



StorageGRID® 11.4

Monitoring and Troubleshooting Guide

December 2020 | 215-14713_2020-12_en-us
doccomments@netapp.com

 **NetApp®**

Contents

- Using the Grid Manager for monitoring..... 5**
 - Web browser requirements..... 5
 - Viewing the Dashboard.....5
 - Viewing the Nodes page..... 8
 - Viewing the Overview tab..... 9
 - Viewing the Hardware tab..... 11
 - Viewing the Network tab..... 12
 - Viewing the Storage tab..... 14
 - Viewing the Events tab..... 15
 - Using the Task tab to reboot a grid node..... 17
 - Viewing the Objects tab..... 18
 - Viewing the ILM tab..... 19
 - Viewing the Load Balancer tab.....20
 - Viewing the Platform Services tab.....22
 - Viewing information about appliance Storage Nodes..... 22
 - Viewing information about appliance Admin Nodes and Gateway Nodes..... 30

- Information you should monitor regularly.....36**
 - Monitoring system health..... 36
 - Monitoring node connection states..... 38
 - Viewing current alerts..... 39
 - Viewing resolved alerts.....41
 - Viewing a specific alert.....43
 - Viewing legacy alarms.....46
 - Monitoring storage capacity..... 48
 - Monitoring storage capacity for the entire grid..... 48
 - Monitoring storage capacity for each Storage Node..... 50
 - Monitoring object metadata capacity for each Storage Node.....53
 - Monitoring information lifecycle management 55
 - Monitoring performance, networking, and system resources.....57
 - Monitoring query latency..... 57
 - Monitoring network connections and performance..... 58
 - Monitoring node-level resources..... 61
 - Monitoring tenant activity..... 62
 - Monitoring archival capacity..... 64
 - Monitoring load balancing operations..... 65
 - Applying hotfixes or upgrading software if necessary..... 66

- Managing alerts and alarms 68**
 - Comparing alerts and alarms..... 68
 - Getting started with the alert system..... 71
 - Managing alerts..... 72
 - What alerts are..... 73
 - Managing alert rules..... 74
 - Managing alert notifications..... 81

- Silencing alert notifications..... 87
- Managing alarms (legacy system)..... 90
 - Alarm classes (legacy system)..... 90
 - Alarm triggering logic (legacy system)..... 91
 - Acknowledging current alarms (legacy system)..... 94
 - Viewing Default alarms (legacy system)..... 95
 - Reviewing historical alarms and alarm frequency (legacy system)..... 96
 - Creating Global Custom alarms (legacy system)..... 97
 - Disabling alarms (legacy system)..... 99
 - Configuring notifications for alarms (legacy system)..... 102

- Using SNMP monitoring..... 110**
 - Configuring the SNMP agent..... 111
 - Updating the SNMP agent..... 118

- Collecting additional StorageGRID data..... 121**
 - Using charts and reports..... 121
 - Types of charts and graphs..... 121
 - Displaying charts and graphs..... 124
 - Generating charts..... 127
 - Types of text reports..... 127
 - Generating text reports..... 128
 - Exporting text reports..... 129
 - Monitoring PUT and GET performance..... 130
 - Monitoring object verification operations..... 131
 - Monitoring events..... 132
 - Reviewing previous events..... 133
 - Resetting event counts..... 134
 - Creating custom syslog events..... 135
 - Resetting the count of custom events to zero..... 137
 - Reviewing audit messages..... 138
 - Collecting log files and system data..... 139
 - Triggering an AutoSupport message..... 141
 - Viewing the Grid Topology tree..... 141
 - Reviewing support metrics..... 142
 - Running diagnostics..... 144
 - Creating custom monitoring applications..... 147

- Troubleshooting a StorageGRID system..... 149**
 - Overview of problem determination..... 149
 - Defining the problem..... 149
 - Assessing the risk and impact on the system..... 150
 - Collecting data..... 150
 - Analyzing data..... 154
 - Escalation information checklist..... 154
 - Troubleshooting object and storage issues..... 155
 - Confirming object data locations..... 155
 - Object store (storage volume) failures..... 157
 - Verifying object integrity..... 159
 - Troubleshooting lost and missing object data..... 165
 - Troubleshooting the Low object data storage alert..... 174

- Troubleshooting the Total Usable Space (Percent) (SAVP) alarm..... 176
- Troubleshooting the Storage Status (SSTS) alarm..... 177
- Troubleshooting delivery of platform services messages (SMTT alarm).....181
- Troubleshooting the S3 multipart part too small alert..... 181
- Troubleshooting metadata issues..... 183
 - Troubleshooting the Low metadata storage alert..... 183
 - Troubleshooting the Services: Status - Cassandra (SVST) alarm..... 185
 - Troubleshooting Cassandra Out of Memory errors (SMTT alarm).....188
- Troubleshooting certificate errors..... 188
- Troubleshooting Admin Node and user interface issues..... 190
 - Troubleshooting sign-on errors 190
 - Troubleshooting user interface issues..... 192
 - Checking the status of an unavailable Admin Node..... 193
- Troubleshooting network, hardware, and platform issues..... 194
 - Troubleshooting "422: Unprocessable Entity" errors..... 194
 - Troubleshooting the Grid Network MTU mismatch alert..... 195
 - Troubleshooting the Network Receive Error (NRER) alarm.....196
 - Troubleshooting time synchronization errors..... 197
 - Linux: Network connectivity issues..... 198
 - Linux: Node status is "orphaned" 198
 - Linux: Troubleshooting IPv6 support..... 199

- Alerts reference..... 201**
 - Commonly used Prometheus metrics..... 217

- Alarms reference (legacy system)..... 222**
 - Alarms that generate SNMP notifications (legacy system)..... 243

- Log files reference..... 246**
 - StorageGRID software logs..... 246
 - Deployment and maintenance logs..... 250
 - Logs for third-party software..... 250
 - About the bycast.log..... 251
 - File rotation for bycast.log..... 251
 - Messages in bycast.log..... 252
 - Message severities in bycast.log..... 252
 - Error codes in bycast.log..... 252

- Copyright and trademark..... 258**
 - Copyright..... 258
 - Trademark..... 258

Using the Grid Manager for monitoring

The Grid Manager is the most important tool for monitoring your StorageGRID system. This section introduces the Grid Manager Dashboard and provides detailed information about the Nodes pages.

Web browser requirements

You must use a supported web browser.

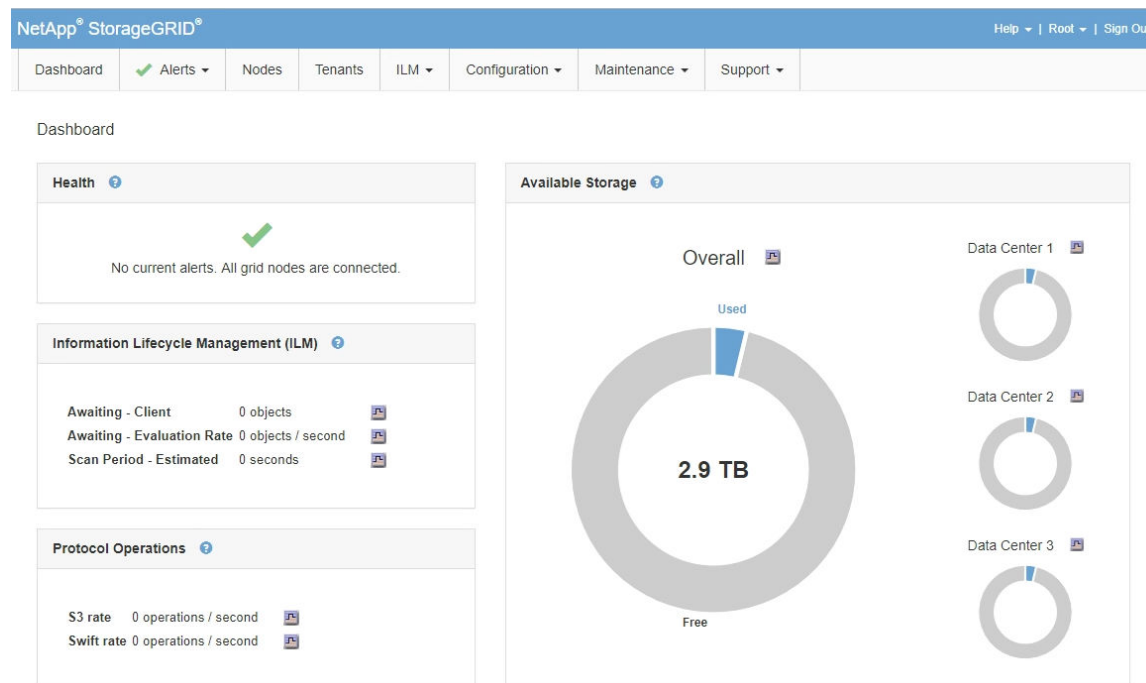
Web browser	Minimum supported version
Google Chrome	74
Microsoft Internet Explorer	11 (Native Mode)
Mozilla Firefox	67

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Viewing the Dashboard


When you first sign in to the Grid Manager, you can use the Dashboard to monitor system activities at a glance. The Dashboard includes information about system health, usage metrics, and operational trends and charts.



Health panel

Description	View additional details	Learn more
<p>Summarizes the system's health. A green checkmark means that there are no current alerts and all grid nodes are connected. Any other icon means that there is at least one current alert or disconnected node.</p>	<p>You might see one or more of the following links:</p> <ul style="list-style-type: none"> • Grid details: Appears if any nodes are disconnected (connection state Unknown or Administratively Down). Click the link, or click the blue or gray icon to determine which node or nodes are affected. • Current alerts: Appears if any alerts are currently active. Click the link, or click Critical, Major, or Minor to see the details on the Alerts > Current page. • Recently resolved alerts: Appears if any alerts triggered in the past week are now resolved. Click the link to see the details on the Alerts > Resolved page. • Legacy alarms: Appears if any alarms (legacy system) are currently active. Click the link to see the details on the Support > Current Alarms page. • License: Appears if there is an issue with the software license for this StorageGRID system. Click the link to see the details on the Maintenance > License page. 	<ul style="list-style-type: none"> • Monitoring node connection states on page 38 • Viewing current alerts on page 39 • Viewing resolved alerts on page 41 • Viewing legacy alarms on page 46 • Administering StorageGRID


Available Storage panel

Description	View additional details	Learn more
<p>Displays the available and used storage capacity in the entire grid, not including archival media.</p> <p>The Overall chart presents grid-wide totals. If this is a multi-site grid, additional charts appear for each data center site.</p> <p>You can use this information to compare the used storage with the available storage. If you have a multi-site grid, you can determine which site is consuming more storage.</p>	<ul style="list-style-type: none"> • To view the capacity, place your cursor over the chart's available and used capacity sections. • To view capacity trends over a date range, click the chart icon  for the overall grid, or for a data center site. • To see details, select Nodes. Then, view the Storage tab for the entire grid, an entire site, or a single Storage Node. 	<ul style="list-style-type: none"> • Viewing the Storage tab on page 14 • Monitoring storage capacity on page 48

Information Lifecycle Management (ILM) panel

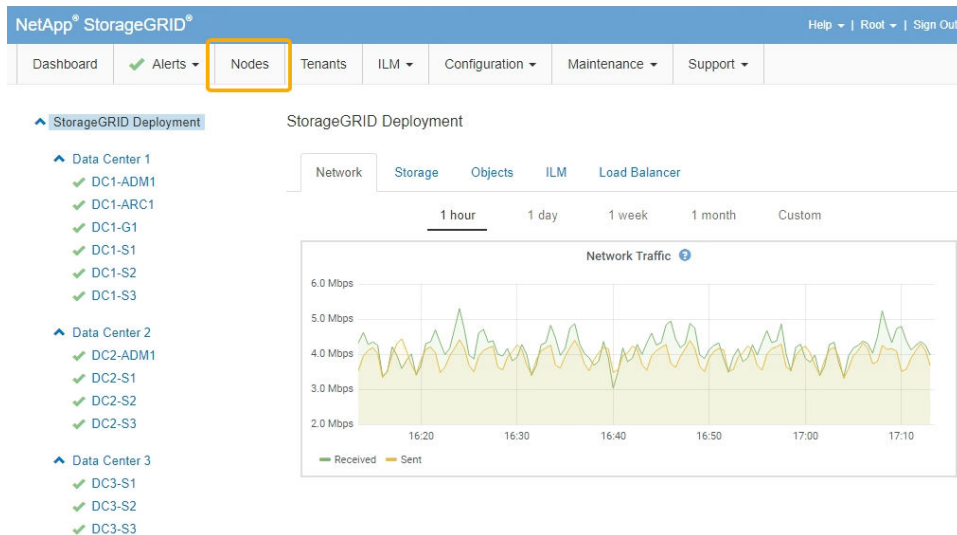
Description	View additional details	Learn more
<p>Displays current ILM operations and ILM queues for your system. You can use this information to monitor your system's workload.</p> <ul style="list-style-type: none"> • Awaiting - Client: The total number of objects awaiting ILM evaluation from client operations (for example, ingest). • Awaiting - Evaluation Rate: The current rate at which objects are evaluated against the ILM policy in the grid. • Scan Period - Estimated: The estimated time to complete a full ILM scan of all objects. <p>Note: A full scan does not guarantee that ILM has been applied to all objects.</p>	<ul style="list-style-type: none"> • To see details, select Nodes. Then, view the ILM tab for the entire grid, an entire site, or a single Storage Node. • To see the existing ILM rules, select ILM > Rules. • To see the existing ILM policies, select ILM > Policies. 	<ul style="list-style-type: none"> • Viewing the ILM tab on page 19 • Administering StorageGRID.

Protocol Operations panel

Description	View additional details	Learn more
<p>Displays the number of protocol-specific operations (S3 and Swift) performed by your system. You can use this information to monitor your system's workloads and efficiencies. Protocol rates are averaged over the last two minutes.</p>	<ul style="list-style-type: none"> • To see details, select Nodes. Then, view the Objects tab for the entire grid, an entire site, or a single Storage Node. • To view trends over a date range, click the chart icon  to the right of the S3 or Swift protocol rate. 	<ul style="list-style-type: none"> • Viewing the Objects tab on page 18 • Implementing S3 client applications • Implementing Swift client applications

Viewing the Nodes page


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




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

Connection state icons

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



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



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

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

Alert icons

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

- **Critical** : An abnormal condition exists that has stopped the normal operations of a StorageGRID node or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved.
- **Major** : An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service.

- **Minor** : The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that do not clear on their own to ensure they do not result in a more serious problem.
- **Normal** : No alerts are active, and the node is connected to the grid.

Viewing details for a system, site, or node

To view the available information, click the appropriate links on the left, as follows:

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

Choices

- [Viewing the Overview tab](#) on page 9
- [Viewing the Hardware tab](#) on page 11
- [Viewing the Network tab](#) on page 12
- [Viewing the Storage tab](#) on page 14
- [Viewing the Events tab](#) on page 15
- [Using the Task tab to reboot a grid node](#) on page 17
- [Viewing the Objects tab](#) on page 18
- [Viewing the ILM tab](#) on page 19
- [Viewing the Load Balancer tab](#) on page 20
- [Viewing the Platform Services tab](#) on page 22
- [Viewing information about appliance Storage Nodes](#) on page 22
- [Viewing information about appliance Admin Nodes and Gateway Nodes](#) on page 30

Viewing 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 grid node.

DC1-S1 (Storage Node)

The screenshot shows the Grid Manager interface for a Storage Node named DC1-S1. At the top, there are tabs for Overview, Hardware, Network, Storage, Objects, ILM, Events, and Tasks. The 'Overview' tab is selected. Below the tabs is a 'Node Information' section with a help icon. It displays the following details: Name (DC1-S1), Type (Storage Node), ID (5bf57bd4-a68d-467e-b866-bfe09a5c6b96), Connection State (Connected with a green checkmark icon), Software Version (11.4.0 (build 20200328.0051.269ac98)), and IP Addresses (10.96.101.111 with a 'Show more' dropdown). Below this is an 'Alerts' section with a help icon, showing a green checkmark icon and the text 'No active alerts'.

The overview information for a node includes the following:

- **Name:** The hostname assigned to the node and displayed in the Grid Manager.
- **Type:** The type of node — Admin Node, Storage Node, Gateway Node, or Archive Node.
- **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.
 - **Not connected - Unknown** : The node is not connected to the grid for an unknown reason. For example, the network connection between nodes has been lost or the power is down. The **Unable to communicate with node** alert might also be triggered. Other alerts might be active as well. This situation requires immediate attention.

Note: A node might appear as Unknown during managed shutdown operations. You can ignore the Unknown state in these cases.
 - **Not connected - Administratively down** : The node is not connected to the grid for an expected reason. For example, the node, or services on the node, has been gracefully shut down, the node is rebooting, or the software is being upgraded. One or more alerts might also be active.
 - **Connected** : The node is connected to the grid.
- **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 Master or the Backup.

DC1-ADM1 (Admin Node)

Overview Hardware Network Storage Load Balancer Events Tasks

Node Information

Name DC1-ADM1
 Type Admin Node
 ID 711b7b9b-8d24-4d9f-877a-be3fa3ac27e8
 Connection State ✔ Connected
 Software Version 11.4.0 (build 20200515.2346.8edcbbf)
HA Groups Fabric Pools, Master
 IP Addresses 192.168.2.208, 10.224.2.208, 47.47.2.208, 47.47.4.219 [Show more](#)

- **IP Addresses:** The node's IP addresses. Click **Show more** to view the node's IPv4 and IPv6 addresses and interface mappings:
 - eth0: Grid Network
 - eth1: Admin Network
 - eth2: Client Network

Alerts

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

Alerts

Name	Severity	Time triggered	Current values
Low installed node memory The amount of installed memory on a node is low.	✘ Critical	18 hours ago	Total RAM size: 8.37 GB

Related tasks

[Monitoring node connection states](#) on page 38

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

[Viewing current alerts](#) on page 39

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. An email notification might also be sent, unless the alert has been silenced.

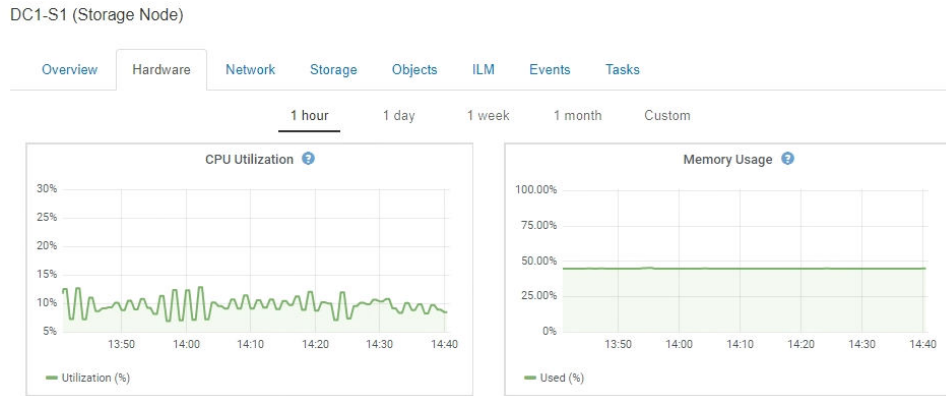
[Viewing a specific alert](#) on page 43

You can view detailed information about an alert that is currently affecting your StorageGRID system or an alert that has been resolved. The details include recommended corrective actions, the time the alert was triggered, and the current value of the metrics related to this alert. Optionally, you can silence a current alert or update the alert rule.

Viewing 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, 1 month, or 1 year. 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, hover your cursor over each graph.



If the node is an appliance Storage Node or an appliance Admin Node or Gateway Node, this tab also includes a section with more information about the appliance hardware.

Related tasks

[Viewing information about appliance Storage Nodes](#) on page 22

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, network resources, network interfaces, network addresses, and receive and transmit data.

[Viewing information about appliance Admin Nodes and Gateway Nodes](#) on page 30

The Nodes page lists information about service health and all computational, disk device, and network resources for each services appliance that is used for 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.

Viewing 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, 1 month, or 1 year. 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.

DC1-S1-226 (Storage Node)

Overview
Hardware
Network
Storage
Objects
ILM
Events

1 hour
1 day
1 week
1 month
1 year
Custom

Network Interfaces

Name	Hardware Address	Speed	Duplex	Auto Negotiate	Link Status
eth0	00:50:56:A8:2A:75	10 Gigabit	Full	Off	Up

Network Communication

Receive

Interface	Data	Packets	Errors	Dropped	Frame Overruns	Frames
eth0	738.858 GB	904,587,345	0	14,340	0	0

Transmit

Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	677.555 GB	465,715,998	0	0	0	0

Related tasks

[Monitoring network connections and performance](#) on page 58

Grid nodes must be able to communicate with one another to permit the grid to operate. The integrity of the network between nodes and sites, and the network bandwidth between sites, are critical to efficient operations.

Viewing the Storage tab

The Storage tab summarizes storage availability and other storage metrics.

