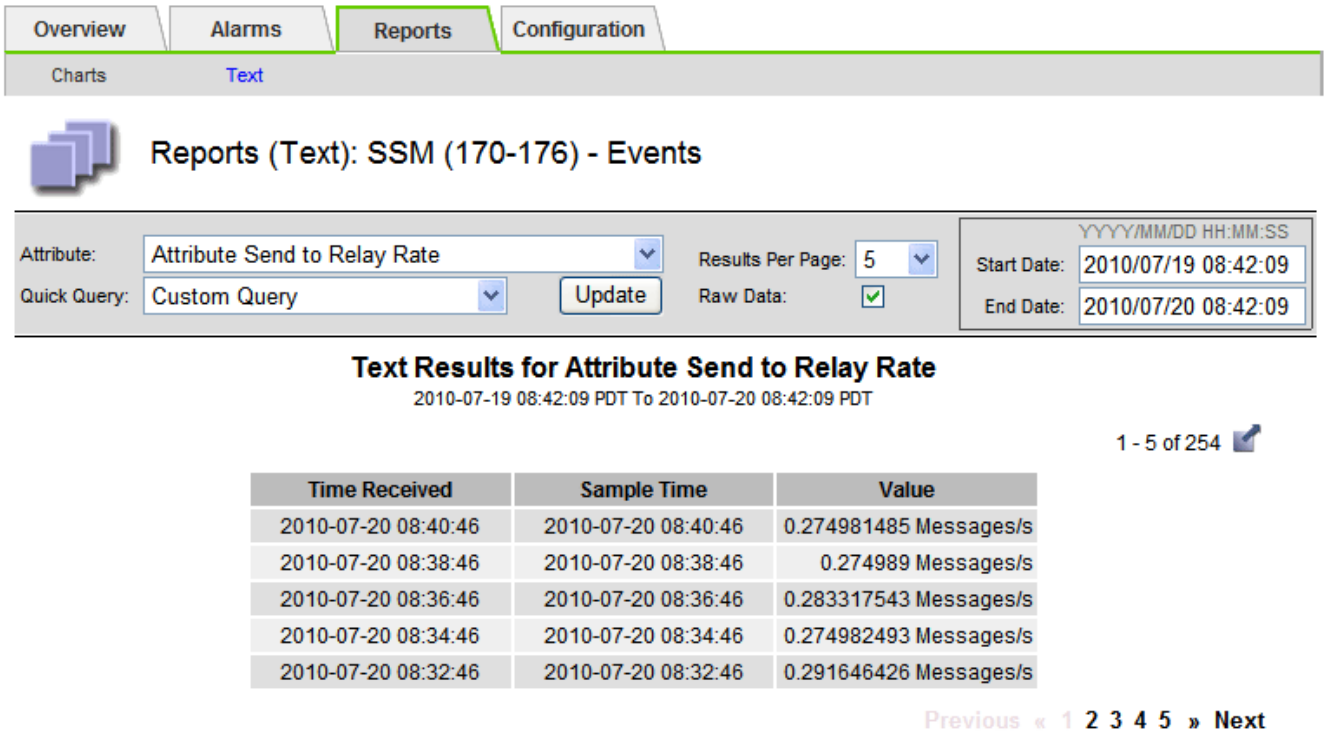


   The Export Text Report window opens displaying the report.

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

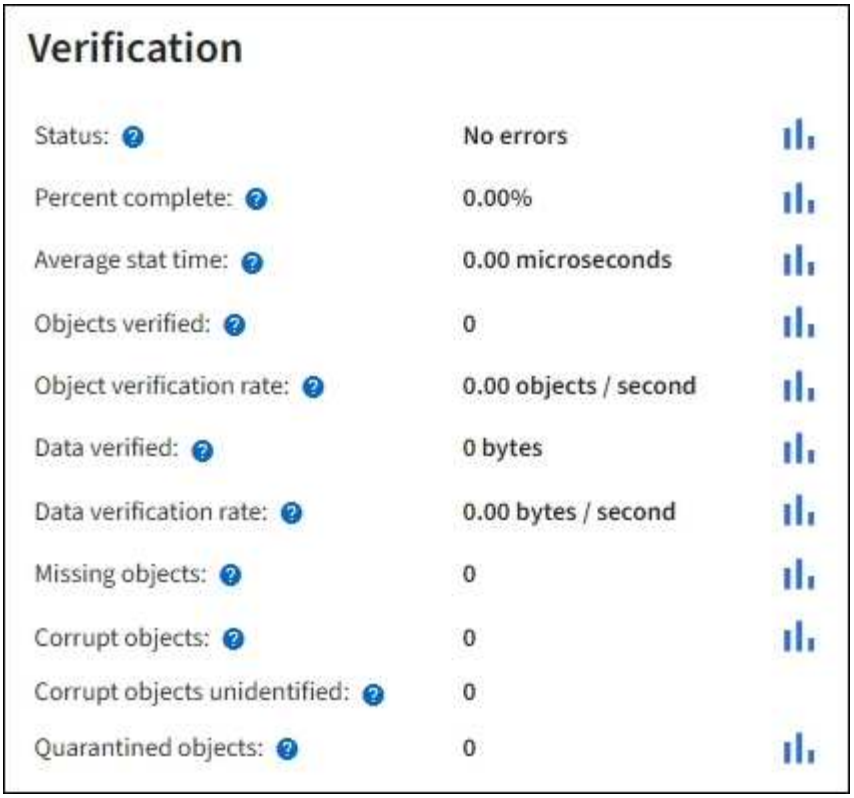


- To check erasure-coded fragment verification, select ***Storage Node*** > **ILM** and look at the attributes in the Erasure coding verification section.