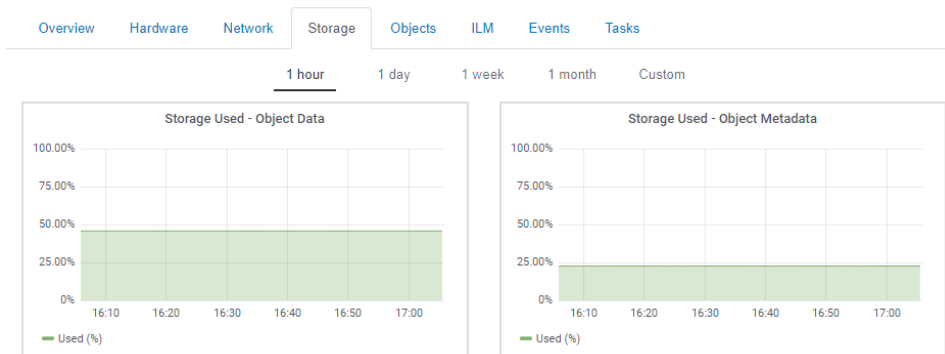
The Storage tab is shown for all nodes, each site, and the entire grid.

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.

Note: The total values for a site or the grid do not include nodes that not have reported metrics for at least five minutes, such as offline nodes.

DC1-SN1-99-88 (Storage Node)



Disk Devices and Volumes charts

For all nodes, the Storage tab also contains details for the disk devices and volumes on the node.

DC1-ADM1-225 (Admin Node)

Overview Hardware Network **Storage** Events

Disk Devices

Name	World Wide Name	I/O Load	Read Rate	Write Rate
croot(8:1,sda1)	N/A	0.44%	0 bytes/s	14 KB/s
cvloc(8:2,sda2)	N/A	8.85%	0 bytes/s	517 KB/s
sdc(8:16,sdb)	N/A	0.11%	0 bytes/s	3 KB/s
sdd(8:32,sdc)	N/A	1.90%	0 bytes/s	2 MB/s

Volumes

Mount Point	Device	Status	Size	Available	Write Cache Status
/	croot	Online	10.50 GB	3.29 GB	Unknown
/var/local	cvloc	Online	96.59 GB	92.19 GB	Unknown
/var/local/audit/export	sdc	Online	214.64 GB	214.38 GB	Enabled
/var/local/mysql_ibdata	sdd	Online	214.64 GB	213.88 GB	Enabled

Related tasks

[Monitoring storage capacity for the entire grid](#) on page 48

You must 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.

[Monitoring storage capacity for each Storage Node](#) on page 50

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

[Monitoring object metadata capacity for each Storage Node](#) on page 53

You must 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.

Viewing the Events tab

The Events tab displays a count of any system error or fault events for a node, including errors such as network errors.

The Events tab is shown for all nodes.

If you experience issues with a particular node, you can use the Events tab to learn more about the issue. Technical support can also use the information on the Events tab to help with troubleshooting.

Events ?	
Last Event	No Events
Description	Count
Abnormal Software Events	0
Account Service Events	0
Cassandra Heap Out Of Memory Errors	0
Cassandra unhandled exceptions	0
Chunk Service Events	0
Custom Events	0
Data-Mover Service Events	0
File System Errors	0
Forced Termination Events	0
Hotfix Installation Failure Events	0
I/O Errors	0
IDE Errors	0
Identity Service Events	0
Kernel Errors	0
Kernel Memory Allocation Failure	0
Keystone Service Events	0
Network Receive Errors	0
Network Transmit Errors	0
Node Errors	0
Out Of Memory Errors	0
Replicated State Machine Service Events	0
SCSI Errors	0
Stat Service Events	0
Storage Hardware Events	0
System Time Events	0

[Reset event counts](#)

You can perform these tasks from the Events tab:

- Use the information shown for the **Last Event** field at the top of the table to determine which event occurred most recently.
- Click the chart icon for a specific event to see when that event occurred over time.
- Reset event counts to zero after resolving any issues.

Related concepts

[Monitoring events](#) on page 132

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.

Related tasks

[Displaying charts and graphs](#) on page 124

The Nodes page contains the graphs and charts 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 > Grid Topology** page to access additional charts.

[Resetting event counts](#) on page 134

After resolving system events, you can reset event counts to zero.

Using the Task tab to reboot a grid node

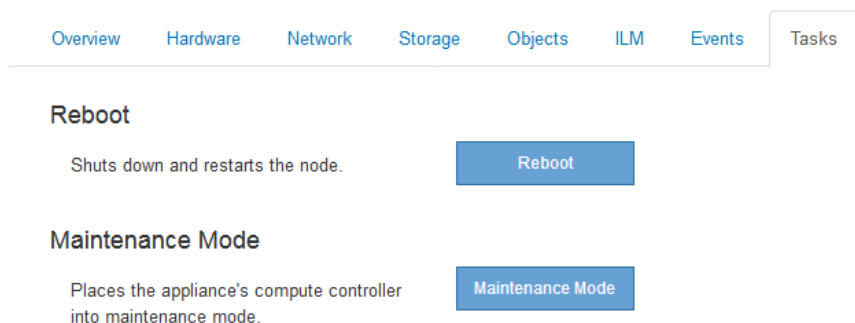
The Task tab allows you to reboot the selected node. The Task tab is shown for all nodes.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must 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:

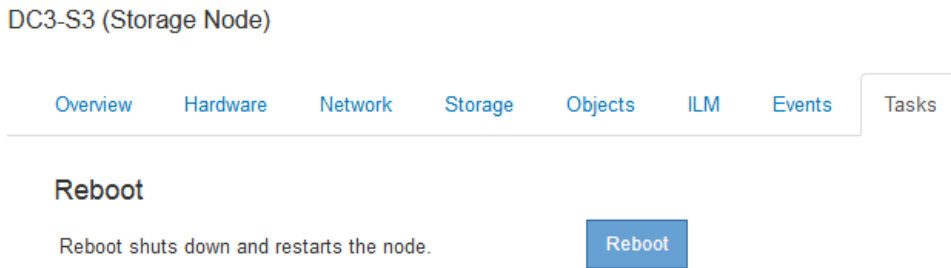
- If an ILM rule specifies an ingest behavior of Dual commit or the rule specifies Balanced and it is not possible to immediately create all required copies, StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site and evaluates ILM later. If you want to reboot 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 the hardware installation and maintenance instructions for the appliance.

Note: Putting an appliance into maintenance mode might make the appliance unavailable for remote access.

Steps

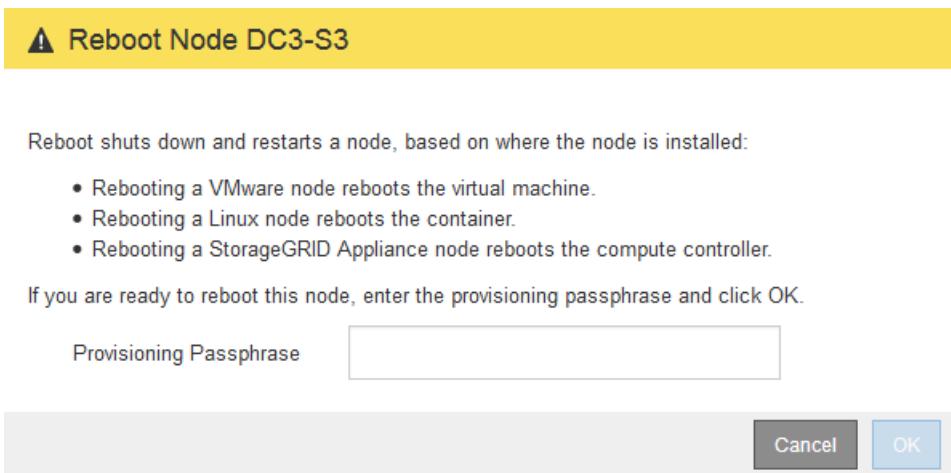
1. Select **Nodes**.
2. Select the grid node you want to reboot.

3. Select the **Tasks** tab.



4. Click **Reboot**.

A confirmation dialog box appears.



Note: 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 click **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, the icon changes back to its original color.

Related information

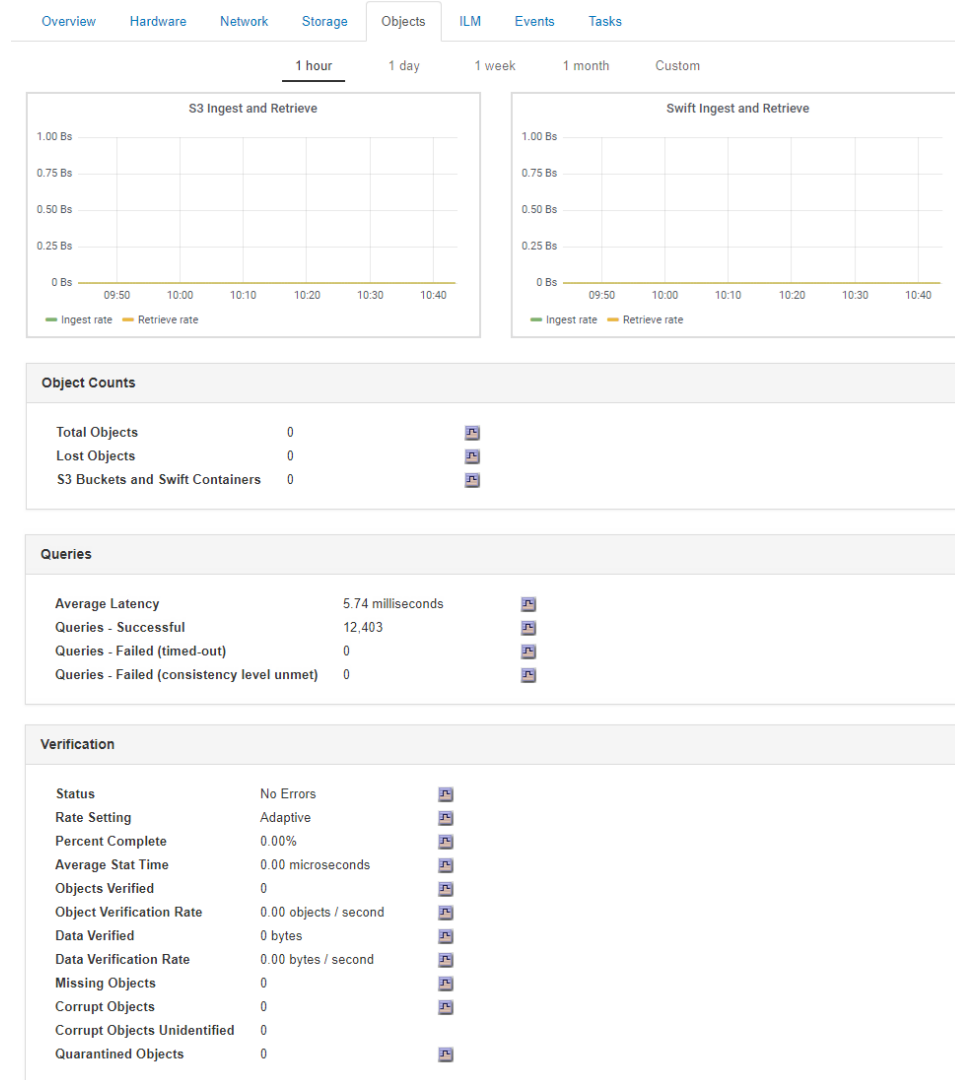
- [SG6000 appliance installation and maintenance](#)
- [SG5700 appliance installation and maintenance](#)
- [SG5600 appliance installation and maintenance](#)
- [SG100 and SG1000 appliance installation and maintenance](#)

Viewing 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)



Related information

- [Implementing S3 client applications](#)
- [Implementing Swift client applications](#)

Viewing 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.







DC1-S1 (Storage Node)

Overview Hardware Network Storage Objects **ILM** Events

Evaluation

Awaiting - All	0 objects	
Awaiting - Client	0 objects	
Evaluation Rate	0.00 objects / second	
Scan Rate	0.00 objects / second	

Erasure Coding Verification

Status	Idle	
Next Scheduled	2018-05-23 10:44:47 MDT	
Fragments Verified	0	
Data Verified	0 bytes	
Corrupt Copies	0	
Corrupt Fragments	0	
Missing Fragments	0	

Related tasks

[Monitoring information lifecycle management](#) on page 55

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.

Related information

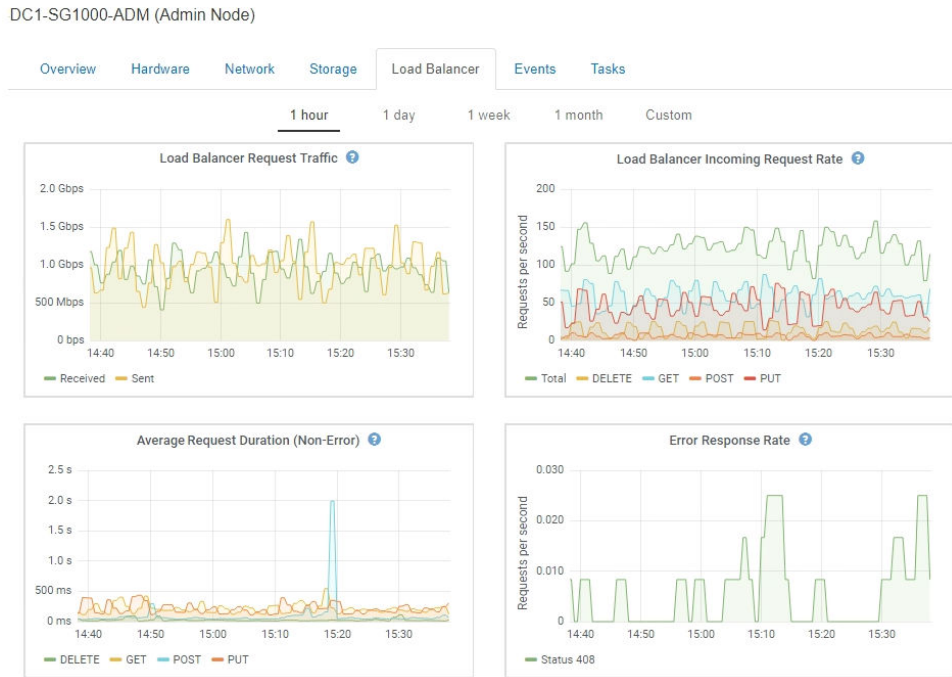
[Administering StorageGRID](#)

Viewing 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."



Load Balancer 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.

Note: 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.

Load Balancer 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 tasks

[Monitoring load balancing operations](#) on page 65

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.

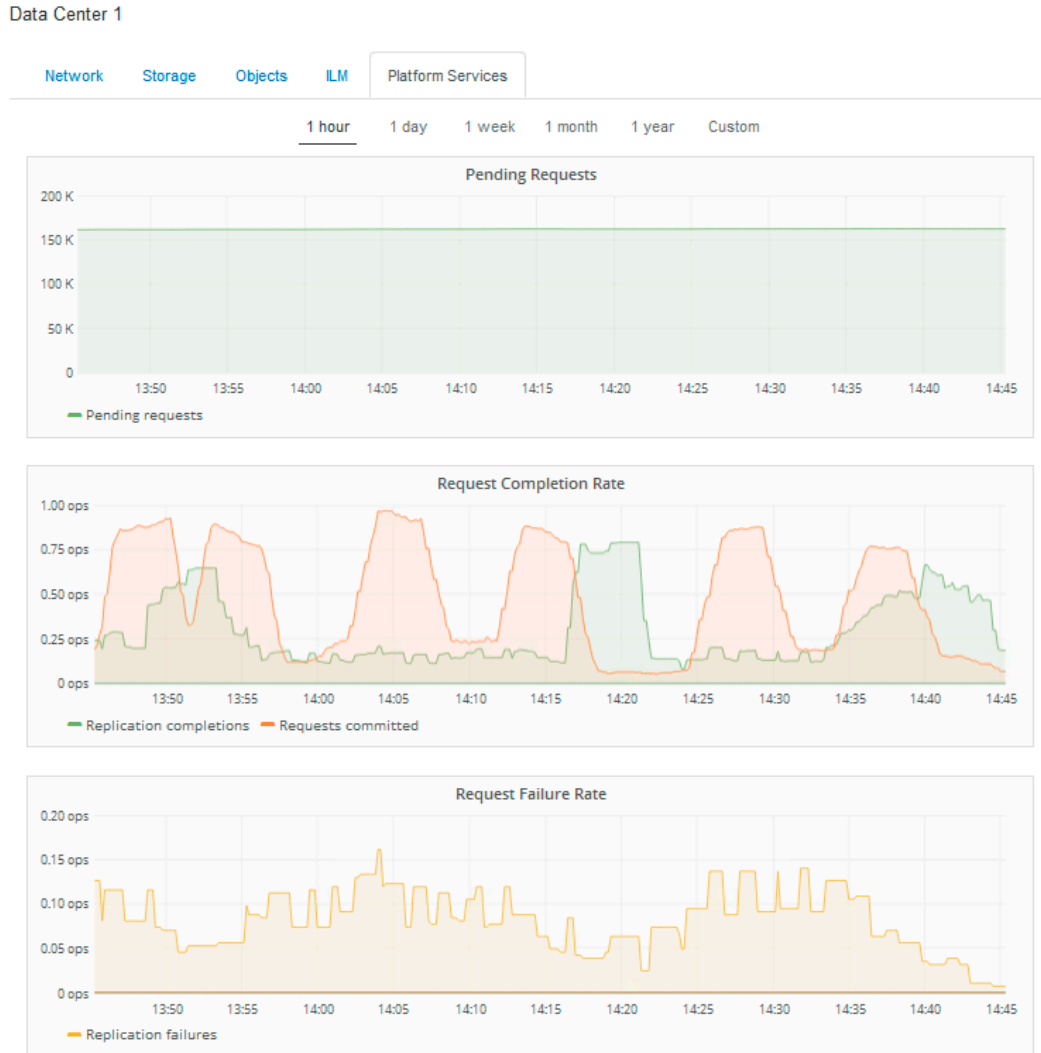
Related information

[Administering StorageGRID](#)

Viewing 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.

Related information

[Administering StorageGRID](#)

Viewing 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, 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 table on the Overview tab displays the node's ID and name, the node type, the software version installed, and the IP addresses associated with the node. The Interface column contains the name of the interface, 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, which can be bonded or aliased and connected to the StorageGRID Admin Network (eth1).

Node Information	
Name	SGA-lab11
Type	Storage Node
ID	0b583829-6659-4c6e-b2d0-31461d22ba67
Connection State	✔ Connected
Software Version	11.4.0 (build 20200527.0043.61839a2)
IP Addresses	192.168.4.138, 10.224.4.138, 169.254.0.1 Show less
Interface	IP Address
eth0	192.168.4.138
eth0	fd20:331:331:0:2a0:98ff:fea1:831d
eth0	fe80::2a0:98ff:fea1:831d
eth1	10.224.4.138
eth1	fd20:327:327:0:280:e5ff:fe43:a99c
eth1	fd20:8b1e:b255:8154:280:e5ff:fe43:a99c
eth1	fe80::280:e5ff:fe43:a99c
hic2	192.168.4.138
hic4	192.168.4.138
mtc1	10.224.4.138
mtc2	169.254.0.1

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, 1 month, or 1 year. 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.

Note: 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.

StorageGRID Appliance	
Appliance Model	SG6060
Storage Controller Name	StorageGRID-SG6060wShelf
Storage Controller A Management IP	10.224.0.18
Storage Controller B Management IP	10.224.0.19
Storage Controller WWID	6d039ea000016f61000000005ec2c245
Storage Appliance Chassis Serial Number	721924500064
Storage Hardware	Nominal 
Storage Controller Failed Drive Count	0 
Storage Controller A	Nominal 
Storage Controller B	Nominal 
Storage Controller Power Supply A	Nominal 
Storage Controller Power Supply B	Nominal 
Storage Data Drive Type	NL-SAS HDD
Storage Data Drive Size	4.00 TB
Storage RAID Mode	RAID6
Storage Connectivity	Nominal 
Overall Power Supply	Nominal 
Compute Controller BMC IP	10.224.0.11
Compute Controller Serial Number	721917500026
Compute Hardware	Nominal 
Compute Controller CPU Temperature	Nominal 
Compute Controller Chassis Temperature	Nominal 

Storage Shelves													
Shelf Chassis Serial Number	Shelf ID	Shelf Status	IOM Status	Power Supply Status	Drawer Status	Fan Status	Drive Slots	Data Drives	Data Drive Size	Cache Drives	Cache Drive Size	Configuration Status	
721924500064	99	Nominal 	N/A	Nominal	Nominal	Nominal	60	58	4.00 TB	2	800.17 GB	Configured (in use)	
721929500041	0	Nominal 	Nominal	Nominal	Nominal	Nominal	60	60	4.00 TB	0	0 bytes	Configured (in use)	
721929500040	1	Nominal 	Nominal	Nominal	Nominal	Nominal	60	60	4.00 TB	0	0 bytes	Configured (in use)	

Field in the Appliance table	Description
Appliance Model	The model number for this StorageGRID appliance shown in SANtricity software.
Storage Controller Name	The name for this StorageGRID appliance shown in SANtricity software.
Storage Controller A Management IP	IP address for management port 1 on storage controller A. You use this IP to access SANtricity software to troubleshoot storage issues.
Storage Controller B Management IP	IP address for management port 1 on storage controller B. You use this IP to access SANtricity software to troubleshoot storage issues. Some appliance models do not have a storage controller B.
Storage Controller WWID	The worldwide identifier of the storage controller shown in SANtricity software.
Storage Appliance Chassis Serial Number	The chassis serial number of the appliance.