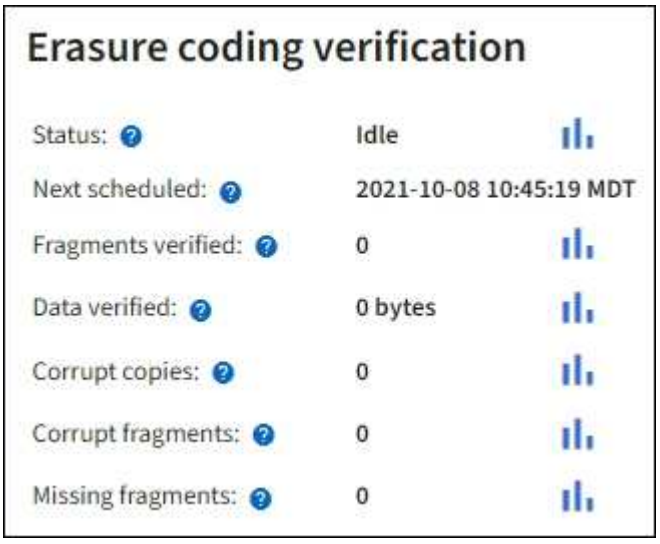


Select the question mark ? next to an attribute's name to display help text.

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.

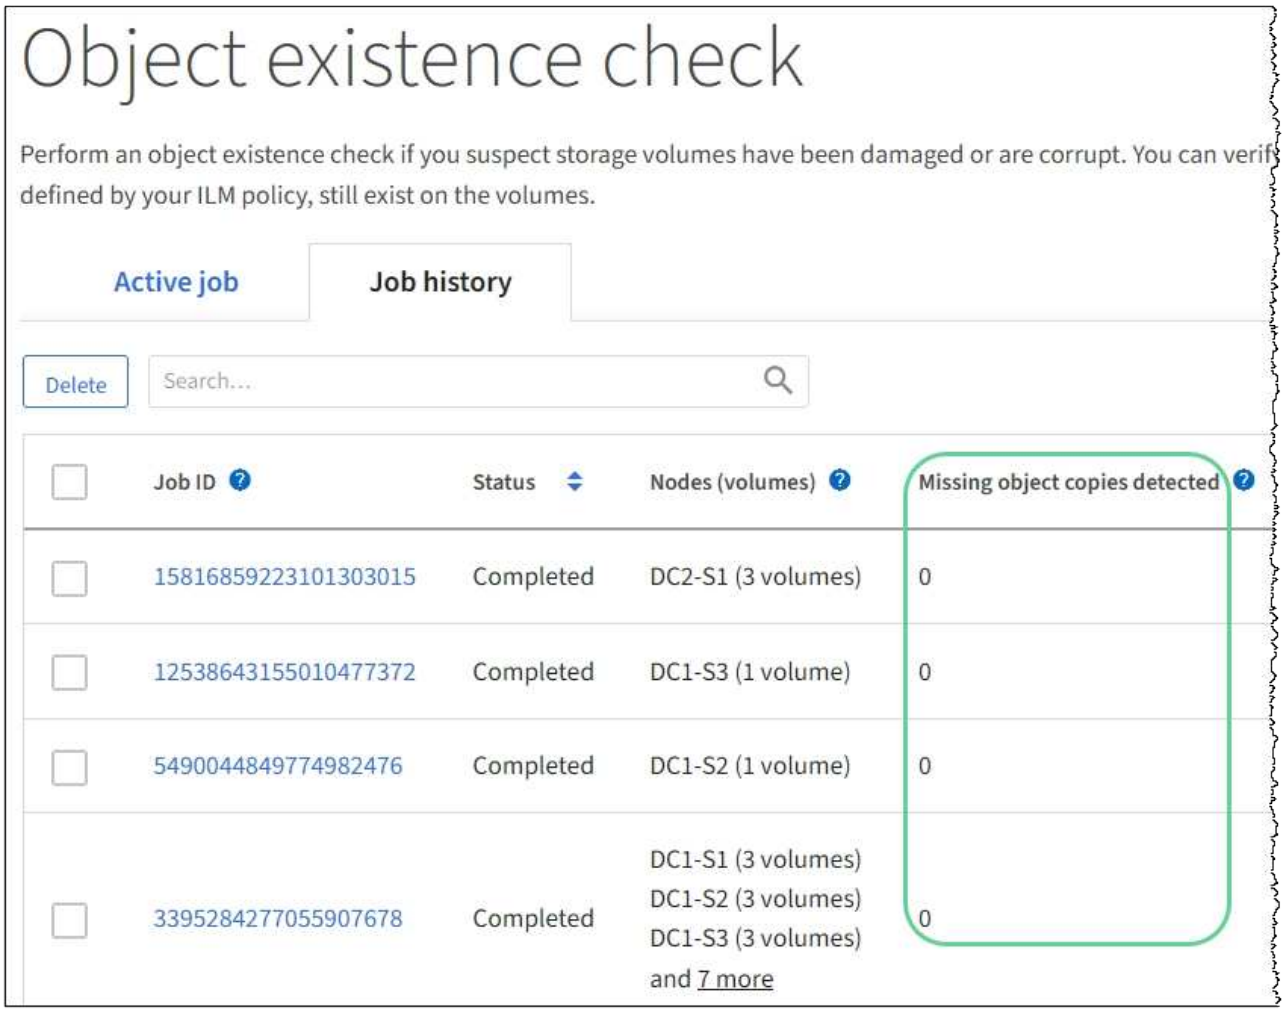


## Monitor events

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**.