Field in the Appliance table	Description
Storage Hardware	The overall status of the storage controller hardware. If SANtricity System Manager reports a status of Needs Attention for the storage hardware, the StorageGRID system also reports this value. If the status is "needs attention," first check the storage controller using SANtricity software. Then, ensure that no other alarms exist that apply to the compute controller.
Storage Controller Failed Drive Count	The number of drives that are not optimal.
Storage Controller A	The status of storage controller A.
Storage Controller B	The status of storage controller B. Some appliance models do not 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 disk drive) or SSD (solid state drive).
Storage Data Drive Size	The total capacity including all data drives in the appliance.
Storage RAID Mode	The RAID mode configured for the appliance.
Storage Connectivity	The storage connectivity state.
Overall Power Supply	The status of all power supplies for the appliance.
Compute Controller BMC IP	The IP address of the baseboard management controller (BMC) port in the compute controller. You use this IP to connect to the BMC interface to monitor and diagnose the appliance hardware. This field is not displayed for appliance models that do not 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 do not 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.


Column in the Storage Shelves table	Description
Shelf ID	The numeric identifier for the storage shelf. <ul style="list-style-type: none"> • 99: Storage controller shelf • 0: First expansion shelf • 1: Second expansion shelf Note: Expansion shelves apply to the SG6060 only.
Shelf Status	The overall status of the storage shelf.
IOM Status	The status of the input/output modules (IOMs) in any expansion shelves. N/A if this is not an expansion shelf.
Power Supply Status	The overall status of the power supplies for the storage shelf.
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 Negotiate	Link Status
eth0	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
eth1	D8:C4:97:2A:E4:9E	Gigabit	Full	Off	Up
eth2	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
hic1	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic2	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic3	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic4	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
mtc1	D8:C4:97:2A:E4:9E	Gigabit	Full	On	Up
mtc2	D8:C4:97:2A:E4:9F	Gigabit	Full	On	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.

Note: 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 the installation and maintenance instructions for your appliance 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	3.250 TB	5,610,578,144	0	8,327	0	0	
eth1	1.205 GB	9,828,095	0	32,049	0	0	
eth2	849.829 GB	186,349,407	0	10,269	0	0	
hic1	114.864 GB	303,443,393	0	0	0	0	
hic2	2.315 TB	5,351,180,956	0	305	0	0	
hic3	1.690 TB	1,793,580,230	0	0	0	0	
hic4	194.283 GB	331,640,075	0	0	0	0	
mtc1	1.205 GB	9,828,096	0	0	0	0	
mtc2	1.168 GB	9,564,173	0	32,050	0	0	

Transmit							
Interface	Data	Packets	Errors	Dropped	Collisions	Carrier	
eth0	5.759 TB	5,789,638,626	0	0	0	0	
eth1	4.563 MB	41,520	0	0	0	0	
eth2	855.404 GB	139,975,194	0	0	0	0	
hic1	289.248 GB	326,321,151	5	0	0	5	
hic2	1.636 TB	2,640,416,419	18	0	0	18	
hic3	3.219 TB	4,571,516,003	33	0	0	33	
hic4	1.687 TB	1,658,180,262	22	0	0	22	
mtc1	4.563 MB	41,520	0	0	0	0	
mtc2	49.678 KB	609	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 software (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	0.03%	0 bytes/s	4 KB/s		
cvloc(8:2,sda2)	N/A	0.37%	0 bytes/s	29 KB/s		
sd(8:16,sdb)	N/A	0.00%	0 bytes/s	0 bytes/s		
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	183 bytes/s		
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	12 bytes/s		

Volumes						
Mount Point	Device	Status	Size	Available	Write Cache Status	
/	croot	Online	10.50 GB	3.46 GB	Unknown	
/var/local	cvloc	Online	96.59 GB	94.99 GB	Unknown	
/var/local/rangedb/0	sdc	Online	53.66 GB	53.57 GB	Enabled	
/var/local/rangedb/1	sdd	Online	53.66 GB	53.57 GB	Enabled	
/var/local/rangedb/2	sde	Online	53.66 GB	53.57 GB	Enabled	

Object Stores						
ID	Size	Available	Object Data	Object Data (%)	Health	
0000	53.66 GB	48.21 GB	976.25 KB	0.00%	No Errors	
0001	53.66 GB	53.57 GB	0 bytes	0.00%	No Errors	
0002	53.66 GB	53.57 GB	0 bytes	0.00%	No Errors	

Related information

- [SG6000 appliance installation and maintenance](#)
- [SG5700 appliance installation and maintenance](#)
- [SG5600 appliance installation and maintenance](#)

Viewing 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 for 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**.

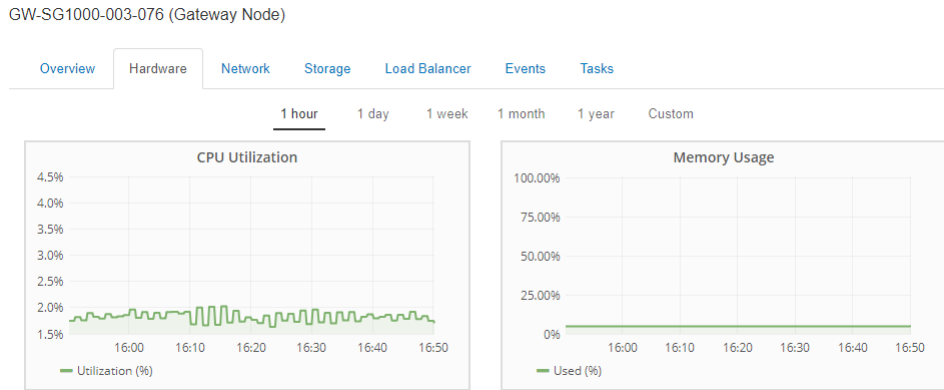
The Node Information table on the Overview tab displays the node's ID and name, the node type, the software version installed, and the IP addresses associated with the node. The Interface column contains the name of the interface, as follows:

- **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, which can be bonded or aliased and connected to the StorageGRID Admin Network (eth1).

Node Information ?	
ID	46702fe0-2bca-4097-8f61-f3fe6b22ed75
Name	GW-SG1000-003-076
Type	Gateway Node
Software Version	11.3.0 (build 20190708.2304.71ba19a)
IP Addresses	169.254.0.1, 172.16.3.76, 10.224.3.76, 47.47.3.76 Show less ▲
Interface	IP Address
adllb	fe80::c020:17ff:fe59:1cf3
adlli	169.254.0.1
adlli	fd20:327:327:0:408f:84ff:fe80:a9
adlli	fd20:8b1e:b255:8154:408f:84ff:fe80:a9
adlli	fe80::408f:84ff:fe80:a9
eth0	172.16.3.76
eth0	fd20:328:328:0:9a03:9bff:fe98:a272
eth0	fe80::9a03:9bff:fe98:a272
eth1	10.224.3.76
eth1	fd20:327:327:0:b6a9:fcff:fe08:4e49
eth1	fd20:8b1e:b255:8154:b6a9:fcff:fe08:4e49
eth1	fe80::b6a9:fcff:fe08:4e49
eth2	47.47.3.76
eth2	fd20:332:332:0:9a03:9bff:fe98:a272
eth2	fe80::9a03:9bff:fe98:a272
hic1	47.47.3.76
hic2	47.47.3.76
hic3	47.47.3.76
hic4	47.47.3.76
mtc1	10.224.3.76
mtc2	10.224.3.76

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, 1 month, or 1 year. 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, and the status of each component.

StorageGRID Appliance		
Appliance Model	SG1000	
Storage Controller Failed Drive Count	0	
Storage Data Drive Type	SSD	
Storage Data Drive Size	960.20 GB	
Storage RAID Mode	RAID1 [healthy]	
Storage Connectivity	Nominal	
Overall Power Supply	Nominal	
Compute Controller BMC IP	10.224.3.95	
Compute Controller Serial Number	721911500171	
Compute Hardware	Nominal	
Compute Controller CPU Temperature	Nominal	
Compute Controller Chassis Temperature	Nominal	


Field in the Appliance table	Description
Appliance Model	The model number for this StorageGRID appliance.
Storage Controller Failed Drive Count	The number of drives that are not optimal.
Storage Data Drive Type	The type of drives in the appliance, such as HDD (hard disk drive) or SSD (solid state drive).
Storage Data Drive Size	The total capacity including all data drives in the appliance.
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.
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.

Field in the Appliance table	Description
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 Negotiate	Link Status
adlib	C2:20:17:59:1C:F3	10 Gigabit	Full	Off	Up
adlli	42:8F:84:80:00:A9	10 Gigabit	Full	Off	Up
eth0	98:03:9B:98:A2:72	400 Gigabit	Full	Off	Up
eth1	B4:A9:FC:08:4E:49	10 Gigabit	Full	Off	Up
eth2	98:03:9B:98:A2:72	400 Gigabit	Full	Off	Up
hic1	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
hic2	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
hic3	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
hic4	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
mtc1	B4:A9:FC:08:4E:49	Gigabit	Full	On	Up
mtc2	B4:A9:FC:08:4E:49	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.

Note: 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	3.250 TB	5,610,578,144	0	8,327	0	0
eth1	1.205 GB	9,828,095	0	32,049	0	0
eth2	849.829 GB	186,349,407	0	10,269	0	0
hic1	114.864 GB	303,443,393	0	0	0	0
hic2	2.315 TB	5,351,180,956	0	305	0	0
hic3	1.690 TB	1,793,580,230	0	0	0	0
hic4	194.283 GB	331,640,075	0	0	0	0
mtc1	1.205 GB	9,828,096	0	0	0	0
mtc2	1.168 GB	9,564,173	0	32,050	0	0

Transmit						
Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	5.759 TB	5,789,638,626	0	0	0	0
eth1	4.563 MB	41,520	0	0	0	0
eth2	855.404 GB	139,975,194	0	0	0	0
hic1	289.248 GB	326,321,151	5	0	0	5
hic2	1.636 TB	2,640,416,419	18	0	0	18
hic3	3.219 TB	4,571,516,003	33	0	0	33
hic4	1.687 TB	1,658,180,262	22	0	0	22
mtc1	4.563 MB	41,520	0	0	0	0
mtc2	49.678 KB	609	0	0	0	0

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

GW-SG1000-003-076 (Gateway Node)

Overview Hardware Network **Storage** Load Balancer Events Tasks

Disk Devices

Name	World Wide Name	I/O Load	Read Rate	Write Rate
croot(253:2,dm-2)	N/A	0.00%	0 bytes/s	8 KB/s
cvloc(253:3,dm-3)	N/A	0.01%	0 bytes/s	405 KB/s

Volumes

Mount Point	Device	Status	Size	Available	Write Cache Status
/	croot	Online	21.00 GB	13.09 GB	Unknown
/var/local	cvloc	Online	903.78 GB	894.55 GB	Unknown

Related information

[SG100 and SG1000 appliance installation and maintenance](#)

Information you should monitor regularly

StorageGRID is a fault-tolerant, distributed storage system that is designed to continue operating even when errors occur, or when nodes or sites are unavailable. You must proactively monitor system health, workloads, and usage statistics so that you can take action to address potential issues before they affect the grid's efficiency or availability.

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

What to monitor	Frequency
The system health data shown on the Grid Manager Dashboard Note if anything has changed from the previous day.	Daily
Rate at which Storage Node object and metadata capacity is being consumed	Weekly
Information lifecycle management operations	Weekly
Performance, networking, and system resources: <ul style="list-style-type: none">• Query latency• Connectivity and networking• Node-level resources	Weekly
Tenant activity	Weekly
Capacity of the external archival storage system	Weekly
Load balancing operations	After the initial configuration and after any configuration changes
Availability of software hotfixes and software upgrades	Monthly

Choices

- [Monitoring system health](#) on page 36
- [Monitoring storage capacity](#) on page 48
- [Monitoring information lifecycle management](#) on page 55
- [Monitoring performance, networking, and system resources](#) on page 57
- [Monitoring tenant activity](#) on page 62
- [Monitoring archival capacity](#) on page 64
- [Monitoring load balancing operations](#) on page 65
- [Applying hotfixes or upgrading software if necessary](#) on page 66

Monitoring system health

You should monitor the overall health of your StorageGRID system on a daily basis.

About this task

The StorageGRID system is fault tolerant and can continue to operate even when parts of the grid are unavailable. The first sign of a potential issue with your StorageGRID system is likely to be an alert or an alarm (legacy system) and not necessarily an issue with system operations. Paying attention to system health can help you detect minor issues before they affect operations or grid efficiency.

The Health panel on the Grid Manager Dashboard provides a summary of any issues that might be affecting your system. You should investigate any issues that are shown on the Dashboard.

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

Steps

1. Sign in to the Grid Manager to view the Dashboard.
2. Review the information in the Health panel.



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

Link	Indicates
Grid details	Appears if any nodes are disconnected (connection state Unknown or Administratively Down). Click the link, or click the blue or gray icon to determine which node or nodes are affected.
Current alerts	Appears if any alerts are currently active. Click the link, or click Critical , Major , or Minor to see the details on the Alerts > Current page.
Recently resolved alerts	Appears if any alerts triggered in the past week are now resolved. Click the link to see the details on the Alerts > Resolved page.
Legacy alarms	Appears if any alarms (legacy system) are currently active. Click the link to see the details on the Support > Current Alarms page. Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.
License	Appears if there is an issue with the software license for this StorageGRID system. Click the link to see the details on the Maintenance > License page.

Related concepts

[Using SNMP monitoring](#) on page 110

If you want to monitor StorageGRID using the Simple Network Management Protocol (SNMP), you must configure the SNMP agent that is included with StorageGRID.

Related tasks

[Setting up email notifications for alerts](#) on page 82

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.

Related information

[Administering StorageGRID](#)

Choices

- [Monitoring node connection states](#) on page 38
- [Viewing current alerts](#) on page 39

- [Viewing resolved alerts](#) on page 41
- [Viewing a specific alert](#) on page 43
- [Viewing legacy alarms](#) on page 46

Monitoring node connection states


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

Before you begin



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

About this task

Nodes can have one of three connection states:

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

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

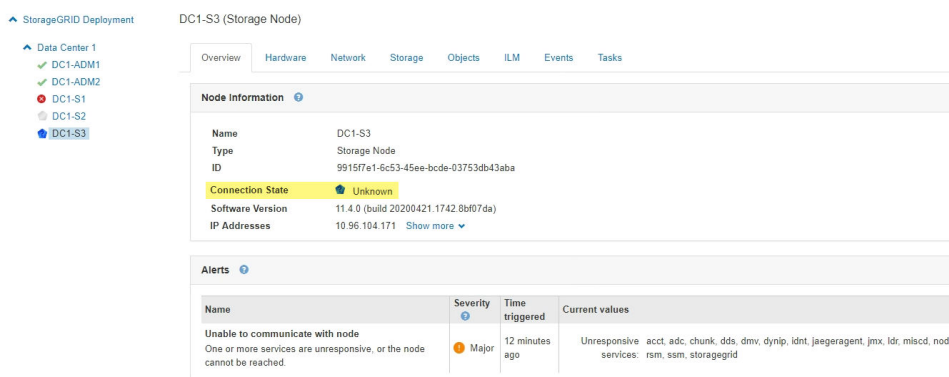
- **Not connected - Administratively down** : The node is not connected to the grid for an expected reason. For example, the node, or services on the node, has been gracefully shut down, the node is rebooting, or the software is being upgraded. One or more alerts might also be active.
- **Connected**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

Steps

1. If a blue or gray icon appears on the Health panel of the Dashboard, click the icon or click **Grid details**. (The blue or gray icons and the **Grid details** link appear only if at least one node is disconnected from the grid.)

The Overview page for the first blue node in the node tree appears. If there are no blue nodes, the Overview page for the first gray node in the tree appears.

In the example, the Storage Node named DC1-S3 has a blue icon. The **Connection State** on the Node Information panel is **Unknown**, and the **Unable to communicate with node** alert is active. The alert indicates that one or more services are unresponsive, or the node cannot be reached.



The screenshot shows the 'StorageGRID Deployment' page for 'DC1-S3 (Storage Node)'. On the left, a tree view shows nodes DC1-ADM1, DC1-ADM2, DC1-S1, DC1-S2, and DC1-S3. DC1-S3 is highlighted with a blue icon. The main panel shows 'Node Information' with fields for Name (DC1-S3), Type (Storage Node), ID (9915f7e1-6c53-45ee-bcde-03753db43aba), Connection State (Unknown), Software Version (11.4.0), and IP Addresses (10.96.104.171). Below this is an 'Alerts' section with a table:

Name	Severity	Time triggered	Current values
Unable to communicate with node One or more services are unresponsive, or the node cannot be reached.	Major	12 minutes ago	Unresponsive acct, adc, chunk, dds, dmv, dynip, idnt, jaegeragent, jmx, ldr, miscd, node, services: rsm, ssm, storagegrid

2. If a node has a blue icon, follow these steps:
 - a. Select each alert in the table, and follow the recommended actions.

For example, you might need to restart a service that has stopped or restart the host for the node.

- b. If you are unable to bring the node back online, contact technical support.
- 3. If a node has a gray icon, follow these steps:

Gray nodes are expected during maintenance procedures and might be associated with one or more alerts. Based on the underlying issue, these "administratively down" nodes often go back online with no intervention.

- a. Review the Alerts section, and determine if any alerts are affecting this node.
- b. If one or more alerts are active, select each alert in the table, and follow the recommended actions.
- c. If you are unable to bring the node back online, contact technical support.

Related reference

[Alerts reference](#) on page 201

The following table lists all default StorageGRID alerts. As required, you can create custom alert rules to fit your system management approach.

Related information

[Recovery and maintenance](#)

Viewing 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. An email notification might also be sent, unless the alert has been silenced.

Before you begin

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

Steps

1. If one or more alerts are active, do either of the following:
 - From the Health panel on the Dashboard, click the alert icon or click **Current alerts**. (An alert icon and the **Current alerts** link appear only if at least one alert is currently active.)
 - Select **Alerts > Current**.

The Current Alerts page appears. It lists all alerts currently affecting your StorageGRID system.

Current Alerts
 View the current alerts affecting your StorageGRID system.

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




By default, 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 are not 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. Review the information in the table.


Column header	Description
Name	The name of the alert and its description.
Severity	<p>The severity of the alert. If multiple alerts are grouped, the title row shows how many instances of that alert are occurring at each severity.</p> <ul style="list-style-type: none"> • Critical : An abnormal condition exists that has stopped the normal operations of a StorageGRID node or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved. • Major : An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service. • Minor : The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that do not clear on their own to ensure they do not result in a more serious problem.
Time triggered	How long ago the alert was triggered. If multiple alerts are grouped, the title row shows times for the most recent instance of the alert (<i>newest</i>) and the oldest instance of the alert (<i>oldest</i>).
Site/Node	The name of the site and node where the alert is occurring. If multiple alerts are grouped, the site and node names are not shown in the title row.
Status	Whether the alert is active or has been silenced. 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.
Current values	<p>The current 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.</p> <p>Note: If multiple alerts are grouped, current values are not shown in the title row.</p>

3. To expand and collapse groups of alerts:


- To show the individual alerts in a group, click the down caret  in the heading, or click the group's name.
- To hide the individual alerts in a group, click the up caret  in the heading, or click the group's name.

Name	Severity	Time triggered	Site / Node	Status	Current values
<input checked="" type="checkbox"/> Group alerts Active ▾					
^ Low object data storage The disk space available for storing object data is low.	▲ 5 Minor	a day ago (newest) a day ago (oldest)		5 Active	
Low object data storage The disk space available for storing object data is low.	▲ Minor	a day ago	DC2 231-236 / DC2-S2-233	Active	Disk space remaining: 525.17 GB Disk space used: 243.06 KB Disk space used (%): 0.000%
Low object data storage The disk space available for storing object data is low.	▲ Minor	a day ago	DC1 225-230 / DC1-S1-226	Active	Disk space remaining: 525.17 GB Disk space used: 325.65 KB Disk space used (%): 0.000%
Low object data storage The disk space available for storing object data is low.	▲ Minor	a day ago	DC2 231-236 / DC2-S3-234	Active	Disk space remaining: 525.17 GB Disk space used: 381.55 KB Disk space used (%): 0.000%
Low object data storage The disk space available for storing object data is low.	▲ Minor	a day ago	DC1 225-230 / DC1-S2-227	Active	Disk space remaining: 525.17 GB Disk space used: 282.19 KB Disk space used (%): 0.000%
Low object data storage The disk space available for storing object data is low.	▲ Minor	a day ago	DC2 231-236 / DC2-S1-232	Active	Disk space remaining: 525.17 GB Disk space used: 189.24 KB Disk space used (%): 0.000%

4. To display individual alerts instead of groups of alerts, unselect the **Group alerts** check box at the top of the table.



5. To sort alerts or alert groups, click 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 unselected, 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.
6. To filter the alerts by status, use the drop-down menu at the top of the table.



- Select **All alerts** to view all current alerts (both active and silenced alerts).
 - Select **Active** to view only the current alerts that are active.
 - Select **Silenced** to view only the current alerts that have been silenced.
7. To view details for a specific alert, select the alert from the table.
 A dialog box for the alert appears. See the instructions for viewing a specific alert.

Related tasks

[Viewing a specific alert](#) on page 43

You can view detailed information about an alert that is currently affecting your StorageGRID system or an alert that has been resolved. The details include recommended corrective actions, the time the alert was triggered, and the current value of the metrics related to this alert. Optionally, you can silence a current alert or update the alert rule.

[Silencing alert notifications](#) on page 87

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

Viewing resolved alerts

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

Before you begin

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

Steps

1. To view resolved alerts, do either of the following:

- From the Health panel on the Dashboard, click **Recently resolved alerts**.
 The **Recently resolved alerts** link appears only if one or more alerts were triggered in the past week and are now resolved.
- Select **Alerts > Resolved**.