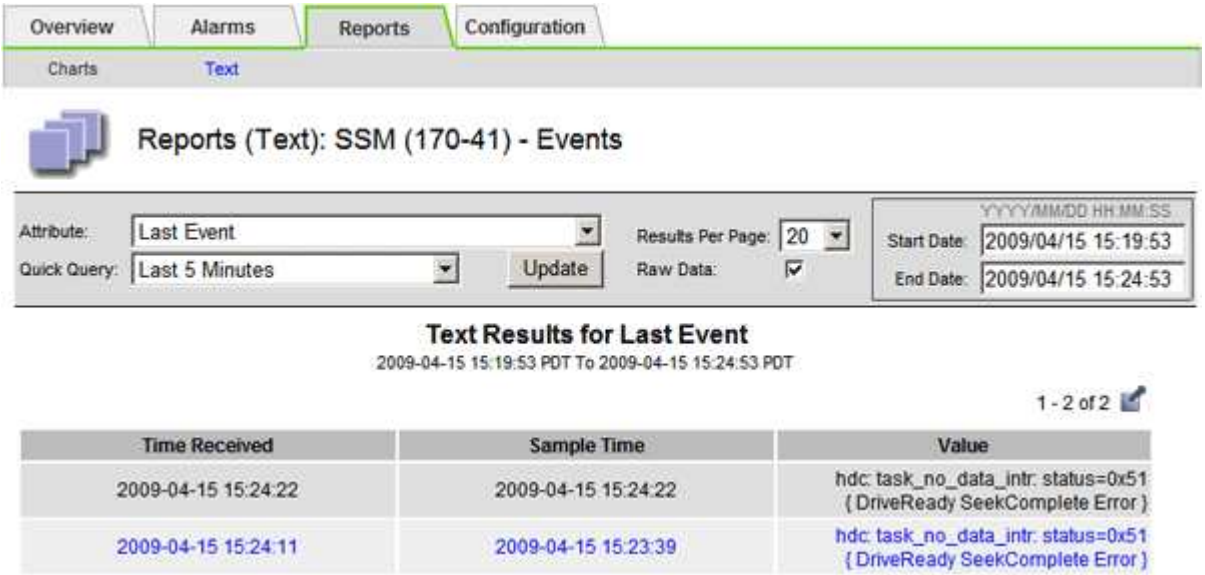


### Create custom syslog events

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

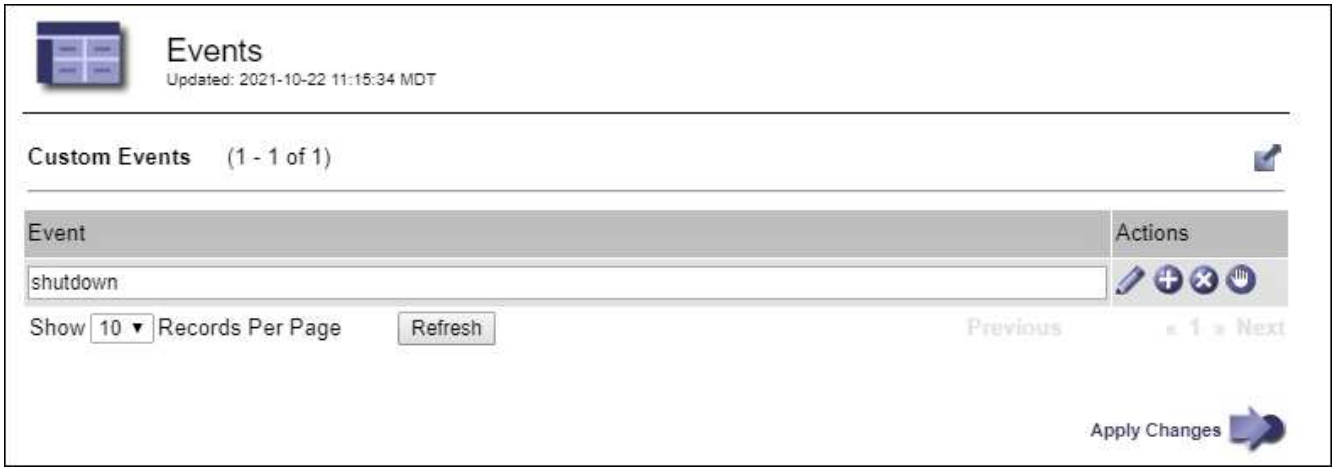


4. Select **Apply Changes**.
5. Select **SUPPORT** > **Tools** > **Grid topology**.
6. Select ***grid node*** > **SSM** > **Events**.
7. Locate the entry for Custom Events in the Events table, and monitor the value for **Count**.

   If the count increases, a custom event you are monitoring is being triggered on that grid node.

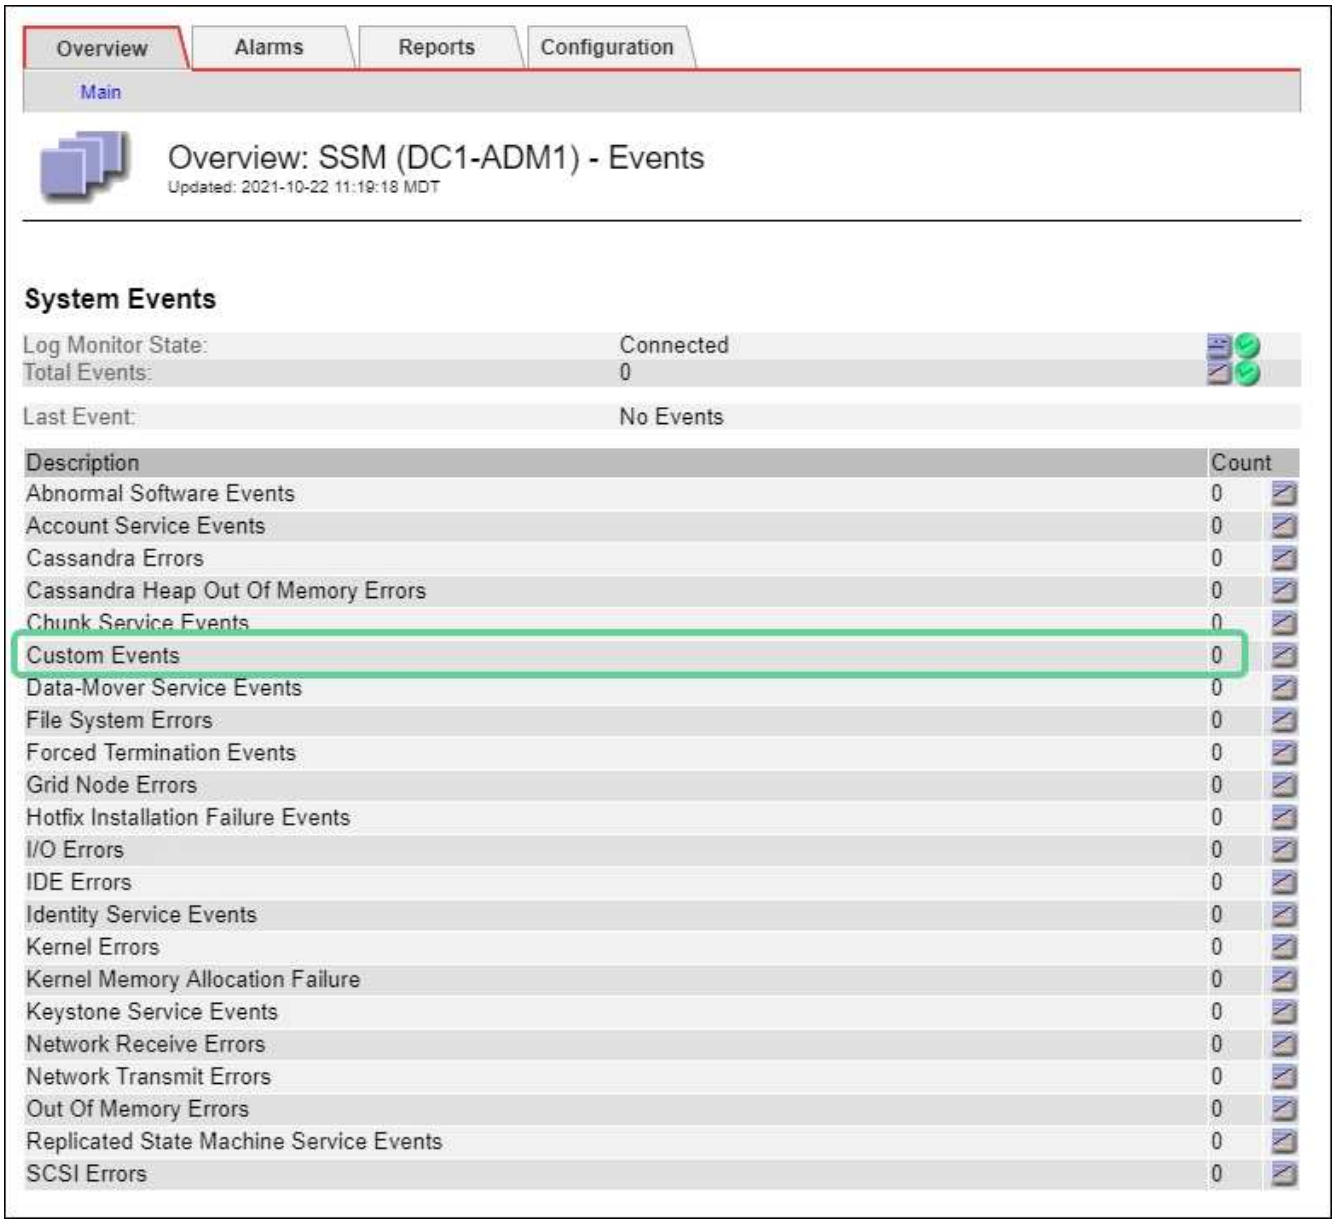


### Reset the count of custom events to zero

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.