The Resolved Alerts page appears. By default, resolved alerts that were triggered in the last week are shown, with the most recently triggered alerts shown first. The alerts on this page were previously shown on the Current Alerts page or in an email notification.

Resolved Alerts
 Search and view alerts that have been resolved.




When triggered ✕ Severity ✕ Alert rule ✕ Node ✕


Last week Filter by severity Filter by rule Filter by node Search

Name	Severity	Time triggered	Time resolved	Site / Node	Triggered values
Low installed node memory The amount of installed memory on a node is low.	Critical	2 days ago	a day ago	Data Center 1 / DC1-S2	Total RAM size: 8.37 GB
Low installed node memory The amount of installed memory on a node is low.	Critical	2 days ago	a day ago	Data Center 1 / DC1-S3	Total RAM size: 8.37 GB
Low installed node memory The amount of installed memory on a node is low.	Critical	2 days ago	a day ago	Data Center 1 / DC1-S4	Total RAM size: 8.37 GB
Low installed node memory The amount of installed memory on a node is low.	Critical	2 days ago	a day ago	Data Center 1 / DC1-ADM1	Total RAM size: 8.37 GB
Low installed node memory The amount of installed memory on a node is low.	Critical	2 days ago	a day ago	Data Center 1 / DC1-ADM2	Total RAM size: 8.37 GB
Low installed node memory The amount of installed memory on a node is low.	Critical	2 days ago	a day ago	Data Center 1 / DC1-S1	Total RAM size: 8.37 GB

Page 1 of 1

2. Review the information in the table.

Column header	Description
Name	The name of the alert and its description.
Severity	The severity of the alert. <ul style="list-style-type: none"> Critical : An abnormal condition exists that has stopped the normal operations of a StorageGRID node or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved. Major : An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service. Minor : The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that do not clear on their own to ensure they do not result in a more serious problem.
Time triggered	How long ago the alert was triggered.
Time resolved	How long ago the alert was resolved.
Site/Node	The name of the site and node where the alert occurred.
Triggered 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.

3. To sort the entire list of resolved alerts, click the up/down arrows  in each column header.
For example, you might want to sort resolved alerts by **Site/Node** to see the alerts that affected a specific node.
4. Optionally, filter the list of resolved alerts by using the drop-down menus at the top of the table.
 - a. Select a time period from the **When triggered** drop-down menu to show resolved alerts based on how long ago they were triggered.
You can search for alerts that were triggered within the following time periods:
 - Last hour
 - Last day
 - Last week (default view)
 - Last month
 - Any time period
 - Custom (allows you to specify the start date and the end date for the time period)
 - b. Select one or more severities from the **Severity** drop-down menu to filter on resolved alerts of a specific severity.
 - c. 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.
 - d. Select one or more nodes from the **Node** drop-down menu to filter on resolved alerts related to a specific node.
 - e. Click **Search**.
5. To view details for a specific resolved alert, select the alert from the table.
A dialog box for the alert appears. See the instructions for viewing a specific alert.

Related tasks

[Viewing a specific alert](#) on page 43

You can view detailed information about an alert that is currently affecting your StorageGRID system or an alert that has been resolved. The details include recommended corrective actions, the time the alert was triggered, and the current value of the metrics related to this alert. Optionally, you can silence a current alert or update the alert rule.

Viewing a specific alert

You can view detailed information about an alert that is currently affecting your StorageGRID system or an alert that has been resolved. The details include recommended corrective actions, the time the alert was triggered, and the current value of the metrics related to this alert. Optionally, you can silence a current alert or update the alert rule.

Before you begin

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

Steps

1. Do one of the following, based on whether you want to view a current or resolved alert:

Option	Description
Current alert	<ul style="list-style-type: none">• From the Health panel on the Dashboard, click the Current alerts link. This link appears only if at least one alert is currently active. This link is hidden if there are no current alerts or if all current alerts have been silenced.• Select Alerts > Current.• From the Nodes page, select the Overview tab for a node that has an alert icon. Then, in the Alerts section, click the alert name.

Option	Description
Resolved alert	<ul style="list-style-type: none"> From the Health panel on the Dashboard, click the Recently resolved alerts link. (This link appears only if one or more alerts were triggered in the past week and are now resolved. This link is hidden if no alerts were triggered and resolved in the last week.) Select Alerts > Resolved.

2. As required, expand a group of alerts and then select the alert you want to view.

Note: Select the alert, not the heading for a group of alerts.

^ Low installed node memory The amount of installed memory on a node is low.		8 Critical	a day ago (newest) a day ago (oldest)	8 Active
Low installed node memory The amount of installed memory on a node is low.	8 Critical	a day ago	Data Center 2 / DC2-S1-99-56	Active Total RAM size: 8.38 GB

A dialog box appears and provides details for the selected alert.

Low installed node memory

The amount of installed memory on a node is low.

Recommended actions

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

See the instructions for your platform:

- [VMware installation](#)
- [Red Hat Enterprise Linux or CentOS installation](#)
- [Ubuntu or Debian installation](#)

Status
Active [\(silence this alert\)](#)

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

Severity
8 Critical

Total RAM size
8.38 GB




Condition
[View conditions](#) | [Edit rule](#)

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

[Close](#)

3. Review the alert details.

Information	Description
<i>title</i>	The name of the alert.
<i>first paragraph</i>	The description of the alert.
Recommended actions	The recommended actions for this alert.
Time triggered	The date and time the alert was triggered in your local time and in UTC.
Time resolved	For resolved alerts only, the date and time the alert was resolved in your local time and in UTC.
Status	The status of the alert: Active, Silenced, or Resolved.
Site/Node	The name of the site and node affected by the alert.

Information	Description
Severity	The severity of the alert. <ul style="list-style-type: none"> • Critical : An abnormal condition exists that has stopped the normal operations of a StorageGRID node or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved. • Major : An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service. • Minor : The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that do not clear on their own to ensure they do not result in a more serious problem.
<i>data values</i>	The current value of the metric for this alert. For some alerts, additional values are shown to help you understand and investigate the alert. For example, the values shown for a Low metadata storage alert include the percent of disk space used, the total amount of disk space, and the amount of disk space used.

- Optionally, click **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.



Attention: 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.

- To view the current conditions for the alert rule:

- From the alert details, click **View conditions**.

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



- To close the pop-up, click anywhere outside of the pop-up.

- Optionally, click **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.



Attention: 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.

- To close the alert details, click **Close**.

Related tasks

[Silencing alert notifications](#) on page 87

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

[Editing an alert rule](#) on page 78

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.

Viewing legacy alarms

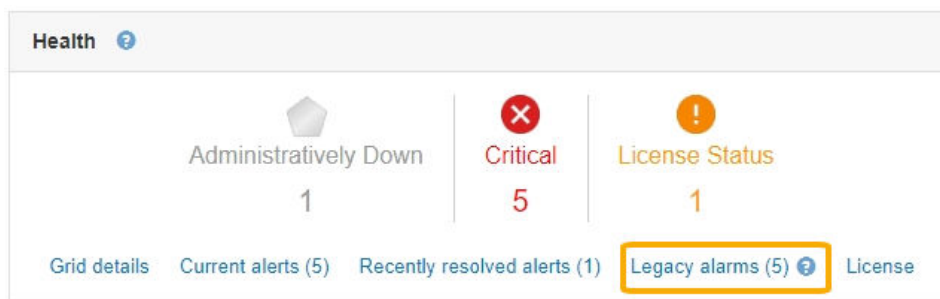
Alarms (legacy system) are triggered when system attributes reach alarm threshold values. You can view the currently active alarms from the Dashboard or the Current Alarms page.

Before you begin

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

About this task

If one or more of the legacy alarms are currently active, the Health panel on the Dashboard includes a **Legacy alarms** link. The number in parentheses indicates how many alarms are currently active.



Because the legacy alarm system continues to be supported in StorageGRID 11.4, the number of alarms shown on the Dashboard is incremented whenever a new alarm occurs. This count is incremented even if you have disabled alarm email notifications. You can typically ignore this number (since alerts provide a better view of the system), or you can view the alarms that are currently active.

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.

Steps

- To view the legacy alarms that are currently active, do one of the following:
 - From the Health panel on the Dashboard, click **Legacy alarms**. This link appears only if at least one alarm is currently active.
 - Select **Support**. Then, in the Alarms (legacy) section of the menu, select **Current Alarms**.

The Current Alarms page appears.

The alarm system is the legacy system. The alert system offers significant benefits and is easier to use. See [Managing alerts and alarms](#) in the instructions for monitoring and troubleshooting StorageGRID.

Current Alarms





Last Refreshed: 2020-05-27 09:41:39 MDT

Show Acknowledged Alarms (1 - 1 of 1)

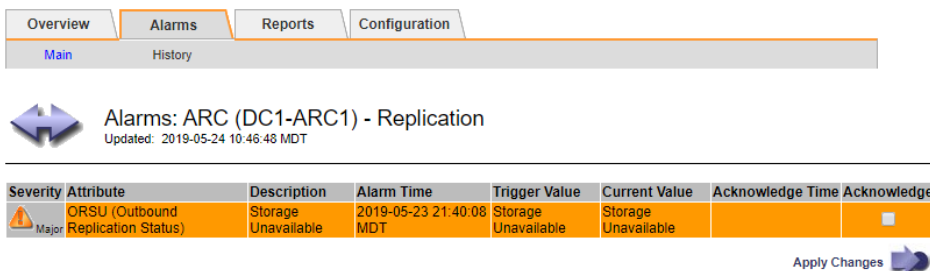
Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value
Major	ORSU (Outbound Replication Status)	Data Center 1/DC1-ARC1/ARC	Storage Unavailable	2020-05-26 21:47:18 MDT	Storage Unavailable	Storage Unavailable

Show Records Per Page Previous « 1 » Next

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.
	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.

- To learn about the attribute that caused the alarm to be triggered, right click the attribute name in the table.
- To view additional details about an alarm, click the service name in the table.
 The Alarms tab for the selected service appears (**Support > Grid Topology > Grid Node > Service > Alarms**).



Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Major	ORSU (Outbound Replication Status)	Storage Unavailable	2019-05-23 21:40:08 MDT	Storage Unavailable	Storage Unavailable		<input type="checkbox"/>

- 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 concepts

[Disabling alarms \(legacy system\)](#) on page 99

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

Related tasks

[Acknowledging current alarms \(legacy system\)](#) on page 94

Legacy alarms are triggered when system attributes reach alarm threshold values. If you want to reduce or clear the count of legacy alarms on the Dashboard, you can acknowledge the alarms.

Related reference

[Alarms reference \(legacy system\)](#) on page 222

Monitoring storage capacity

You must monitor the total usable space available on Storage Nodes 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.

Related concepts

[Viewing the Storage tab](#) on page 14

The Storage tab summarizes storage availability and other storage metrics.

Choices

- [Monitoring storage capacity for the entire grid](#) on page 48
- [Monitoring storage capacity for each Storage Node](#) on page 50
- [Monitoring object metadata capacity for each Storage Node](#) on page 53

Monitoring storage capacity for the entire grid

You must 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.

Before you begin

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

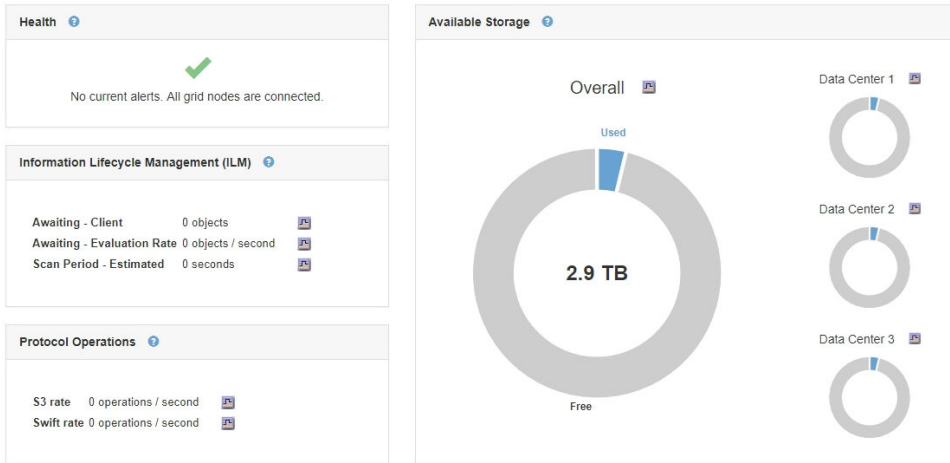
About this task

The Dashboard in the Grid Manager 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**.
 - b. In the **Available Storage** panel, note the overall summary of free and used storage capacity.

Note: The summary does not include archival media.



- c. Place your cursor over the chart's Free or Used capacity sections to see exactly how much space is free or used.

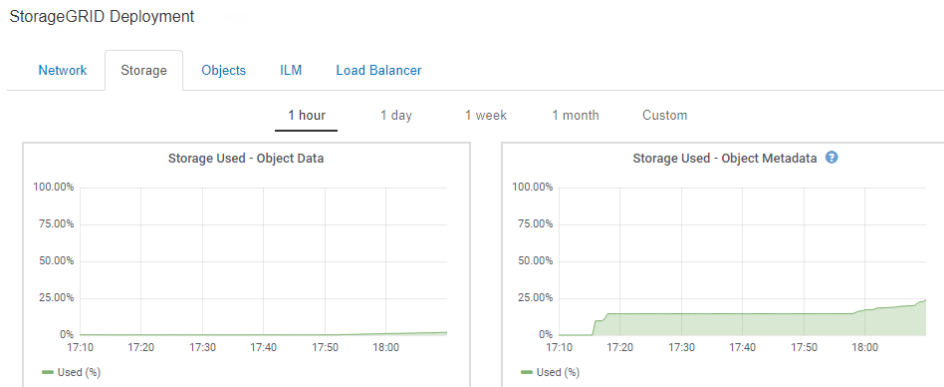


- d. For multi-site grids, review the chart for each data center.
- e. Click the chart icon  for the overall chart or for an individual data center to view a graph showing capacity usage over time.

A graph showing Percentage Storage Capacity Used (%) vs. Time appears.

- 2. Determine how much storage has been used and how much storage remains available for object data and object metadata.

- a. Select **Nodes**.
- b. Select *grid* > **Storage**.

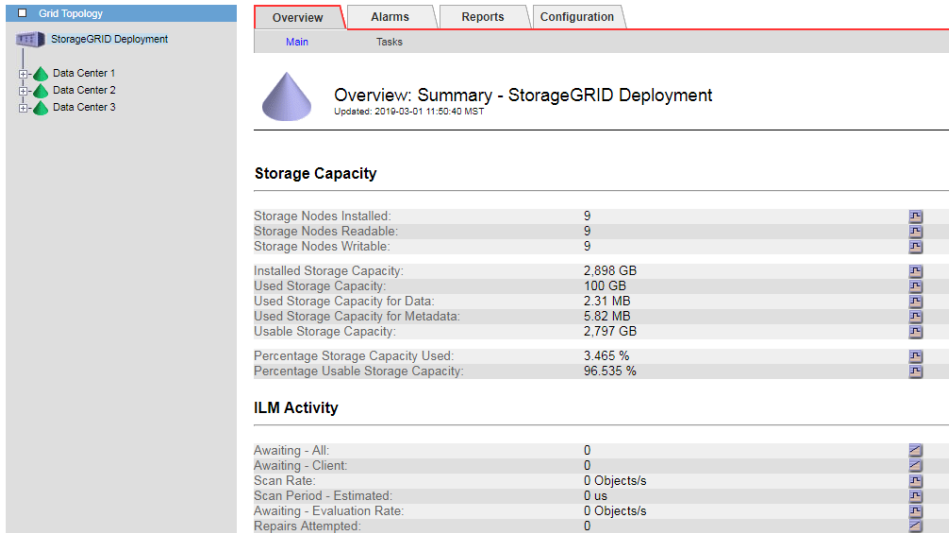


- c. Hover your cursor over the Storage Used - Object Data and the Storage Used - Object Metadata charts to see how much object storage and object metadata storage is available for the entire grid, and how much has been used over time.

Note: The total values for a site or the grid do not include nodes that not have reported metrics for at least five minutes, such as offline nodes.

- 3. As directed by technical support, view additional details about the storage capacity for your grid.

- a. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
- b. Select **grid > Overview > Main**.



- 4. 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.

Note: 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 on planning a storage expansion, see the instructions for expanding StorageGRID.

Related information

[Expanding a StorageGRID system](#)

Monitoring storage capacity for each Storage Node

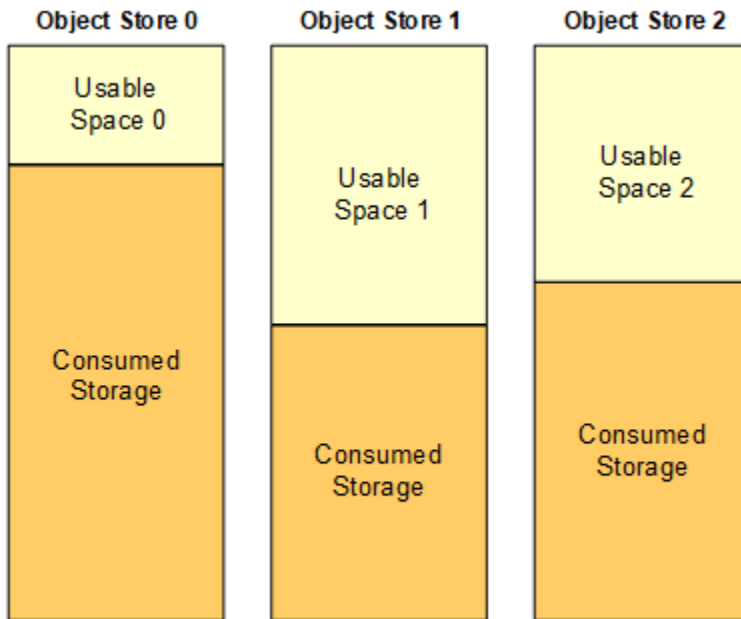
You must monitor the total usable space for each Storage Node to ensure that the node has enough space for new object data.

Before you begin

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

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.



Total Usable Space = Usable Space 0 + Usable Space 1 + Usable Space 2


Steps

1. Select **Nodes > Storage Node > Storage**.


The graphs and tables for the node appear.

2. Hover your cursor over the **Storage Used - Object Data** graph.

A pop-up displays Used (%), Used, and Total values for the Storage Node. The Total value is the sum of the available space on all object stores within the node (STAS attribute).



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

Note: 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	0.03%		0 bytes/s		4 KB/s
cvloc(8:2,sda2)	N/A	0.37%		0 bytes/s		29 KB/s
sdc(8:16,sdb)	N/A	0.00%		0 bytes/s		0 bytes/s
sdd(8:32,sdc)	N/A	0.00%		0 bytes/s		183 bytes/s
sde(8:48,sdd)	N/A	0.00%		0 bytes/s		12 bytes/s

Volumes						
Mount Point	Device	Status	Size	Available		Write Cache Status
/	croot	Online	10.50 GB	3.46 GB		Unknown
/var/local	cvloc	Online	96.59 GB	94.99 GB		Unknown
/var/local/rangedb/0	sdc	Online	53.66 GB	53.57 GB		Enabled
/var/local/rangedb/1	sdd	Online	53.66 GB	53.57 GB		Enabled
/var/local/rangedb/2	sde	Online	53.66 GB	53.57 GB		Enabled

Object Stores							
ID	Size	Available		Object Data		Object Data (%)	Health
0000	53.66 GB	48.21 GB		976.25 KB		0.00%	No Errors
0001	53.66 GB	53.57 GB		0 bytes		0.00%	No Errors
0002	53.66 GB	53.57 GB		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.

Note: 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 on planning a storage expansion, see the instructions for expanding StorageGRID.

The **Low object data storage** alert and the legacy Storage Status (SSTS) alarm are triggered when insufficient space remains for storing object data on a Storage Node.

Related tasks

[Troubleshooting the Low object data storage alert](#) on page 174

The **Low object data storage** alert monitors how much space is available for storing object data on each Storage Node. It is related to the Storage Status (SSTS) legacy alarm, but it is not exactly equivalent.

Related information

[Administering StorageGRID](#)

[Expanding a StorageGRID system](#)

Monitoring object metadata capacity for each Storage Node

You must 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.

Before you begin

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

About this task

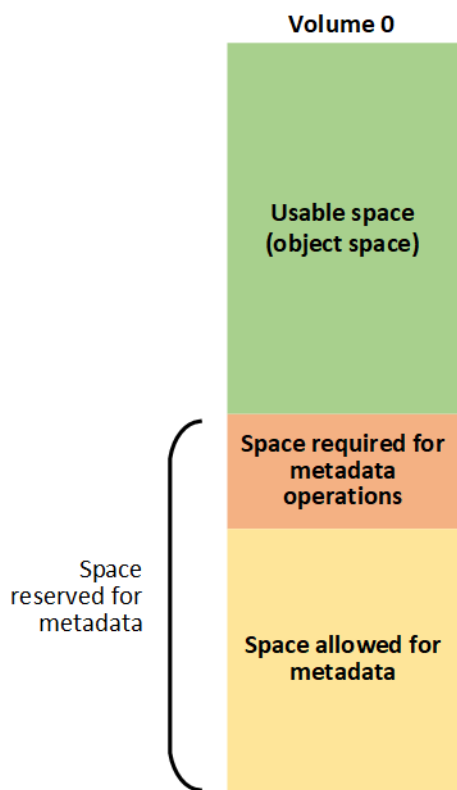
StorageGRID maintains object metadata for each object or object version that it stores. The metadata for an object or object version can include the following types of information:

- System metadata, including the object ID (UUID), the object name, the name of the container, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
- Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object.
- For replicated object copies, the current storage location of each copy.
- For erasure-coded object copies, the current storage location of each fragment.
- For object copies in a Cloud Storage Pool, the location of the object, including the name of the external bucket and the object's unique identifier.
- For segmented objects and multipart objects, segment identifiers and data sizes.