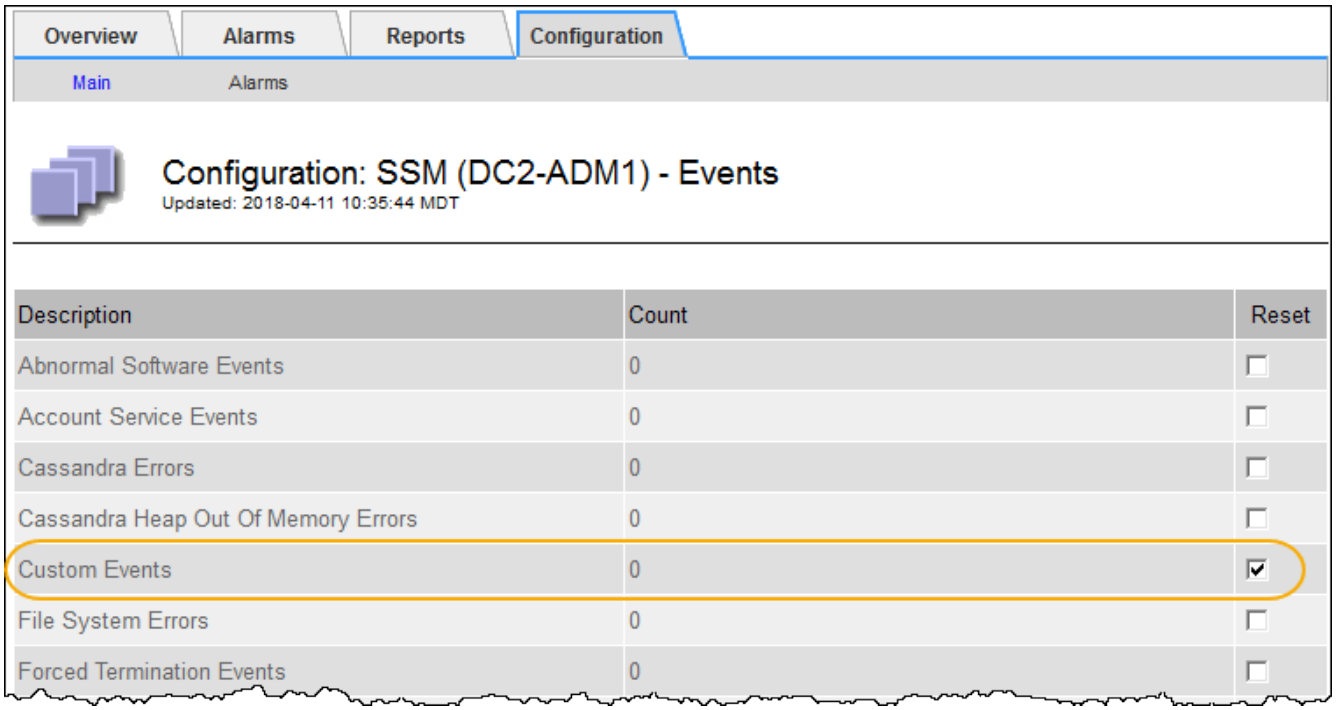


4. Select **Apply Changes**.

## 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 must 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**.

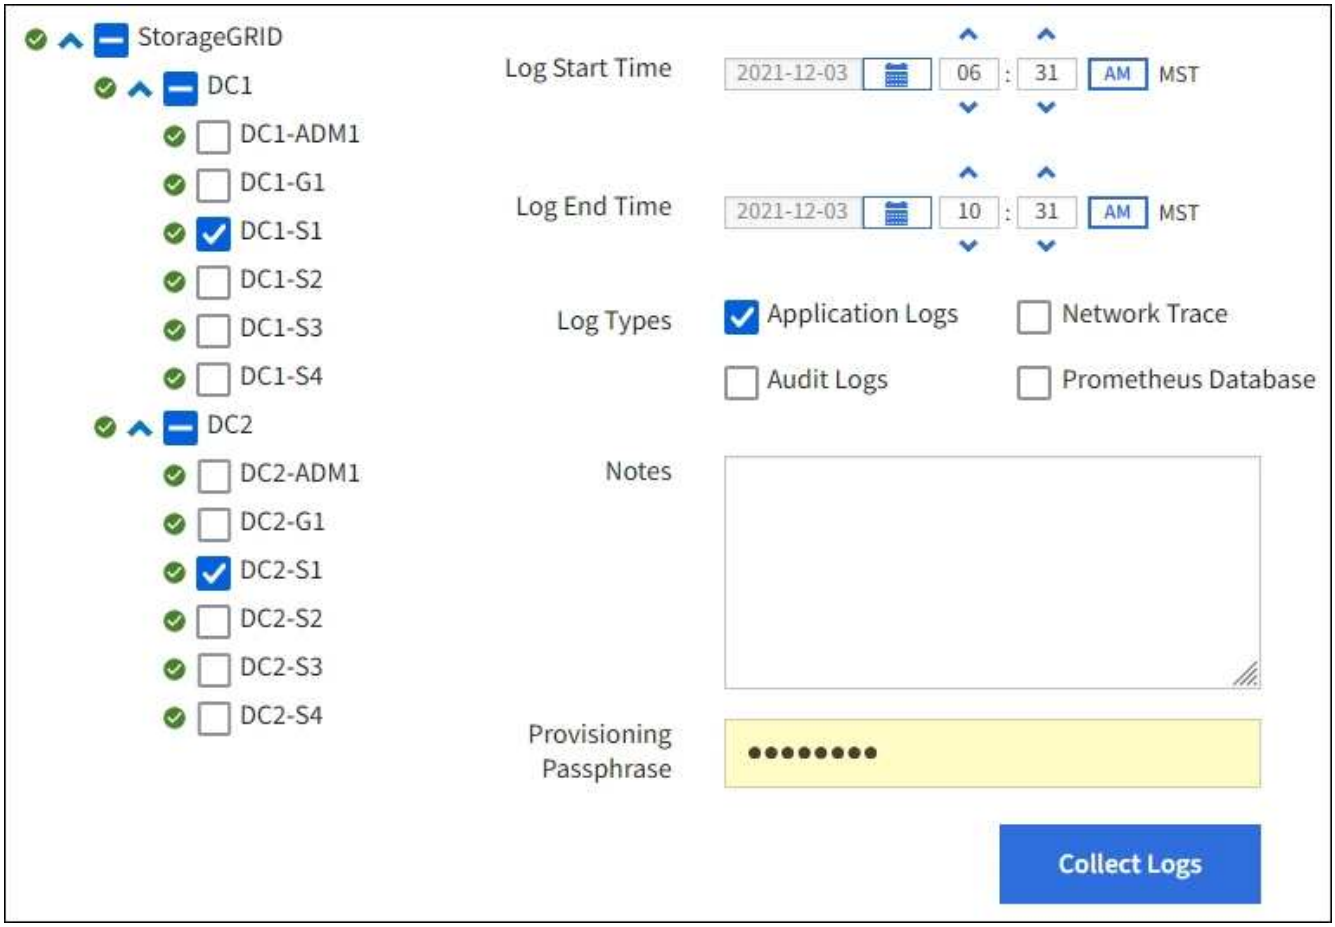


2. Select the grid nodes for which you want to collect log files.

   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 message

To assist technical support in troubleshooting issues with your StorageGRID system, you can manually trigger an AutoSupport message 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 **Settings** tab, select **Send User-Triggered AutoSupport**.

   StorageGRID attempts to send an AutoSupport message to technical support. 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 message again.

   After sending an User-triggered AutoSupport message, 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**.

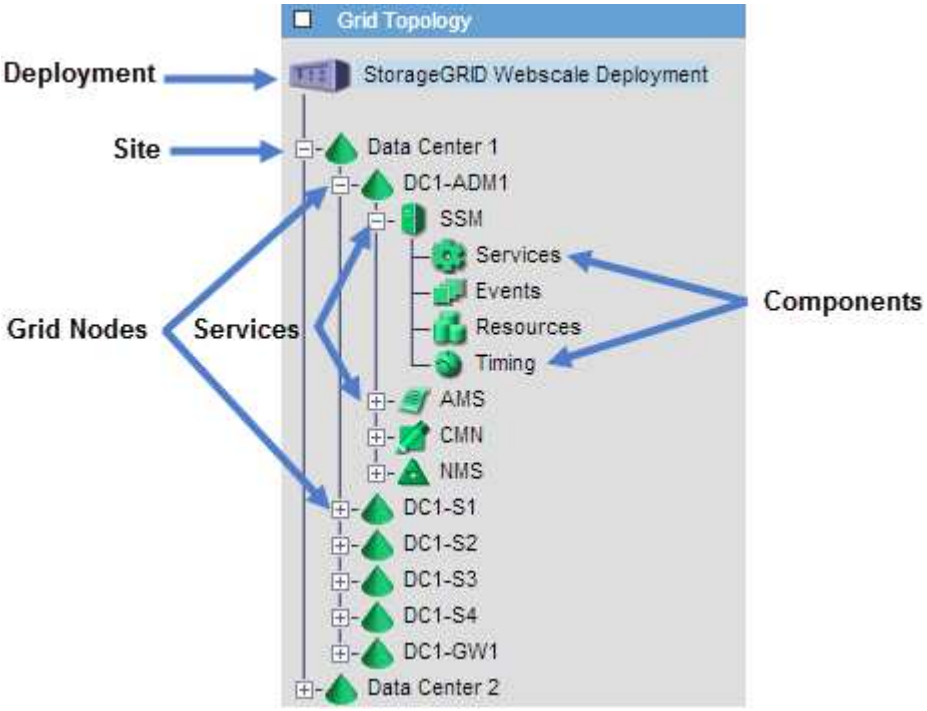


To expand or collapse the Grid Topology tree, click ⊞ or ⊟ at the site, node, or service level. To expand or collapse all items in the entire site or in each node, hold down the **<Ctrl>** key and click.

### 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 must 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:

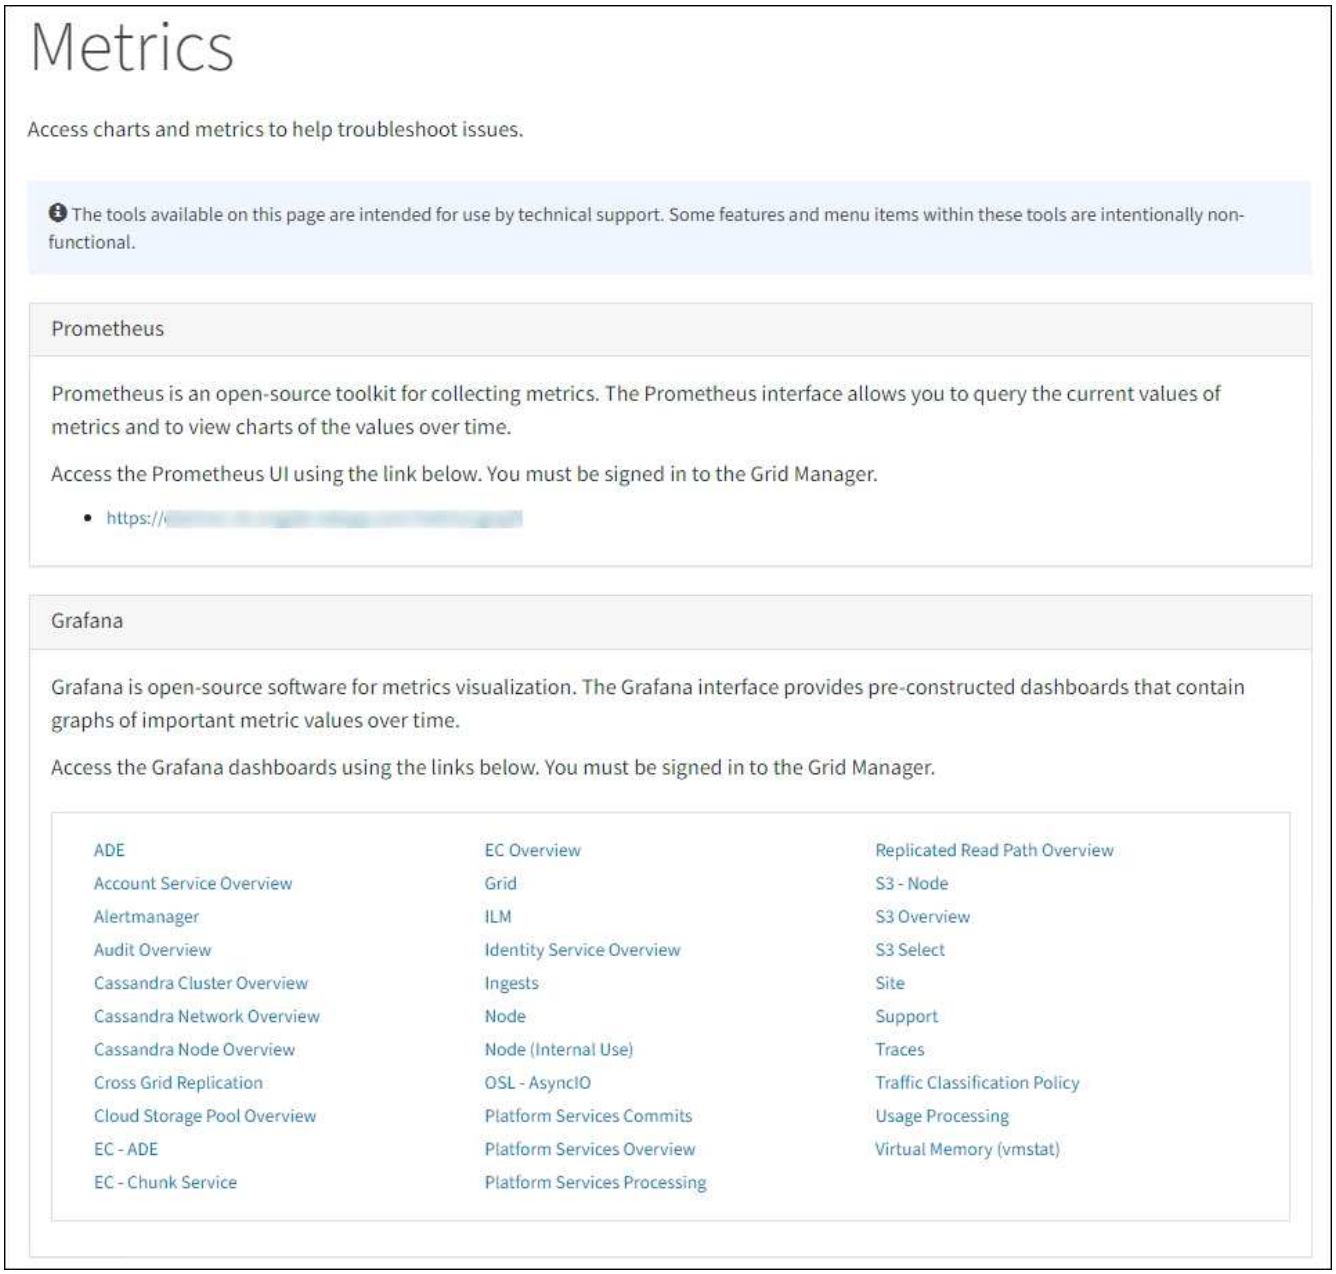


2. To query the current values of StorageGRID metrics and to view graphs of the values over time, click the link in the Prometheus section.