Three copies of object metadata are maintained 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 space reserved on storage volume 0 of each Storage Node.

The total space reserved for metadata on volume 0 of each Storage Node is a system-wide setting known as the Metadata Reserved Space. (To see this value, go to **Configuration > Storage Options > Overview**.)

The Metadata Reserved Space is subdivided into the space available for object metadata (the Metadata Allowed Space) and the space required for essential database operations, such as compaction and repair.



StorageGRID uses the following Prometheus expression to measure how full the Metadata Allowed Space is:

```
storagegrid_storage_utilization_metadata_bytes/  
storagegrid_storage_utilization_metadata_allowed_bytes
```

When this Prometheus expression reaches certain thresholds, the **Low metadata storage** alert is triggered.

Note: This Prometheus expression is equivalent to the legacy CDLP (Metadata Used Space (Percent)) attribute, which triggers the legacy CDLP alarm at the same thresholds.

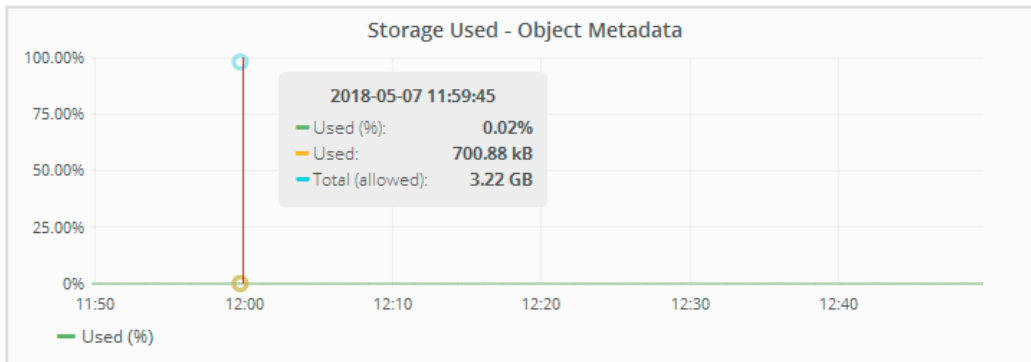
Some of the factors that can increase metadata usage include the size and quantity of user metadata and tags, the total number of parts in a multipart upload, and the frequency of changes to ILM storage locations.

To ensure that adequate space remains for object metadata, follow these steps to monitor the metadata capacity of a Storage Node.

Steps

1. Select **Nodes** > **Storage Node** > **Storage**.
2. Hover your cursor over the **Storage Used - Object Metadata** graph to see the percentage of allowed space consumed by object metadata.

The value for Used % is the current value of the Prometheus expression shown above.



3. If the **Used %** value is 70% or higher, expand the StorageGRID system by adding Storage Nodes.

When you add the new nodes, the system automatically rebalances object metadata across all Storage Nodes within the site.

Related tasks

[Troubleshooting the Low metadata storage alert](#) on page 183

If the **Low metadata storage** alert or the legacy CDLP alarm is triggered, you must add new Storage Nodes.

Related information

[Administering StorageGRID](#)

[Expanding a StorageGRID system](#)

Monitoring 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.

Before you begin


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

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 cannot 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 in the Information Lifecycle Management (ILM) panel.
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. Hover 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.

Note: 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.
 - Note:** 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.

Note: 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.

Monitoring performance, networking, and system resources

You should monitor performance, networking, and system resources to determine whether StorageGRID can handle its current load and to ensure that client performance does not degrade over time.

Choices

- [Monitoring query latency](#) on page 57
- [Monitoring network connections and performance](#) on page 58
- [Monitoring node-level resources](#) on page 61

Monitoring query latency

Client actions such as storing, retrieving, or deleting objects create queries to the grid's distributed database of object metadata. You should monitor trends in query latency to ensure that grid resources are adequate for the current load.

Before you begin

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

About this task





Temporary increases in query latency are normal and can be caused by a sudden increase in ingest requests. Failed queries are also normal and can result from transient network issues or nodes that are temporarily unavailable. However, if the average time to perform a query increases, overall grid performance declines.

If you notice that query latency is increasing over time, you should consider adding additional Storage Nodes in an expansion procedure to satisfy future workloads .


The **High latency for metadata queries** alert is triggered if the average time for queries is too long.

Steps

1. Select **Nodes** > **Storage Node** > **Objects**.
2. Scroll down to the Queries table and view the value for Average Latency.

Queries	
Average Latency	1.22 milliseconds 
Queries - Successful	1,349,103,223 
Queries - Failed (timed-out)	12022 
Queries - Failed (consistency level unmet)	560925 

3. Click the chart icon  to chart the value over time.



The example chart shows spikes in query latency during normal grid operation.

Related information

[Expanding a StorageGRID system](#)

Monitoring network connections and performance

Grid nodes must be able to communicate with one another to permit the grid to operate. The integrity of the network between nodes and sites, and the network bandwidth between sites, are critical to efficient operations.

Before you begin

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

About this task

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 simply to poor ingest performance and ILM backlogs.

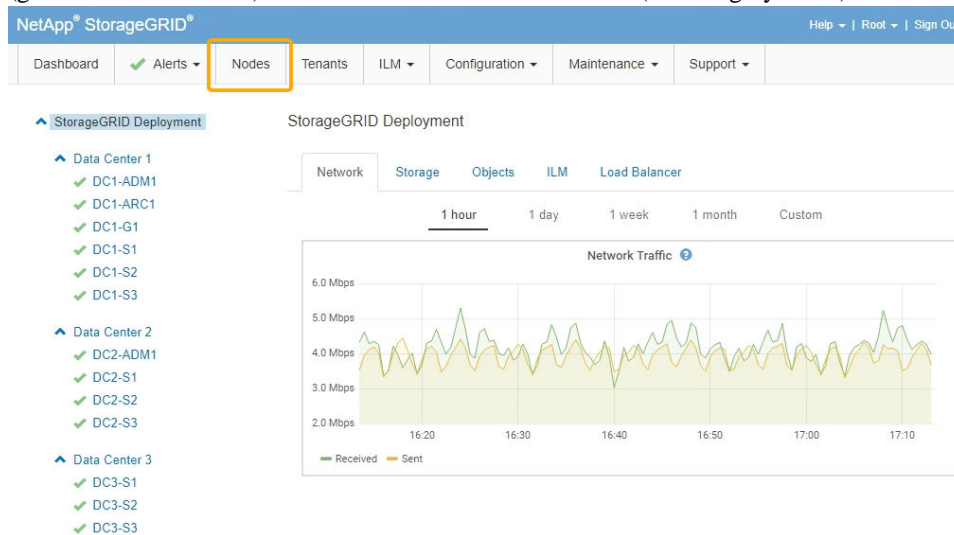
You can use the Grid Manager to monitor connectivity and network performance, so you can address any issues promptly.

Additionally, consider creating network traffic classification policies to provide monitoring and limiting for traffic related to specific tenants, buckets, subnets, or load balancer endpoints. See the instructions for administering StorageGRID.

Steps


1. Select **Nodes**.

The Nodes page appears. The node icons indicate at a glance which nodes are connected (green checkmark icon) and which nodes are disconnected (blue or gray icons).



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 Negotiate	Link Status
eth0	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
eth1	D8:C4:97:2A:E4:9E	Gigabit	Full	Off	Up
eth2	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
hic1	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic2	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic3	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic4	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
mtc1	D8:C4:97:2A:E4:9E	Gigabit	Full	On	Up
mtc2	D8:C4:97:2A:E4:9F	Gigabit	Full	On	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	3.250 TB	5,610,578,144	0	8,327	0	0
eth1	1.205 GB	9,828,095	0	32,049	0	0
eth2	849.829 GB	186,349,407	0	10,269	0	0
hic1	114.864 GB	303,443,393	0	0	0	0
hic2	2.315 TB	5,351,180,956	0	305	0	0
hic3	1.690 TB	1,793,580,230	0	0	0	0
hic4	194.283 GB	331,640,075	0	0	0	0
mtc1	1.205 GB	9,828,096	0	0	0	0
mtc2	1.168 GB	9,564,173	0	32,050	0	0

Transmit						
Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	5.759 TB	5,789,638,626	0	0	0	0
eth1	4.563 MB	41,520	0	0	0	0
eth2	855.404 GB	139,975,194	0	0	0	0
hic1	289.248 GB	326,321,151	5	0	0	5
hic2	1.636 TB	2,640,416,419	18	0	0	18
hic3	3.219 TB	4,571,516,003	33	0	0	33
hic4	1.687 TB	1,658,180,262	22	0	0	22
mtc1	4.563 MB	41,520	0	0	0	0
mtc2	49.678 KB	609	0	0	0	0

3. Use the metrics associated with your traffic classification policies to monitor network traffic.
 - a. Select **Configuration**. Then, in the System Settings section of the menu, select **Traffic Classification**.

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

Traffic Classification Policies

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

+ Create ✎ Edit ✕ Remove 📊 Metrics		
Name	Description	ID
<input type="radio"/> ERP Traffic Control	Manage ERP traffic into the grid	cd9afb7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/> Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b

Displaying 2 traffic classification policies.

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

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

To create, edit, or delete traffic classification policies, see the instructions for administering StorageGRID.

Related concepts

[Viewing the Network tab](#) on page 12

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.

Related tasks

[Monitoring node connection states](#) on page 38

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

Related information

[Administering StorageGRID](#)

Monitoring node-level resources

You should monitor individual grid nodes to check their resource utilization levels.

Before you begin

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

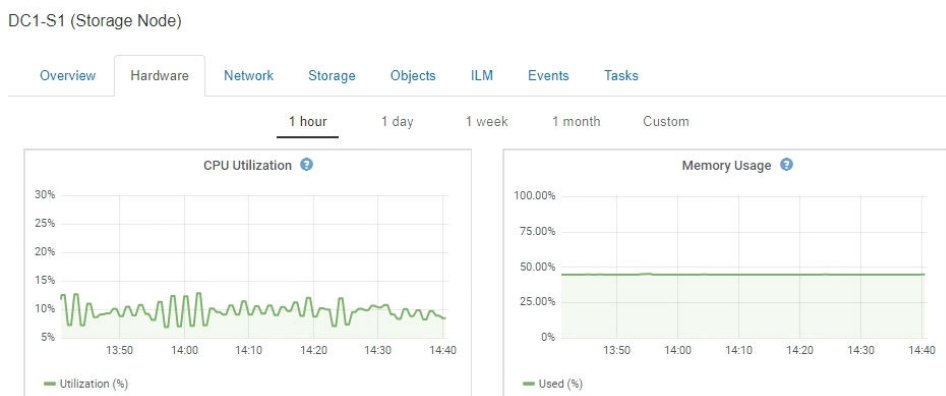
About this task

If nodes are consistently overloaded, more nodes might be required for efficient operations.

Step

To view information about hardware utilization of a grid node:

- a. From the **Nodes** page, select the node.
- b. Select the **Hardware** tab to display graphs of CPU Utilization and Memory Usage.



- c. 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, 1 month, or 1 year. You can also set a custom interval, which allows you to specify date and time ranges.

- d. 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 tasks

Viewing information about appliance Storage Nodes on page 22

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, network resources, network interfaces, network addresses, and receive and transmit data.

Viewing information about appliance Admin Nodes and Gateway Nodes on page 30

The Nodes page lists information about service health and all computational, disk device, and network resources for each services appliance that is used for 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.

Monitoring tenant activity

All client activity is associated with a tenant account. You can use the Grid Manager to monitor a tenant's storage usage or network traffic, or you can use the audit log or Grafana dashboards to gather more detailed information about how tenants are using StorageGRID.

Before you begin

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

About this task



Attention: The storage values and quotas displayed are estimates. These estimates are affected by the timing of ingests, network connectivity, and node status.



Attention: Quotas are based on estimated storage usage 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.

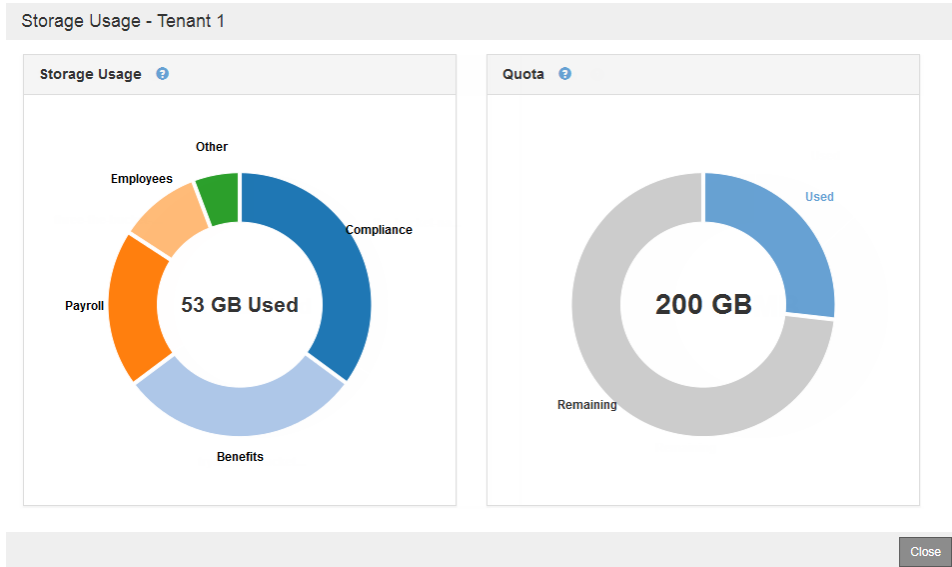
Steps

1. Review the amount of storage used by a tenant:
 - a. Select **Tenants**.
The Tenant Accounts page appears.

Tenant Accounts

+ Create Edit Account Change Root Password Remove Usage			
Display Name	ID	Protocol	Tenant sign in
<input type="radio"/> S3 tenant	68218524085409783911	S3	Sign in
<input type="radio"/> Swift tenant	29382982121425257063	Swift	Sign in
Show <input type="text" value="20"/> rows per page			◀ ▶

- b. Select the tenant whose usage you want to display, then click **Usage**.
The Storage Usage panel appears.



If a storage quota was specified for the tenant, the panel shows both the total amount of storage used and the amount of the quota that has been consumed.

2. If traffic classification policies are in place for a tenant, review the network traffic for that tenant.
 - a. Select **Configuration**. Then, in the System Settings section of the menu, select **Traffic Classification**.

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

Traffic Classification Policies

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

<input type="button" value="+ Create"/> <input type="button" value="Edit"/> <input type="button" value="Remove"/> <input type="button" value="Metrics"/>		
Name	Description	ID
<input type="radio"/> ERP Traffic Control	Manage ERP traffic into the grid	cd9afb7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/> Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b

Displaying 2 traffic classification policies.

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

To create, edit, or delete traffic classification policies, see the instructions for administering StorageGRID.

3. Optionally, 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 the instructions for understanding audit messages for more information.

4. Optionally, use Prometheus metrics to report on tenant activity:

- In the Grid Manager, select **Support**. Then, in the Tools section of the menu, select **Metrics**. You can use existing dashboards, such as S3 Overview, to review client activities.



Attention: 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.

- Select **Help > 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.

Related tasks

[Reviewing support metrics](#) on page 142

When troubleshooting an issue, you can work with technical support to review detailed metrics and charts for your StorageGRID system.

Related information

[Understanding audit messages](#)

[Administering StorageGRID](#)

Monitoring archival capacity

You cannot 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.

Before you begin

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

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**. Then, in the Tools section of the menu, select **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.

ARC State:	Online		
ARC Status:	No Errors		
Tivoli Storage Manager State:	Online		
Tivoli Storage Manager Status:	No Errors		
Store State:	Online		
Store Status:	No Errors		
Retrieve State:	Online		
Retrieve Status:	No Errors		
Inbound Replication Status:	No Errors		
Outbound Replication Status:	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.

Related information

[Administering StorageGRID](#)

Monitoring 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.

Before you begin

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

About this task

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

Note: The CLB service is deprecated.

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.

See the section on configuring client connections in the instructions for administering StorageGRID.

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 Master.

Nodes with the role of Master and nodes that are not 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 do not 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.

2. If S3 or Swift clients connect using the CLB service, perform the following checks:

- a. Select **Nodes**.
- b. Select a Gateway 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 Master.

Nodes with the role of Master and nodes that are not in an HA group should be actively distributing requests to clients.

- d. For each Gateway Node that should be actively distributing client requests, select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
- e. Select **Gateway Node > CLB > HTTP > Overview > Main**.
- f. Review the number of **Incoming Sessions - Established** to verify that the Gateway Node has been actively handling requests.

3. 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.

Related concepts

[Viewing the Load Balancer tab](#) on page 20

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

Related information

[Administering StorageGRID](#)

Applying hotfixes or upgrading 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 the NetApp Downloads page for StorageGRID.

[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 pull-down menu, and click **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 the "What's new" topic in the instructions for upgrading StorageGRID.
 - 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 StorageGRID.
 - For a hotfix, locate the hotfix procedure in the recovery and maintenance instructions

Related information

[*Upgrading StorageGRID*](#)

[*Recovery and maintenance*](#)

Managing alerts and alarms

The StorageGRID alert system is designed to inform you about operational issues that require your attention. As required, you can also use the legacy alarm system to monitor your system.

StorageGRID 11.4 includes two systems for informing you about issues.

Alert system

New for StorageGRID 11.4, 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 > 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

The alarm system is supported in StorageGRID 11.4, but is considered to be a legacy 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.

When an alarm is triggered, the following actions occur:

- The count of legacy alarms on the Dashboard is incremented.
- The alarm appears on the **Support > 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 are not sent for all alarms or alarm severities.)

Comparing alerts and alarms

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

While the alarm system continues to be supported in StorageGRID 11.4, you should transition from the legacy alarm system to the new alert system. As you begin using alerts, refer to the following table to learn how to perform similar operations.