   The Prometheus interface appears. You can use this interface to execute queries on the available StorageGRID metrics and to graph StorageGRID metrics over time.

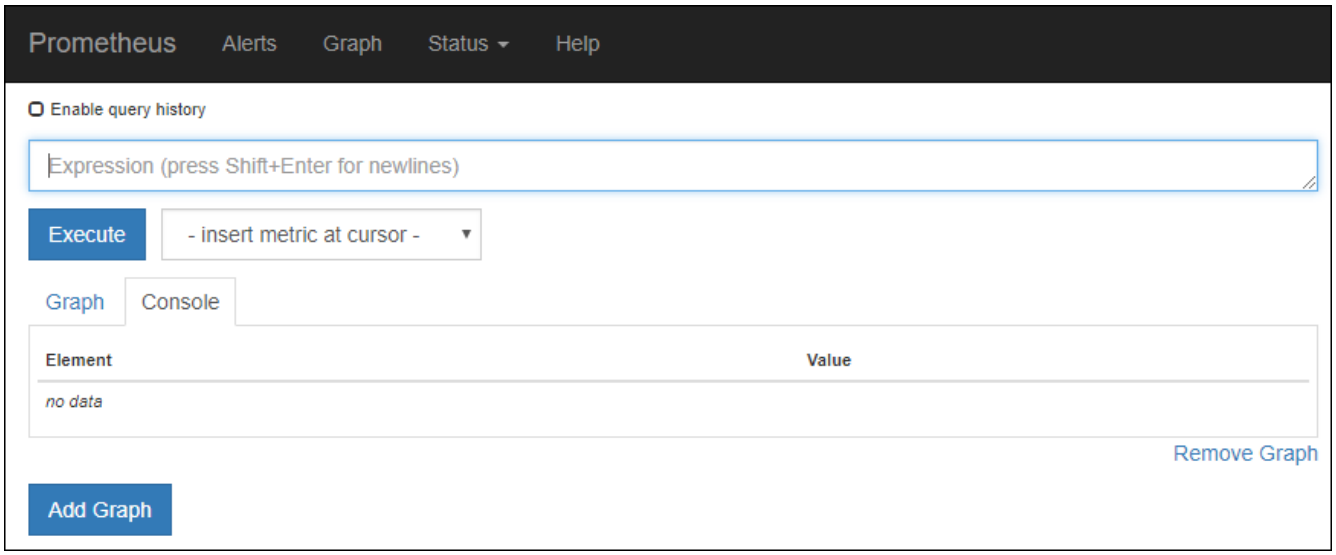


> Metrics that include *private* in their names are intended for internal use only and are subject to change between StorageGRID releases without notice.

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.

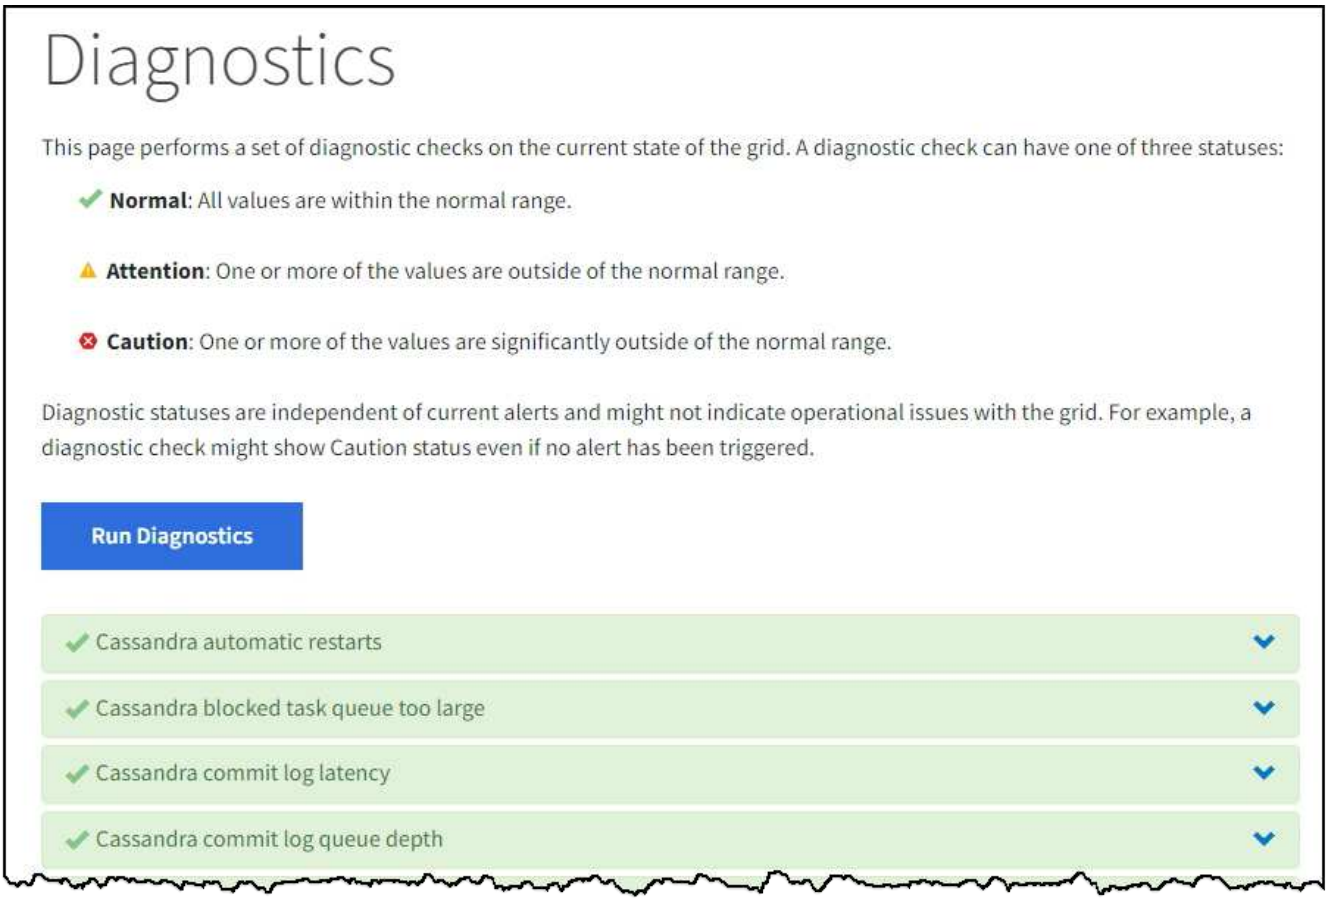


2. To learn more about a specific diagnostic, click anywhere in the row.

   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.

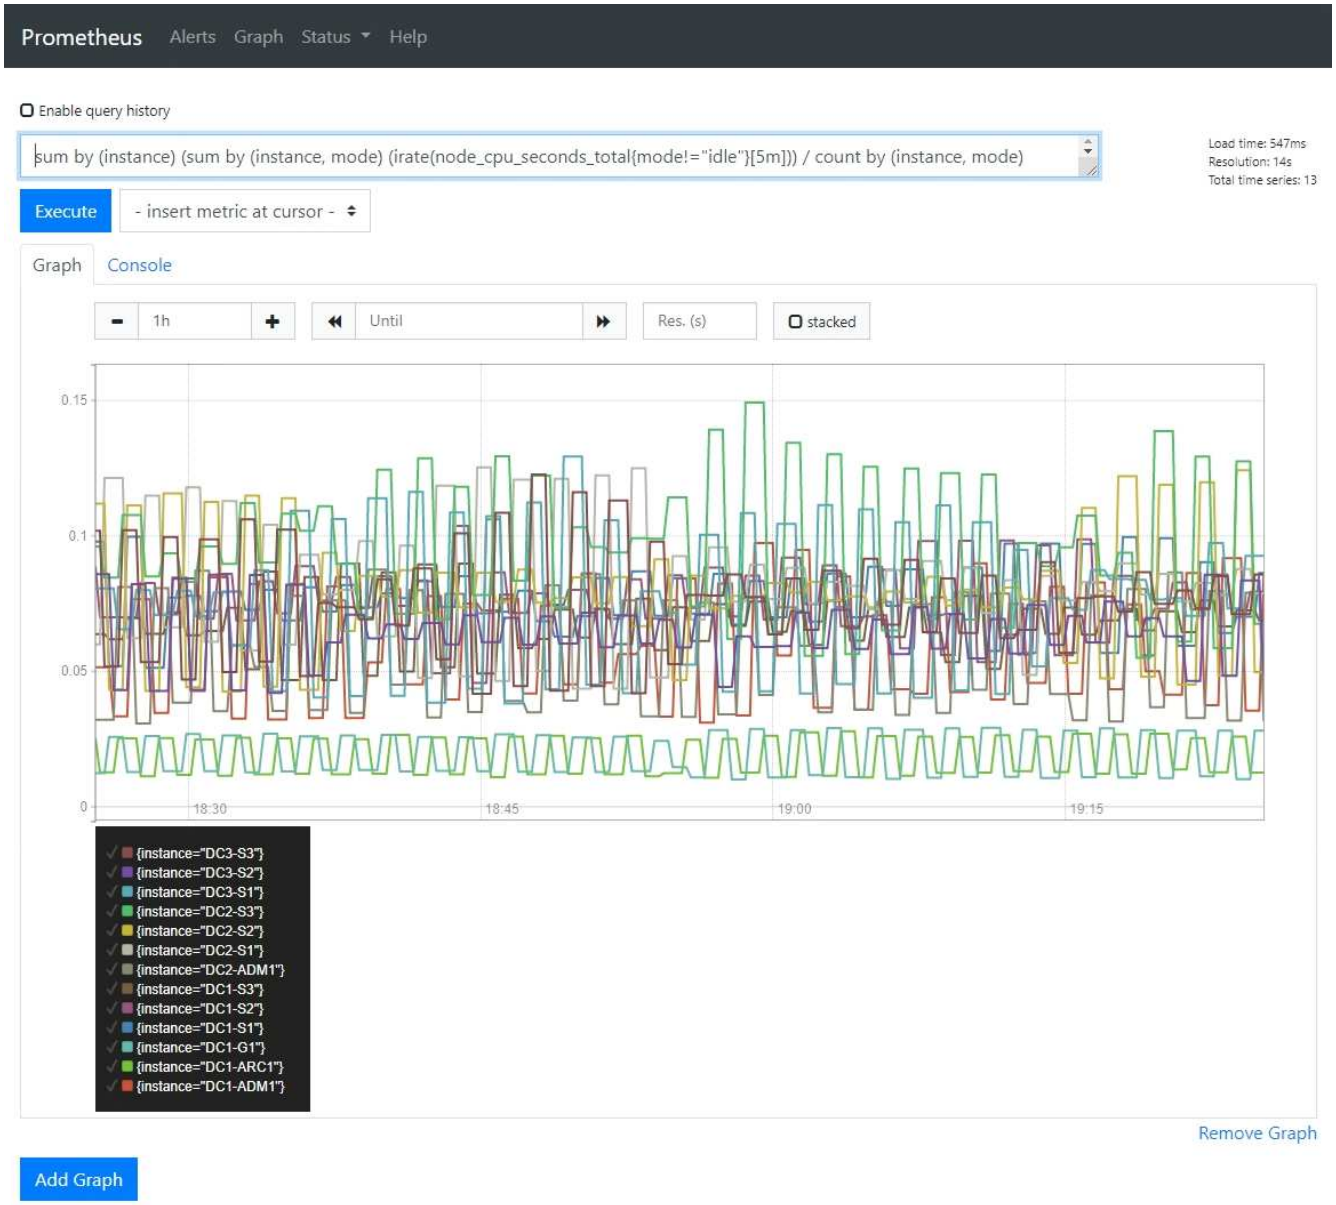


## Create custom monitoring applications

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.

**Copyright information**

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

**Trademark information**

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