	Alerts (new system)	Alarms (legacy system)
How do I see which alerts or alarms are active?	<ul style="list-style-type: none">• Click the Current alerts link on the Dashboard.• Click the alert on the Nodes > Overview page.• Select Alerts > Current. <p><i>Viewing current alerts</i> on page 39</p>	<ul style="list-style-type: none">• Click the Legacy alarms link on the Dashboard.• Select Support. Then, in the Alarms (legacy) section of the menu, select Current Alarms. <p><i>Viewing legacy alarms</i> on page 46</p>

	Alerts (new system)	Alarms (legacy system)
What causes an alert or an alert to be triggered?	Alerts are triggered when a Prometheus expression in an alert rule returns a specific response or threshold value. <i>Viewing alert rules</i> on page 74	Alarms are triggered when a StorageGRID attribute reaches a threshold value. <i>Alarm triggering logic (legacy system)</i> on page 91
If an alert or alarm is triggered, how do I resolve the underlying problem?	The recommended actions for an alert are included in email notifications and are available from the Alerts pages in the Grid Manager. As required, additional information is provided in the StorageGRID documentation. <i>Alerts reference</i> on page 201	You can learn about an alarm by clicking the attribute name, or you can search for an alarm code in the StorageGRID documentation. <i>Alarms reference (legacy system)</i> on page 222
Where can I see a list of alerts or alarms have been resolved?	<ul style="list-style-type: none"> Click the Recently resolved alerts link on the Dashboard. Select Alerts > Resolved. <i>Viewing resolved alerts</i> on page 41	Select Support . Then, in the Alarms (legacy) section of the menu, select Historical Alarms . <i>Reviewing historical alarms and alarm frequency (legacy system)</i> on page 96
Where do I manage the settings?	Select Alerts . Then, use the options in the Alerts menu. <i>Managing alerts</i> on page 72	Select Support . Then, use the options in the Alarms (legacy) section of the menu. <i>Managing alarms (legacy system)</i> on page 90
What user group permissions do I need?	<ul style="list-style-type: none"> Anyone who can sign in to the Grid Manager can view current and resolved alerts. You must have the Manage Alerts permission to manage silences, alert notifications, and alert rules. <i>Administering StorageGRID</i>	<ul style="list-style-type: none"> Anyone who can sign in to the Grid Manager can view legacy alarms. You must have the Acknowledge Alarms permission to acknowledge alarms. You must have both the Grid Topology Page Configuration and Other Grid Configuration permissions to manage global alarms and email notifications. <i>Administering StorageGRID</i>
How do I manage email notifications?	Select Alerts > Email Setup . Note: Because alarms and alerts are independent systems, the email setup used for alarm and AutoSupport notifications is not used for alert notifications. However, you can use the same mail server for all notifications. <i>Managing alert notifications</i> on page 81	Select Support . Then, in the Alarms (legacy) section of the menu, select Email Setup . <i>Configuring notifications for alarms (legacy system)</i> on page 102

	Alerts (new system)	Alarms (legacy system)
How do I manage SNMP notifications?	<p>Select Configuration. Then, in the Monitoring section of the menu, select SNMP Agent.</p> <p><i>Using SNMP monitoring</i> on page 110</p>	<p>Select Configuration. Then, in the Monitoring section of the menu, select SNMP Agent.</p> <p><i>Using SNMP monitoring</i> on page 110</p> <p>Note: SNMP notifications are not sent for every alarm or alarm severity.</p> <p><i>Alarms that generate SNMP notifications (legacy system)</i> on page 243</p>
How do I control who receives notifications?	<ol style="list-style-type: none"> 1. Select Alerts > Email Setup. 2. In the Recipients section, enter an email address for each email list or person who should receive an email when an alert occurs. <p><i>Setting up email notifications for alerts</i> on page 82</p>	<ol style="list-style-type: none"> 1. Select Support. Then, in the Alarms (legacy) section of the menu, select Email Setup. 2. Creating a mailing list. 3. Select Notifications. 4. Select the mailing list. <p><i>Creating mailing lists for alarm notifications (legacy system)</i> on page 106</p> <p><i>Configuring email notifications for alarms (legacy system)</i> on page 107</p>
Which Admin Nodes send notifications?	<p>A single Admin Node (the "preferred sender").</p> <p><i>Administering StorageGRID</i></p>	<p>A single Admin Node (the "preferred sender").</p> <p><i>Administering StorageGRID</i></p>
How do I suppress some notifications?	<ol style="list-style-type: none"> 1. Select Alerts > Silences. 2. Select the alert rule you want to silence. 3. Specify a duration for the silence. 4. Select the severity of alert you want to silence. 5. Select to apply the silence to the entire grid, a single site, or a single node. <p>Note: If you have enabled the SNMP agent, silences also suppress SNMP traps and informs.</p> <p><i>Silencing alert notifications</i> on page 87</p>	<ol style="list-style-type: none"> 1. Select Support. Then, in the Alarms (legacy) section of the menu, select Email Setup. 2. Select Notifications. 3. Select a mailing list, and select Suppress. <p><i>Suppressing alarm notifications for a mailing list (legacy system)</i> on page 107</p>
How do I suppress all notifications?	<p>Select Alerts > Silences. Then, select All rules.</p> <p>Note: If you have enabled the SNMP agent, silences also suppress SNMP traps and informs.</p> <p><i>Silencing alert notifications</i> on page 87</p>	<ol style="list-style-type: none"> 1. Select Configuration > Display Options. 2. Select the Notification Suppress All check box. <p>Note: Suppressing email notifications system wide also suppresses event-triggered AutoSupport emails.</p> <p><i>Suppressing email notifications system wide</i> on page 108</p>


	Alerts (new system)	Alarms (legacy system)
How do I customize the conditions and triggers?	<ol style="list-style-type: none"> 1. Select Alerts > Rules. 2. Select a default rule to edit, or select Create custom rule. <p><i>Editing an alert rule</i> on page 78 <i>Creating custom alert rules</i> on page 75</p>	<ol style="list-style-type: none"> 1. Select Support. Then, in the Alarms (legacy) section of the menu, select Global Alarms. 2. Create a Global Custom alarm to override a Default alarm or to monitor an attribute that does not have a Default alarm. <p><i>Creating Global Custom alarms (legacy system)</i> on page 97</p>
How do I disable an individual alert or alarm?	<ol style="list-style-type: none"> 1. Select Alerts > Rules. 2. Select the rule, and click Edit rule. 3. Unselect the Enabled check box. <p><i>Disabling an alert rule</i> on page 80</p>	<ol style="list-style-type: none"> 1. Select Support. Then, in the Alarms (legacy) section of the menu, select Global Alarms. 2. Select the rule, and click the Edit icon. 3. Unselect the Enabled check box. <p><i>Disabling a Default alarm (legacy system)</i> on page 100 <i>Disabling Global Custom alarms (legacy system)</i> on page 101</p>

Getting started with the alert system

If you are new to StorageGRID, you can start using the alert system exclusively. If you upgraded to StorageGRID 11.4 from a previous version, you must perform some steps to transition from the legacy alarm system to the new alert system.

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.

Which kind of user are you	Do this
I am a new user, starting with StorageGRID 11.4	<p>Use the alert system:</p> <ol style="list-style-type: none"> 1. Monitor current alerts on the Dashboard and on the Nodes page. <i>Viewing current alerts</i> on page 39 2. Set up email notifications for alerts. <i>Setting up email notifications for alerts</i> on page 82 3. Optionally, configure the StorageGRID SNMP agent, so you can receive alert notifications. <i>Using SNMP monitoring</i> on page 110 4. Optionally, edit default alert rules, or create custom alert rules to meet specific monitoring requirements. <i>Editing an alert rule</i> on page 78 <i>Creating custom alert rules</i> on page 75 5. Do not set up alarm email templates or mailing lists. 6. Ignore the Legacy alarm count on the Dashboard.

Which kind of user are you	Do this
<p>I upgraded to StorageGRID 11.4 from a previous StorageGRID version</p>	<p>Transition from the alarm system to the alert system:</p> <ol style="list-style-type: none"> 1. Start monitoring current alerts on the Dashboard and on the Nodes page. Viewing current alerts on page 39 2. Set up email notifications for alerts. Note: Because alerts and alarms are independent systems, the email setup used for alarm and AutoSupport notifications is not used for alert notifications. However, you can use the same email server for all notifications. Setting up email notifications for alerts on page 82 3. Optionally, configure the StorageGRID SNMP agent, so you can receive alert notifications. Note: If you have an existing SNMP configuration for StorageGRID, your existing trap destinations will apply to both the legacy alarm system and the new alert system. Using SNMP monitoring on page 110 4. Optionally, edit default alert rules, or create custom alert rules to meet specific monitoring requirements. Note: If you have created Global Custom alarms, consider creating equivalent custom alert rules using Prometheus metrics. Editing an alert rule on page 78 Creating custom alert rules on page 75 5. Optionally, view the Legacy alarm count on the Dashboard and acknowledge or disable specific alarms. Viewing legacy alarms on page 46 Acknowledging current alarms (legacy system) on page 94 Disabling alarms (legacy system) on page 99 6. When you have completed the transition to the alert system, suppress email notifications for each mailing list you configured for the alarm system.  Attention: Do not suppress email notifications system wide because you will also suppress event-triggered AutoSupport messages. Suppressing alarm notifications for a mailing list (legacy system) on page 107

Managing alerts

Alerts allow you to monitor various events and conditions within your StorageGRID system. You can manage alerts by creating custom alerts, editing or disabling the default alerts, setting up email notifications for alerts, and silencing alert notifications.

Related tasks

[Viewing current alerts](#) on page 39

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. An email notification might also be sent, unless the alert has been silenced.

[Viewing resolved alerts](#) on page 41

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

[Viewing a specific alert](#) on page 43

You can view detailed information about an alert that is currently affecting your StorageGRID system or an alert that has been resolved. The details include recommended corrective actions, the time the alert was triggered, and the current value of the metrics related to this alert. Optionally, you can silence a current alert or update the alert rule.

Related reference

[Alerts reference](#) on page 201

The following table lists all default StorageGRID alerts. As required, you can create custom alert rules to fit your system management approach.

Choices

- [What alerts are](#) on page 73
- [Managing alert rules](#) on page 74
- [Managing alert notifications](#) on page 81
- [Silencing alert notifications](#) on page 87

What alerts are

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

- The alert system focuses on actionable problems in the system. Unlike some alarms in the legacy system, alerts are triggered for events that require your immediate attention, not for events that can safely be ignored.
- The Current Alerts page provides a user-friendly interface for viewing current problems. You can sort the listing by individual alerts and alert groups. For example, you might want to sort all alerts by node/site to see which alerts are affecting a specific node. Or, you might want to sort the alerts in a group by time triggered to find the most recent instance of a specific alert.
- The Resolved Alerts page provides similar information as on the Current Alerts page, but it allows you to search and view a history of the alerts that have been resolved, including when the alert was triggered and when it was resolved.
- Multiple alerts of the same type are grouped into one email to reduce the number of notifications. In addition, multiple alerts of the same type are shown as a group on the Alerts page. You can expand and collapse alert groups to show or hide the individual alerts. For example, if several nodes report the **Unable to communicate with node** alert at about the same time, only one email is sent and the alert is shown as a group on the Alerts page.
- Alerts use intuitive names and descriptions to help you quickly understand the problem. Alert notifications include details about the node and site affected, the alert severity, the time when the alert rule was triggered, and the current value of metrics related to the alert.
- Alert emails notifications and the alert listings on the Current Alerts and Resolved Alerts pages provide recommended actions for resolving an alert. These recommended actions often include direct links to the StorageGRID documentation center to make it easier to find and access more detailed troubleshooting procedures.
- If you need to temporarily suppress the notifications for an alert at one or more severity levels, you can easily silence a specific alert rule for a specified duration and for the entire grid, a single site, or a single node. You can also silence all alert rules, for example, during a planned maintenance procedure such as a software upgrade.
- You can edit the default alert rules as required. You can disable an alert rule completely, or change its trigger conditions and duration.
- You can create custom alert rules to target the specific conditions that are relevant to your situation and to provide your own recommended actions. To define the conditions for a custom

alert, you create expressions using the Prometheus metrics available from the Metrics section of the Grid Management API.

Managing 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.

Steps

1. [Viewing alert rules](#) on page 74
2. [Creating custom alert rules](#) on page 75
3. [Editing an alert rule](#) on page 78
4. [Disabling an alert rule](#) on page 80
5. [Removing a custom alert rule](#) on page 81

Viewing 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.

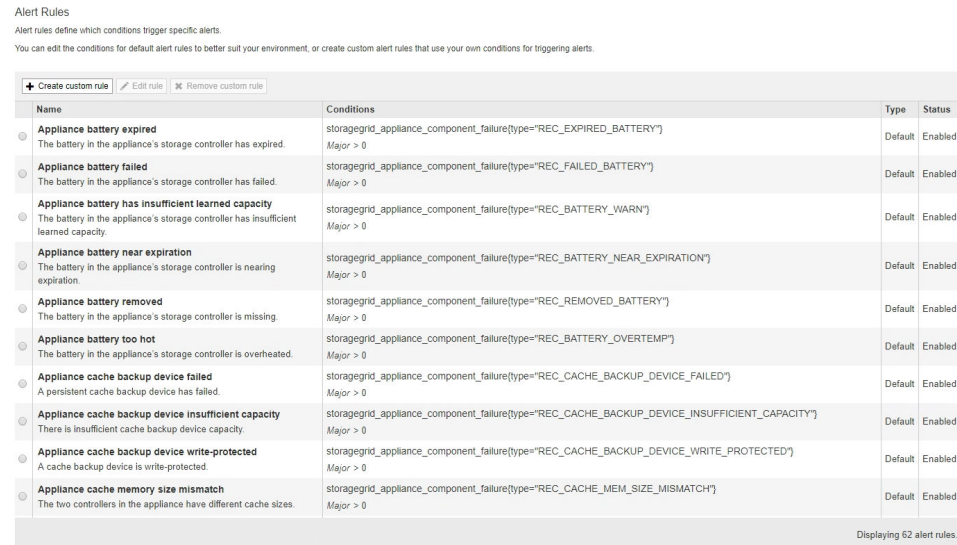
Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Manage Alerts or Root Access permission.

Steps




1. Select **Alerts > Rules**.

The Alert Rules page appears.



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.

Column header	Description
Conditions	<p>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.</p> <ul style="list-style-type: none"> • Critical : An abnormal condition exists that has stopped the normal operations of a StorageGRID node or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved. • Major : An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service. • Minor : The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that do not clear on their own to ensure they do not result in a more serious problem.
Type	<p>The type of alert rule:</p> <ul style="list-style-type: none"> • 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 cannot remove a default alert rule. • Default*: A default alert rule that includes an edited condition or duration. As required, you can easily revert a modified condition back to the original default. • Custom: An alert rule that you created. You can disable, edit, and remove custom alert rules.
Status	<p>Whether this alert rule is currently enabled or disabled. The conditions for disabled alert rules are not evaluated, so no alerts are triggered.</p>

Related reference

[Alerts reference](#) on page 201

The following table lists all default StorageGRID alerts. As required, you can create custom alert rules to fit your system management approach.

Creating custom alert rules

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

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Manage Alerts or Root Access permission.

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 simply 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 < 2400000000`, 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.

- Do not 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.

Create Custom Rule

Enabled

Unique Name

Description

Recommended Actions (optional)

Conditions ?

Minor

Major

Critical

Enter the amount of time a condition must continuously remain in effect before an alert is triggered.

Duration

3. Select or unselect the **Enabled** check box to determine if this alert rule is currently enabled.
If an alert rule is disabled, its expressions are not 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.
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.

- 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.

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

To learn about using the Grid Management API, see the instructions for administering StorageGRID. For details on the syntax of Prometheus queries, see the documentation for Prometheus.

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
```

- 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.

- Click **Save**.

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

Related reference

[Commonly used Prometheus metrics](#) on page 217

The Prometheus service on Admin Nodes collects time series metrics from the services on all nodes. While Prometheus collects more than a thousand metrics, a relatively small number are required to monitor the most critical StorageGRID operations.

Related information

[Administering StorageGRID](#)

[Prometheus: Query basics](#)

Editing an alert rule

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



Attention: 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 cannot be edited.

Edit Rule - Low installed node memory

Enabled

Unique Name

Description

Recommended Actions (optional) VMware installation- [Red Hat Enterprise Linux or CentOS installation](#)
- [Ubuntu or Debian installation](#)
"/>

Conditions ⓘ

Minor

Major

Critical

Enter the amount of time a condition must continuously remain in effect before an alert is triggered.

Duration

4. Select or unselect the **Enabled** check box to determine if this alert rule is currently enabled.

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

Note: 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.



Attention: 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.

Note: You cannot 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.

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

Conditions 

Minor	<input type="text"/>
Major	<input type="text" value="node_memory_MemTotal_bytes < 2400000000"/>
Critical	<input type="text" value="node_memory_MemTotal_bytes <= 1400000000"/>




Note: 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.

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

To learn about using the Grid Management API, see the instructions for administering StorageGRID. For details on the syntax of Prometheus queries, see the documentation for Prometheus.

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. Click **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.

Related reference

[Commonly used Prometheus metrics](#) on page 217

The Prometheus service on Admin Nodes collects time series metrics from the services on all nodes. While Prometheus collects more than a thousand metrics, a relatively small number are required to monitor the most critical StorageGRID operations.

Related information

[Administering StorageGRID](#)

[Prometheus: Query basics](#)

Disabling an alert rule

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

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Manage Alerts or Root Access permission.

About this task

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



Attention: 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 unselect the **Enabled** check box to determine if this alert rule is currently enabled.

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

Note: 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. Click **Save**.

Disabled appears in the **Status** column.

Removing a custom alert rule

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

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must 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 cannot remove a default alert rule.
3. Click **Remove custom rule**.
A confirmation dialog box appears.
4. Click **OK** to remove the alert rule.
Any active instances of the alert will be resolved within 10 minutes.

Managing alert notifications

When an alert is triggered, StorageGRID can send email notifications and Simple Network Management Protocol (SNMP) notifications (traps).

Steps

1. [Setting up SNMP notifications for alerts](#) on page 81
2. [Setting up email notifications for alerts](#) on page 82
3. [Information included in alert email notifications](#) on page 85
4. [How StorageGRID groups alerts in email notifications](#) on page 86
5. [Troubleshooting alert email notifications](#) on page 87

Setting 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.

About this task

You can use the **Configuration > 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 the section for using 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 do not require acknowledgment by the management system. Traps serve to notify the management system that something has happened within StorageGRID, such as an alert being triggered. Traps are supported in all three versions of SNMP
- Informs are similar to traps, but they require acknowledgment by the management system. If the SNMP agent does not receive an acknowledgment within a certain amount of time, it resends the inform until an acknowledgment is received or the maximum retry value has been reached. Informs are supported in SNMPv2c and SNMPv3.

Trap and inform notifications are sent 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. Alert notifications are sent by whichever Admin Node is configured to be the preferred sender. By

default, the primary Admin Node is selected. For details, see the instructions for administering StorageGRID.

Note: Trap and inform notifications are also sent when certain alarms (legacy system) are triggered at specified severity levels or higher; however, SNMP notifications are not sent for every alarm or every alarm severity.

Related concepts

[Using SNMP monitoring](#) on page 110

If you want to monitor StorageGRID using the Simple Network Management Protocol (SNMP), you must configure the SNMP agent that is included with StorageGRID.

Related tasks

[Silencing alert notifications](#) on page 87

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

Related reference

[Alarms that generate SNMP notifications \(legacy system\)](#) on page 243

Related information

[Administering StorageGRID](#)

Setting 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 must be signed in to the Grid Manager using a supported browser.
- You must 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, you can select which Admin Node should be the preferred sender of alert notifications. The same "preferred sender" is also used for alarm notifications and AutoSupport messages. By default, the primary Admin Node is selected. For details, see the instructions for administering StorageGRID.

Steps

1. Select Alerts > Email Setup.

The Email Setup page appears.

Email Setup

You can configure the email server for alert notifications, define filters to limit the number of notifications, and enter email addresses for alert recipients.

Use these settings to define the email server used for alert notifications. These settings are not used for alarm notifications and AutoSupport. See [Managing alerts and alarms](#) in the instructions for monitoring and troubleshooting StorageGRID.

Enable Email Notifications

Save

2. Select the **Enable Email Notifications** check box 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. You must also require TLS and provide a CA certificate.

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

Email (SMTP) Server

Mail Server

Port

Username (optional)

Password (optional)

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.

Click the plus icon **+** to add recipients.

Email Addresses

Sender Email Address

Recipient 1 **x**

Recipient 2 **+ x**

5. In the **Transport Layer Security (TLS)** section, select the **Require TLS** check box if Transport Layer Security (TLS) is required for communications with the SMTP server.
 - 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 click **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** check box 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 click **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 click **Browse** and select the file.

Note: If you need to edit the email setup, click the pencil icon to update this field.

Transport Layer Security (TLS)

Require TLS

CA Certificate

```
-----BEGIN CERTIFICATE-----
abcdefghijklmnopqrstuvwxyz123456780ABCDEFGHIJKL
123456/7890ABCDEFGHIabcdefghijklmnopqrstuvwxyzLABCD
-----END CERTIFICATE-----
```

Browse

Send Client Certificate

Client Certificate

```
-----BEGIN CERTIFICATE-----
abcdefghijklmnopqrstuvwxyz123456780ABCDEFGHIJKL
123456/7890ABCDEFGHIabcdefghijklmnopqrstuvwxyzLABCD
-----END CERTIFICATE-----
```

Browse

Private Key

```
-----BEGIN PRIVATE KEY-----
abcdefghijklmnopqrstuvwxyz123456780ABCDEFGHIJKL
123456/7890ABCDEFGHIabcdefghijklmnopqrstuvwxyzLABCD
-----END 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 are not sent for minor alerts.
Critical only	An email notification is sent only when the critical condition for an alert rule is met. Notifications are not sent for minor or major alerts.

Filters

Severity Minor, major, critical Major, critical Critical only

- 7. When you are ready to test your email settings, perform these steps:

- a. Click **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.
Note: 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.
Note: When you test alert notifications, you must sign in to every Admin Node to verify connectivity. This is in contrast to testing alarm notifications and AutoSupport messages, where all Admin Nodes send the test email.

8. Click Save.

Sending a test email does not save your settings. You must click **Save**.
 The email settings are saved.

Related tasks

[Troubleshooting alert email notifications](#) on page 87

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.

Related information

[Recovery and maintenance](#)

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.

Email notifications include the following information:

NetApp StorageGRID

Low object data storage (6 alerts) 1

The space available for storing object data is low. 2

Recommended actions 3

Perform an expansion procedure. You can add storage volumes (LUNs) to existing Storage Nodes, or you can add new Storage Nodes. See the instructions for expanding a StorageGRID system.

DC1-S1-226

Node	DC1-S1-226	4
Site	DC1 225-230	
Severity	Minor	
Time triggered	Fri Jun 28 14:43:27 UTC 2019	
Job	storagegrid	
Service	ldr	

DC1-S2-227

Node	DC1-S2-227
Site	DC1 225-230
Severity	Minor
Time triggered	Fri Jun 28 14:43:27 UTC 2019
Job	storagegrid
Service	ldr

Sent from: DC1-ADM1-225 5

Callout	Description
1	The name of the alert, followed by the number of active instances of this alert.

Callout	Description
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.

Related tasks

[Silencing alert notifications](#) on page 87

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

How StorageGRID groups alerts in email notifications

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.	<ul style="list-style-type: none"> Alert A is triggered on two nodes at the same time. Only one notification is sent. Alert A is triggered on node 1, and Alert B is triggered on node 2 at the same time. Two notifications are sent—one for each alert.
For a specific alert on a specific node, if the thresholds are reached for more than one severity, a notification is sent only for the most severe alert.	<ul style="list-style-type: none"> 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.	<ol style="list-style-type: none"> Alert A is triggered on node 1 at 08:00. No notification is sent. Alert A is triggered on node 2 at 08:01. No notification is sent. 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.	<ol style="list-style-type: none"> Alert A is triggered on node 1 at 08:00. A notification is sent at 08:02. 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.

Behavior	Example
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.	<ol style="list-style-type: none"> 1. Alert A is triggered for node 1. A notification is sent. 2. Alert A is triggered for node 2. A second notification is sent. 3. Alert A is resolved for node 2, but it remains active for node 1. 4. Alert A is triggered again for node 2. No new notification is sent because the alert is still active for node 1.
StorageGRID continues to send email notifications once every 7 days until all instances of the alert are resolved or the alert rule is silenced.	<ol style="list-style-type: none"> 1. Alert A is triggered for node 1 on March 8. A notification is sent. 2. Alert A is not resolved or silenced. Additional notifications are sent on March 15, March 22, March 29, and so on.

Troubleshooting 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 must be signed in to the Grid Manager using a supported browser.
- You must 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 are not 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.

Related tasks

[Collecting log files and system data](#) on page 139

You can use the Grid Manager to retrieve log files and system data (including configuration data) for your StorageGRID system.

Silencing alert notifications

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

Before you begin

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



Attention: 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.

Note: Because alarms and alerts are independent systems, you cannot use this functionality to suppress alarm notifications.

Steps

1. Select Alerts > Silences.

The Silences page appears.

Silences

You can configure silences to temporarily suppress alert notifications. Each silence suppresses the notifications for an alert rule at one or more severities. You can suppress an alert rule on the entire grid, a single site, or a single node.

Alert Rule	Description	Severity	Time Remaining	Nodes
No results found.				

2. Select Create.

The Create Silence dialog box appears.

Create Silence

Alert Rule

Description (optional)

Duration

Severity Minor only Minor, major Minor, major, critical

Nodes StorageGRID Deployment

- Data Center 1
 - DC1-ADM1
 - DC1-G1
 - DC1-S1
 - DC1-S2
 - DC1-S3

3. Select or enter the following information:

Field	Description
Alert Rule	The name of the alert rule you want to silence. You can select any default or custom alert rule, even if the alert rule is disabled. Note: Select All rules if you want to silence all alert rules using the criteria specified in this dialog box.
Description	Optionally, a description of the silence. For example, describe the purpose of this silence.
Duration	How long you want this silence to remain in effect, in minutes, hours, or days. A silence can be in effect from 5 minutes to 1,825 days (5 years). Note: You should not silence an alert rule for an extended amount of time. If an alert rule is silenced, you might not detect an underlying problem until it prevents a critical operation from completing. However, you might need to use an extended silence if an alert is triggered by a specific, intentional configuration, such as might be the case for the Services appliance link down alerts and the Storage appliance link down alerts.
Severity	Which alert severity or severities should be silenced. If the alert is triggered at one of the selected severities, no notifications are sent.
Nodes	Which node or nodes you want this silence to apply to. You can suppress an alert rule or all rules on the entire grid, a single site, or a single node. If you select the entire grid, the silence applies to all sites and all nodes. If you select a site, the silence applies only to the nodes at that site. Note: You cannot 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. Click **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	<ol style="list-style-type: none"> a. Select Alerts > Silences. b. From the table, select the radio button for the silence you want to edit. c. Click Edit. d. Change the description, the amount of time remaining, the selected severities, or the affected node. e. Click Save.
Remove a silence	<ol style="list-style-type: none"> a. Select Alerts > Silences. b. From the table, select the radio button for the silence you want to remove. c. Click Remove. d. Click OK to confirm you want to remove this silence. <p>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.</p>

Related tasks

[Configuring the SNMP agent](#) on page 111

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.

Managing alarms (legacy system)

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

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.

Related tasks

[Viewing legacy alarms](#) on page 46

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

Related reference

[Alarms reference \(legacy system\)](#) on page 222

Related information

[Administering StorageGRID](#)

Steps

1. [Alarm classes \(legacy system\)](#) on page 90
2. [Alarm triggering logic \(legacy system\)](#) on page 91
3. [Acknowledging current alarms \(legacy system\)](#) on page 94
4. [Viewing Default alarms \(legacy system\)](#) on page 95
5. [Reviewing historical alarms and alarm frequency \(legacy system\)](#) on page 96
6. [Creating Global Custom alarms \(legacy system\)](#) on page 97
7. [Disabling alarms \(legacy system\)](#) on page 99
8. [Configuring notifications for alarms \(legacy system\)](#) on page 102

Alarm classes (legacy system)

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

Default alarms

Default alarms are provided with each StorageGRID system and cannot be modified. However, you can disable Default alarms or override them by defining Global Custom alarms.

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.

Related tasks

[Viewing Default alarms \(legacy system\)](#) on page 95

You can view the list of all Default legacy alarms.

[Disabling a Default alarm \(legacy system\)](#) on page 100

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

[Creating Global Custom alarms \(legacy system\)](#) on page 97





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 do not have a Default alarm.

Disabling Global Custom alarms (legacy system) on page 101

You can disable a legacy Global Custom alarm for the entire 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 are not 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.

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

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

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

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

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

Alarms of same severity

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

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



Global Alarms
 Updated: 2016-03-17 16:05:31 PDT

Global Custom Alarms (0 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	Under 50	=	5000		
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	under100	<=	1000		

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



Global Custom Alarms (0 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	under10i	<=	1000		
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	Under 50	=	5000		

Default Alarms

Filter by Disabled Defaults

0 Result(s)

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
---------	---------	-----------	----------	---------	----------	-------	---------

Apply Changes

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 cannot be disabled for individual rows in a table.

For example, the following table (from **Support > Grid Topology > Storage Node > SSM > Resources**) shows two critical Entries Available (VMFI) alarms. 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 cannot 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

Acknowledging current alarms (legacy system)

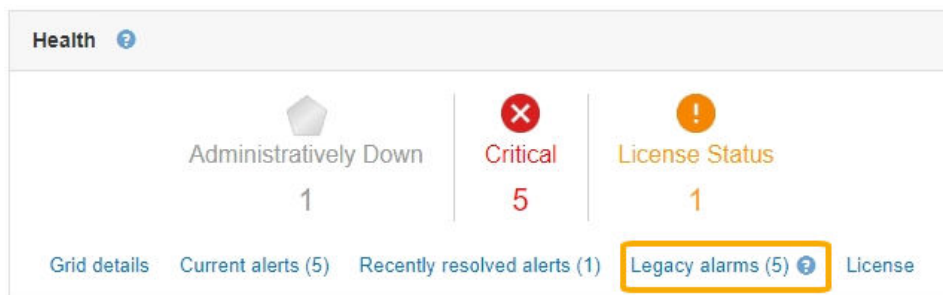
Legacy alarms are triggered when system attributes reach alarm threshold values. If you want to reduce or clear the count of legacy alarms on the Dashboard, you can acknowledge the alarms.

Before you begin

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

About this task

If an alarm from the legacy system is currently active, the Health panel on the Dashboard includes a **Legacy alarms** link. The number in parentheses indicates how many legacy alarms are currently active.



Because the legacy alarm system continues to be supported in StorageGRID 11.4, the number of legacy alarms shown on the Dashboard is incremented whenever a new alarm occurs. This count is incremented even if email notifications are no longer being sent for alarms. You can typically just ignore this number (since alerts provide a better view of the system), or you can acknowledge the alarms.

Note: 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 included in the count of legacy alarms unless the alarm is triggered at the next severity level or it is resolved and occurs again.

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.

Steps

1. To view the alarm, do one of the following:
 - From the Health panel on the Dashboard, click **Legacy alarms**. This link appears only if at least one alarm is currently active.
 - Select **Support**. Then, in the Alarms (legacy) section of the menu, select **Current Alarms**.

The Current Alarms page appears.

The alarm system is the legacy system. The alert system offers significant benefits and is easier to use. See [Managing alerts and alarms](#) in the instructions for monitoring and troubleshooting StorageGRID.

Current Alarms

Last Refreshed: 2020-05-27 09:41:39 MDT

Show Acknowledged Alarms (1 - 1 of 1)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value
Major	ORSU (Outbound Replication Status)	Data Center 1/DC1-ARC1/ARC	Storage Unavailable	2020-05-26 21:47:18 MDT	Storage Unavailable	Storage Unavailable

Show 50 Records Per Page Refresh Previous 1 Next

- Click the service name in the table. The Alarms tab for the selected service appears (**Support > Grid Topology > Grid Node > Service > Alarms**).

Overview | **Alarms** | Reports | Configuration

Main | History

Alarms: ARC (DC1-ARC1) - Replication
 Updated: 2019-05-24 10:46:48 MDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Major	ORSU (Outbound Replication Status)	Storage Unavailable	2019-05-23 21:40:08 MDT	Storage Unavailable	Storage Unavailable		<input type="checkbox"/>

Apply Changes

- Select the **Acknowledge** check box for the alarm, and click **Apply Changes**. The alarm no longer appears on the Dashboard or the Current Alarms page.

Note: 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.
- As required, view acknowledged alarms.
 - Select **Support**. Then, in the Alarms (legacy) section of the menu, select **Current Alarms**.
 - Select **Show Acknowledged Alarms**. Any acknowledged alarms are shown.

The alarm system is the legacy system. The alert system offers significant benefits and is easier to use. See [Managing alerts and alarms](#) in the instructions for monitoring and troubleshooting StorageGRID.

Current Alarms

Last Refreshed: 2020-05-27 17:38:58 MDT

Show Acknowledged Alarms (1 - 1 of 1)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time
Major	ORSU (Outbound Replication Status)	Data Center 1/DC1-ARC1/ARC	Storage Unavailable	2020-05-26 21:47:18 MDT	Storage Unavailable	Storage Unavailable	2020-05-27 17:38:14 MDT

Show 50 Records Per Page Refresh Previous 1 Next

Related reference

[Alarms reference \(legacy system\)](#) on page 222

Viewing 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 browser.
- You must have specific access permissions.

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.





Steps

1. Select **Support**. Then, in the Alarms (legacy) section of the menu, select **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.




Global Alarms
 Updated: 2019-03-01 15:13:02 MST

























Global Custom Alarms (0 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input type="checkbox"/>								   

Default Alarms

Filter by Attribute Code ▼ equals * 

221 Result(s)

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
<input checked="" type="checkbox"/>		IQSZ (Number of Objects)	 Major	Greater than 10,000,000	>=	10000000	 
<input checked="" type="checkbox"/>		IQSZ (Number of Objects)	 Minor	Greater than 1,000,000	>=	1000000	 
<input checked="" type="checkbox"/>		IQSZ (Number of Objects)	 Notice	Greater than 150,000	>=	150000	 
<input checked="" type="checkbox"/>		XCVF (% Completion)	 Notice	Foreground Verification Completed	=	100	 
<input checked="" type="checkbox"/>	ADC	ADCA (ADC Status)	 Minor	Error	>=	10	 
<input checked="" type="checkbox"/>	ADC	ADCE (ADC State)	 Notice	Standby	=	10	 
<input checked="" type="checkbox"/>	ADC	ALIS (Inbound Attribute Sessions)	 Notice	Over 100	>=	100	 
<input checked="" type="checkbox"/>	ADC	ALOS (Outbound Attribute Sessions)	 Notice	Over 200	>=	200	 

Reviewing 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 browser.
- You must have specific access permissions.

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new 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**. Then, in the Alarms (legacy) section of the menu, select **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**. Then, in the Tools section of the menu, select **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**.

Related reference

[Alarms reference \(legacy system\)](#) on page 222

Creating 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 do not have a Default alarm.

Before you begin

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

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.



About this task

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.



Attention: Be very 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**. Then, in the Alarms (legacy) section of the menu, select **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** .



Global Custom Alarms (0 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input checked="" type="checkbox"/>	ARC	ARCE (ARC State)	Notice	Standby	=	10		
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)	Minor	At least 6000	>=	6000		
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)	Notice	At least 3000	>=	3000		

Default Alarms

Filter by Attribute Code equals AR*






9 Result(s)

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
<input checked="" type="checkbox"/>	ARC	ARCE (ARC State)	Notice	Standby	=	10	
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)	Minor	At least 6000	>=	6000	
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)	Notice	At least 3000	>=	3000	
<input checked="" type="checkbox"/>	ARC	ARRF (Request Failures)	Major	At least 1	>=	1	
<input checked="" type="checkbox"/>	ARC	ARRV (Verification Failures)	Major	At least 1	>=	1	
<input checked="" type="checkbox"/>	ARC	ARVF (Store Failures)	Major	At least 1	>=	1	
<input checked="" type="checkbox"/>	NMS	ARRC (Remaining Capacity)	Notice	Below 10	<=	10	
<input checked="" type="checkbox"/>	NMS	ARRS (Repository Status)	Major	Disconnected	<=	9	
<input checked="" type="checkbox"/>	NMS	ARRS (Repository Status)	Notice	Standby	<=	19	

Apply Changes

- To modify a Default alarm, search for the Default alarm.
 - a. Under Filter by, select either **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**.
 - d. 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 unselect the check box to enable or disable the alarm.
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).

Heading	Description
Operator	Operators for testing the current attribute value against the Value threshold: <ul style="list-style-type: none"> • = equals • > greater than • < less than • >= greater than or equal to • <= less than or equal to • ≠ not equal to
Value	The alarm’s threshold value used to test against the attribute’s actual value using the operator. The entry can be a single number, a range of numbers specified with a colon (1:3), or a comma-delineated list of numbers and ranges.
Additional Recipients	A supplementary list of email addresses to be notified when the alarm is triggered. This is in addition to the mailing list configured on the Alarms > Email Setup page. Lists are comma delineated. <p>Note: Mailing lists require SMTP server setup in order to operate. Before adding mailing lists, confirm that SMTP is configured.</p> Notifications for Custom alarms can override notifications from Global Custom or Default alarms.
Actions	Control buttons to: <ul style="list-style-type: none">  Edit a row  Insert a row  Delete a row  Drag-and-drop a row up or down  Copy a row

4. Click **Apply Changes**.

Related tasks

[Configuring email server settings for alarms \(legacy system\)](#) on page 103

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 cannot receive email.

Disabling alarms (legacy system)

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

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.

Steps

1. [Disabling a Default alarm \(legacy system\)](#) on page 100
2. [Disabling Global Custom alarms \(legacy system\)](#) on page 101
3. [Clearing triggered alarms \(legacy system\)](#) on page 101

Disabling 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 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.



CAUTION: Do not 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**. Then, in the Alarms (legacy) section of the menu, select **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**.





Note: 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.




Global Alarms
Updated: 2017-03-30 15:47:43 MDT










Global Custom Alarms (0 Result(s))


Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input type="checkbox"/>								   

Default Alarms

Filter by Attribute Code equals U* 

3 Result(s)

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	 Critical	Under 10000000	<=	10000000	 
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	 Major	Under 50000000	<=	50000000	 
<input type="checkbox"/>	SSM	UMEM (Available Memory)	 Minor	Under 100000000	<=	100000000	 

Apply Changes 

The **Enabled** check box for the selected alarm becomes active.

4. Unselect the **Enabled** check box.
5. Click **Apply Changes**.

The Default alarm is disabled.

Disabling Global Custom alarms (legacy system)

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


Before you begin

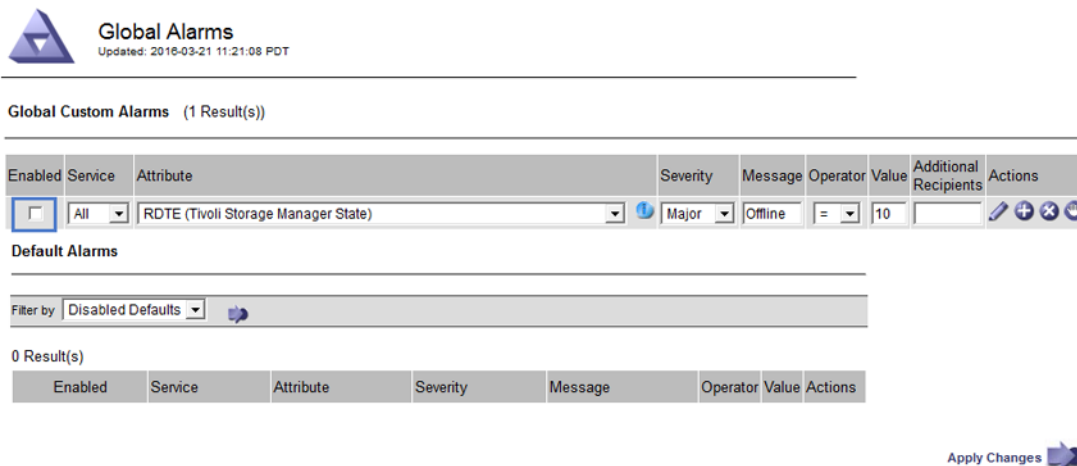
- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

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**. Then, in the Alarms (legacy) section of the menu, select **Global Alarms**.
2. In the Global Custom Alarms table, click **Edit**  next to the alarm you want to disable.
3. Unselect the **Enabled** check box.



4. Click **Apply Changes**.
The Global Custom alarm is disabled.

Clearing 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.

About this task

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).

Steps

1. Disable the alarm.

2. From the service laptop, 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.

Related concepts

[Disabling alarms \(legacy system\)](#) on page 99

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

Configuring 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 are not sent. For email notifications, you must configure the email server and specify the email recipients. For SNMP notifications, you must configure the SNMP agent.

Related concepts

[Using SNMP monitoring](#) on page 110

If you want to monitor StorageGRID using the Simple Network Management Protocol (SNMP), you must configure the SNMP agent that is included with StorageGRID.

Steps

1. [Types of alarm notifications \(legacy system\)](#) on page 102
2. [Configuring email server settings for alarms \(legacy system\)](#) on page 103
3. [Creating alarm email templates \(legacy system\)](#) on page 105
4. [Creating mailing lists for alarm notifications \(legacy system\)](#) on page 106
5. [Configuring email notifications for alarms \(legacy system\)](#) on page 107
6. [Suppressing alarm notifications for a mailing list \(legacy system\)](#) on page 107
7. [Suppressing email notifications system wide](#) on page 108

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 sent when a service enters or leaves one of the following service states:

- Unknown
- Administratively Down

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

Related tasks

[Configuring email notifications for alarms \(legacy system\)](#) on page 107

In order 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.

Configuring 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 cannot receive email.

Before you begin

- You must be signed in to the Grid Manager using a supported 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 are not used for alert notifications.

Note: 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**. Then, in the Alarms (legacy) section of the menu, select **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.

Use these settings to define the email server used for alarm notifications and for AutoSupport messages. These settings are not used for alert notifications. See *Managing alerts and alarms* in the instructions for monitoring and troubleshooting StorageGRID.



Email Server
 Updated: 2018-03-17 11:11:59 PDT

E-mail Server (SMTP) Information

Mail Server	<input type="text"/>
Port	<input type="text"/>
Authentication	Off <input type="button" value="v"/>
Authentication Credentials	Username: <input type="text" value="root"/> Password: <input type="password" value="....."/>
From Address	<input type="text"/>
Test E-mail	To: <input type="text"/> <input type="checkbox"/> Send Test E-mail

Apply Changes

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 are not 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.

Related information

[Administering StorageGRID](#)

Creating 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 browser.
- You must have specific access permissions.

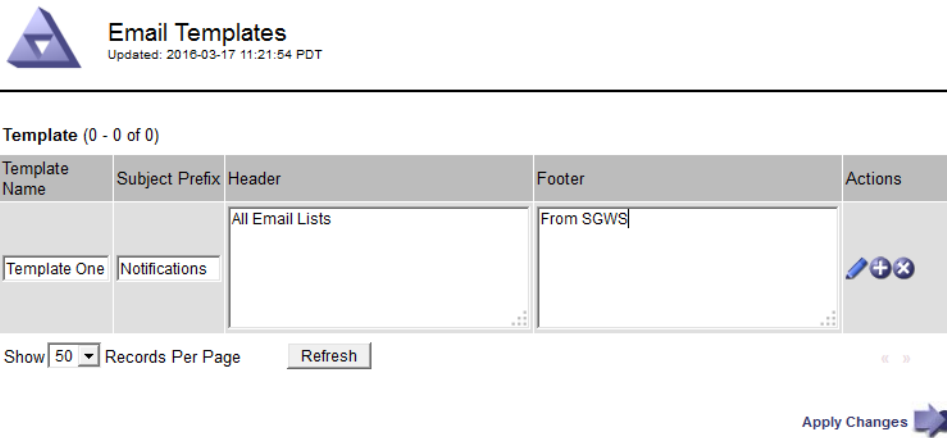
About this task

Use these settings to define the email templates used for legacy alarm notifications. These settings are not used for alert notifications.

Different mailing lists might require different contact information. Templates do not include the body text of the email message.

Steps

1. Select **Support**. Then, in the Alarms (legacy) section of the menu, select **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 cannot 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.

Creating 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

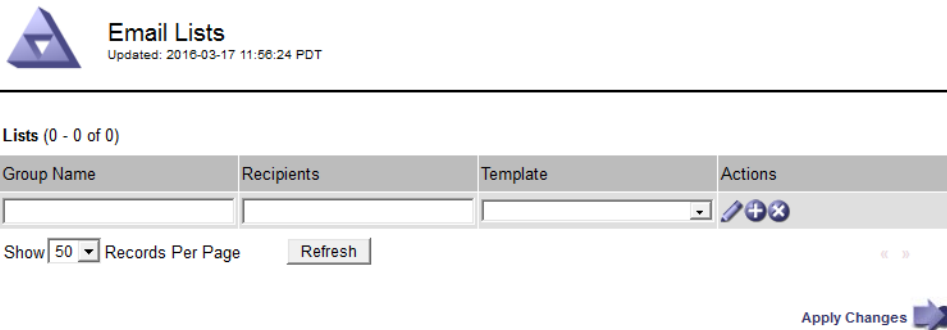
- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- If you want to specify an email template for the mailing list (custom header, footer, and subject line), you must have already created the template.

About this task

Use these settings to define the mailing lists used for legacy alarm email notifications. These settings are not used for alert notifications.

Steps

1. Select **Support**. Then, in the Alarms (legacy) section of the menu, select **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 cannot be duplicated. Note: If you change the name of a mailing list, the change is not propagated to the other locations that use the mailing list name. You must manually update all configured notifications to use the new mailing list name.
Recipients	Single email address, a previously configured mailing list, or a comma-delineated list of email addresses and mailing lists to which notifications will be sent. Note: If an email address belongs to multiple mailing lists, only one email notification is sent when a notification triggering event occurs.
Template	Optionally, select an email template to add a unique header, footer, and subject line to notifications sent to all recipients of this mailing list.

5. Click **Apply Changes**.
A new mailing list is created.

Related tasks

[Creating alarm email templates \(legacy system\)](#) on page 105

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.

Configuring email notifications for alarms (legacy system)

In order 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 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 are not 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**. Then, in the Alarms (legacy) section of the menu, select **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.

Related concepts

[Types of alarm notifications \(legacy system\)](#) on page 102

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

Related tasks

[Creating mailing lists for alarm notifications \(legacy system\)](#) on page 106

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.

Suppressing 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 browser.


- You must have specific access permissions.

About this task

Use these settings to suppress email notifications for the legacy alarm system. These settings do not apply to alert email notifications.

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.

Steps

1. Select **Support**. Then, in the Alarms (legacy) section of the menu, select **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 check box 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.

Suppressing email notifications system wide

You can block the StorageGRID system's ability to send alarm email notifications and event-triggered AutoSupport email notifications.

Before you begin

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

About this task

Use this option to suppress alarm email notifications as well as event-triggered AutoSupport email notifications. For example, you might want to suppress notifications when you upgrade StorageGRID software or apply a hotfix.


Note: This option does not suppress alert email notifications.

Suppressing email notifications system wide also suppresses event-triggered AutoSupport emails, even when the Enabled check box is selected on the Event-Triggered AutoSupport page (**Support > AutoSupport > Event-triggered**).

Suppressing email notifications system wide does not suppress weekly or user-triggered AutoSupport messages.

Steps

1. Select **Configuration > Display Options**.
2. From the Display Options menu, select **Options**.
3. Select **Notification Suppress All**.



Display Options


Updated: 2017-03-23 18:03:48 MDT

Current Sender: ADMIN-DC1-ADM1

Preferred Sender: ADMIN-DC1-ADM1


GUI Inactivity Timeout: 900

Notification Suppress All:

Apply Changes 

4. Click **Apply Changes**.

The Notifications page (**Configuration > Notifications**) displays the following message:






Notifications


Updated: 2016-03-17 14:06:48 PDT

All e-mail notifications are now suppressed.

Notifications (0 - 0 of 0)

E-mail List	Suppress	Severity Levels				Service States		Actions
	<input checked="" type="checkbox"/>	Notice	Minor	Major	Critical	Unknown	Administratively Down	
	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	  

Show Records Per Page

Apply Changes 

Related information

- [Administering StorageGRID](#)
- [Upgrading StorageGRID](#)

Using SNMP monitoring

If you want to monitor StorageGRID using the Simple Network Management Protocol (SNMP), you must configure the SNMP agent that is included with StorageGRID.

Capabilities

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

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

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

Informs are supported in SNMPv2c and SNMPv3.

Trap and inform notifications are sent in the following cases:

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

Note: SNMP notifications are not 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.

Accessing the MIB

You can access the MIB definition file at the following location on any StorageGRID node:

```
/usr/share/snmp/mibs/NETAPP-STORAGEGRID-MIB.txt
```

Related tasks

[Silencing alert notifications](#) on page 87

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

Related reference

[Alerts reference](#) on page 201

The following table lists all default StorageGRID alerts. As required, you can create custom alert rules to fit your system management approach.

[Alarms reference \(legacy system\)](#) on page 222

[Alarms that generate SNMP notifications \(legacy system\)](#) on page 243

Configuring 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 must be signed in to the Grid Manager using a supported browser.
- You must 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**. Then, in the Monitoring section of the menu, select **SNMP Agent**. The SNMP Agent page appears.

SNMP Agent
You can configure SNMP for read-only MIB access and notifications. SNMPv1, SNMPv2c, SNMPv3 are supported. For SNMPv3, only User Security Model (USM) authentication is supported. All nodes in the grid share the same SNMP configuration.

Enable SNMP

Save

2. To enable the SNMP agent on all grid nodes, select the **Enable SNMP** check box. The fields for configuring an SNMP agent appear.

SNMP Agent

You can configure SNMP for read-only MIB access and notifications. SNMPv1, SNMPv2c, SNMPv3 are supported. For SNMPv3, only User Security Model (USM) authentication is supported. All nodes in the grid share the same SNMP configuration.

Enable SNMP

System Contact

System Location

Enable SNMP Agent Notifications

Enable Authentication Traps

Community Strings

Default Trap Community

Read-Only Community

String 1 +

Other Configurations

Agent Addresses (0) USM Users (0) Trap Destinations (0)

+ Create Edit Remove

Internet Protocol	Transport Protocol	StorageGRID Network	Port
No results found.			

Save

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

If this check box is unselected, the SNMP agent supports read-only MIB access, but it does not send any SNMP notifications.
- Select the **Enable Authentication Traps** check box if you want the StorageGRID SNMP agent to send an authentication trap if it receives an improperly authenticated protocol message.
- 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 do not apply to SNMPv3.

 - 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 in step 10 on page 116.

Default Trap Community can be a maximum of 32 characters and cannot 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. Click 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 cannot contain whitespace characters. Up to five strings are allowed.

Note: To ensure the security of your StorageGRID system, do not use "public" as the community string. If you do not 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 do not configure an agent address, the default listening address is UDP port 161 on all StorageGRID networks.

- a. Click **Create**.

The Create Agent Address dialog box appears.

The screenshot shows a dialog box titled "Create Agent Address". It has two rows of radio buttons: "Internet Protocol" with "IPv4" selected and "IPv6" unselected; and "Transport Protocol" with "UDP" selected and "TCP" unselected. Below these is a dropdown menu for "StorageGRID Network" with "Grid, Admin, and Client Networks" selected. At the bottom is a text field for "Port" containing "161". At the bottom right are "Cancel" and "Create" buttons.

- 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

Note: 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.

Note: When you save the SNMP agent, StorageGRID automatically opens the agent address ports on the internal firewall. You must ensure that any external firewalls allow access to these ports.

- f. Click **Create**.

The agent address is created and added to the table.

Other Configurations

Agent Addresses (2) **USM Users (2)** Trap Destinations (2)

	Internet Protocol	Transport Protocol	StorageGRID Network	Port
<input type="radio"/>	IPv4	UDP	Grid Network	161
<input checked="" type="radio"/>	IPv4	UDP	Admin Network	161

- 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.

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

- a. Click **Create**.

The Create USM User dialog box appears.

Create USM User

Username

Read-Only MIB Access

Authoritative Engine ID

Security Level authPriv authNoPriv

Authentication

Protocol

Password

Confirm Password

Privacy

Protocol

Password

Confirm Password

- b. Enter a unique **Username** for this USM user.
Usernames have a maximum of 32 characters and cannot contain whitespace characters. The username cannot be changed after the user is created.
- c. Select the **Read-Only MIB Access** check box 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.
Note: USM users who have read-only MIB access cannot have engine IDs.
- d. If this user will be used in an inform destination, enter the **Authoritative Engine ID** for this user.
Note: SNMPv3 *inform* destinations must have users with engine IDs. SNMPv3 *trap* destination cannot 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.
Note: 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.
Note: The only privacy protocol supported is AES.
- h. Click **Create**.
The USM user is created and added to the table.

Other Configurations

Agent Addresses (2) USM Users (3) Trap Destinations (2)

+ Create ✎ Edit ✕ Remove				
	Username	Read-Only MIB Access	Security Level	Authoritative Engine ID
<input type="radio"/>	user2	✓	authNoPriv	
<input type="radio"/>	user1		authNoPriv	B3A73C2F3D6
<input checked="" type="radio"/>	user3		authPriv	59D39E801256

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 click **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. Click **Create**.
The Create Trap Destination dialog box appears.

Create Trap Destination

Version SNMPv1 SNMPv2C SNMPv3

Type Trap

Host

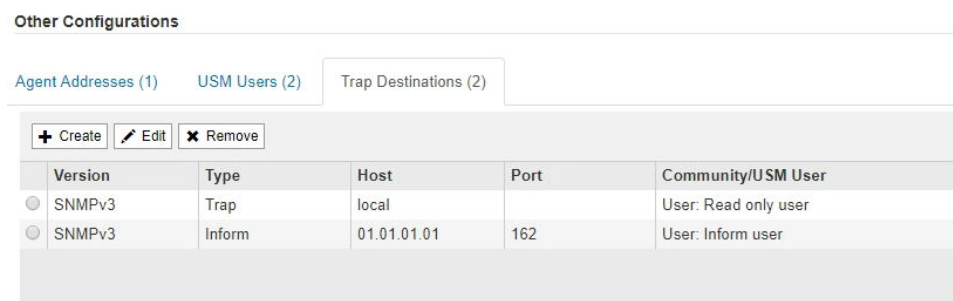
Port (optional)

Community String Use the default trap community: Default Community
 Use a custom community string

- 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	<p>Note: For SNMPv1, the SNMP agent can only send traps. Informs are not supported.</p> <ul style="list-style-type: none"> i. In the Host field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap. Only UDP is supported. ii. Optionally, enter the Port that will receive the trap. The default is port 162. iii. Use the default trap community, if one was specified on the SNMP Agent page, or enter a custom community string for this trap destination. The custom community string can be a maximum of 32 characters and cannot contain whitespace.
SNMPv2c	<ul style="list-style-type: none"> i. Select whether the destination will be used for traps or informs. ii. In the Host field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap. Only UDP is supported. iii. Optionally, enter the Port that will receive the trap. The default is port 162. iv. Use the default trap community, if one was specified on the SNMP Agent page, or enter a custom community string for this trap destination. The custom community string can be a maximum of 32 characters and cannot contain whitespace.
SNMPv3	<ul style="list-style-type: none"> i. Select whether the destination will be used for traps or informs. ii. In the Host field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap. Only UDP is supported. iii. Optionally, enter the Port that will receive the trap. The default is port 162. iv. Select the USM user that will be used for authentication. <ul style="list-style-type: none"> • If you selected Trap, only USM users without authoritative engine IDs are shown. • If you selected Inform, only USM users with authoritative engine IDs are shown.

- d. Click **Create**.
The trap destination is created and added to the table.



11. When you have completed the SNMP agent configuration, click **Save**
The new SNMP agent configuration becomes active.

Related tasks

[Silencing alert notifications](#) on page 87

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

Updating 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

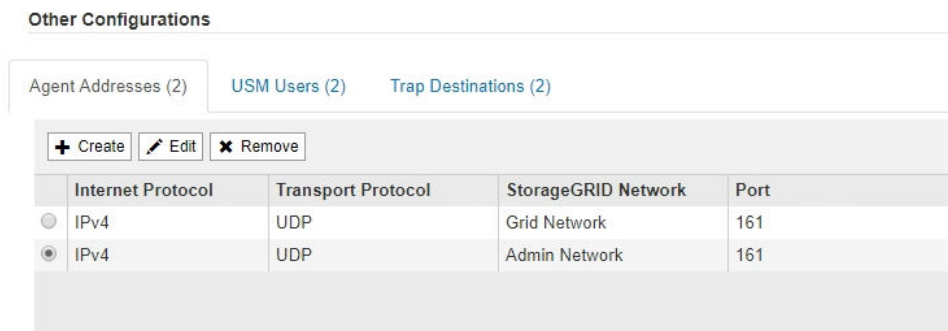
- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

About this task

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

Steps

1. Select **Configuration**. Then, in the Monitoring section of the menu, select **SNMP Agent**.
The SNMP Agent page appears.
2. If you want to disable the SNMP agent on all grid nodes, unselect the **Enable SNMP** check box, and click **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, unselect the **Enable SNMP Agent Notifications** check box if you no longer want the StorageGRID SNMP agent to send trap and inform notifications.
When this check box is unselected, the SNMP agent supports read-only MIB access, but it does not send any SNMP notifications.
5. Optionally, unselect the **Enable Authentication Traps** check box 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 do not apply to SNMPv3.
Note: 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, click **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 click **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 click **Remove**. Then, click **OK** to confirm that you want to remove this address.
 - d. To commit your changes, click **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.

Other Configurations

Agent Addresses (2) **USM Users (3)** Trap Destinations (2)

	Username	Read-Only MIB Access	Security Level	Authoritative Engine ID
<input type="radio"/>	user2	<input checked="" type="checkbox"/>	authNoPriv	
<input type="radio"/>	user1		authNoPriv	B3A73C2F3D6
<input checked="" type="radio"/>	user3		authPriv	59D39E801256

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

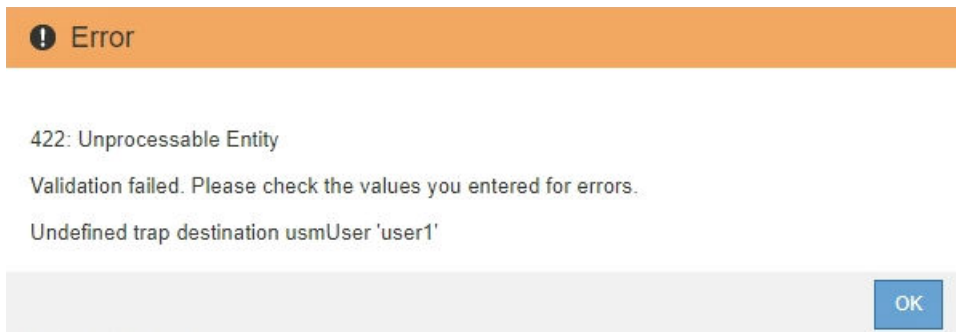
- a. To add a USM user, click **Create**. Then, refer to the step for USM users in the instructions for configuring the SNMP agent.
- b. To edit a USM user, select the radio button for the user, and click **Edit**. Then, refer to the step for USM users in the instructions for configuring the SNMP agent.

The username for an existing USM user cannot be changed. If you need to change a username, you must remove the user and create a new one.

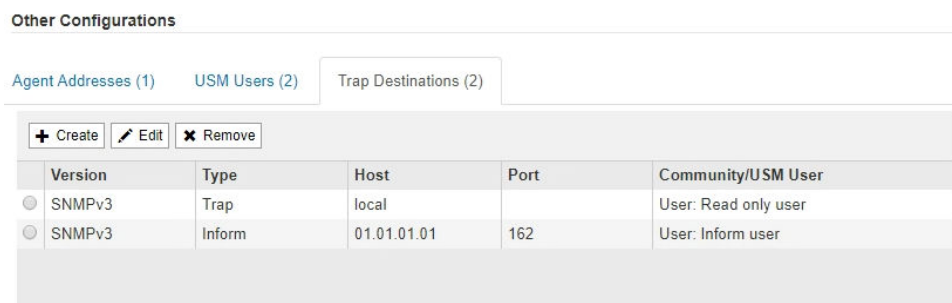
Note: If you add or remove a user's authoritative engine ID and that user is currently selected for a destination, you must edit or remove the destination, as described in step 9 on page 120. Otherwise, a validation error occurs when you save the SNMP agent configuration.

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

Note: If the user you removed is currently selected for a trap destination, you must edit or remove the destination, as described in step 9 on page 120. Otherwise, a validation error occurs when you save the SNMP agent configuration.



- d. To commit your changes, click **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 click **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, click **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 click **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 click **Remove**. Then, click **OK** to confirm that you want to remove this destination.
 - d. To commit your changes, click **Save** at the bottom of the SNMP Agent page.
10. When you have updated the SNMP agent configuration, click **Save**.

Related tasks

[Configuring the SNMP agent](#) on page 111

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.

Collecting additional StorageGRID data

There are a number of additional ways to collect and analyze data that can be useful when investigating the state of your StorageGRID system or when working with technical support to resolve issues.

Choices

- [Using charts and reports](#) on page 121
- [Monitoring PUT and GET performance](#) on page 130
- [Monitoring object verification operations](#) on page 131
- [Monitoring events](#) on page 132
- [Reviewing audit messages](#) on page 138
- [Collecting log files and system data](#) on page 139
- [Triggering an AutoSupport message](#) on page 141
- [Viewing the Grid Topology tree](#) on page 141
- [Reviewing support metrics](#) on page 142
- [Running diagnostics](#) on page 144
- [Creating custom monitoring applications](#) on page 147


Using charts and reports

You can use charts and reports to monitor the state of the StorageGRID system and troubleshoot problems. The types of charts and reports available in the Grid Manager include pie charts (on the Dashboard only), graphs, and text reports.

Types of charts and graphs

Charts and graphs summarize the values of specific StorageGRID metrics and attributes.

The Grid Manager Dashboard includes pie (doughnut) charts to summarize available storage for the grid and each site.



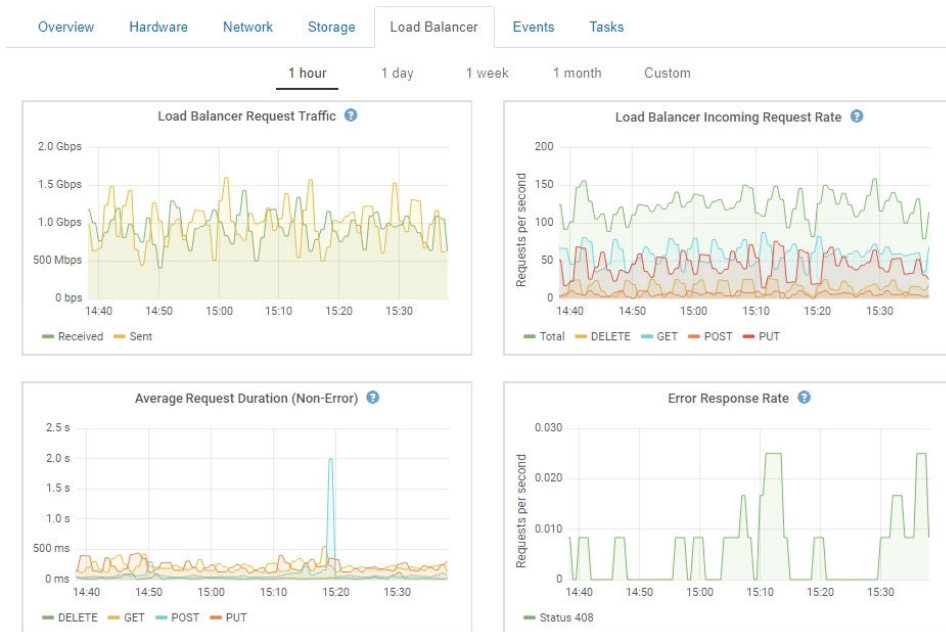
The Tenant Manager Dashboard includes a pie chart to show which S3 buckets or Swift containers are consuming the most storage and a pie chart to show how much of the tenant's quota has been used.

In addition, graphs that show how StorageGRID metrics and attributes change over time are available from the Nodes page and from the **Support > 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 > Load Balancer** tab for an Admin Node includes four Grafana charts.


DC1-SG1000-ADM (Admin Node)




Note: Grafana charts are also included on the pre-constructed dashboards available from the **Support > Metrics** page.


- **Line graphs:** Available from the Nodes page and from the **Support > Grid Topology** page (click 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 > Grid Topology** page (click 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.



- **State graph:** Available from the **Support > Grid Topology** page (click 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 concepts

[Viewing the Nodes page](#) on page 8

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.

[Viewing the Grid Topology tree](#) on page 141

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.







Related tasks

[Reviewing support metrics](#) on page 142

When troubleshooting an issue, you can work with technical support to review detailed metrics and charts for your StorageGRID system.

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.

Displaying charts and graphs

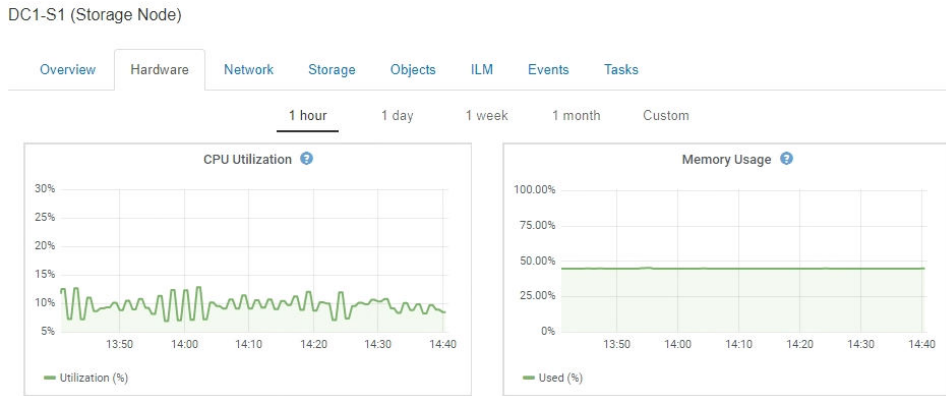
The Nodes page contains the graphs and charts you should access regularly to monitor attributes such as storage capacity and throughput. In some cases, especially when working with technical support, you can use the **Support > Grid Topology** page to access additional charts.

Before you begin


You must be signed in to the Grid Manager using a supported 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** page for node includes two Grafana charts.



3. If a Grafana chart is displayed, optionally hover your cursor over the chart to see more detailed values for a particular point in time.

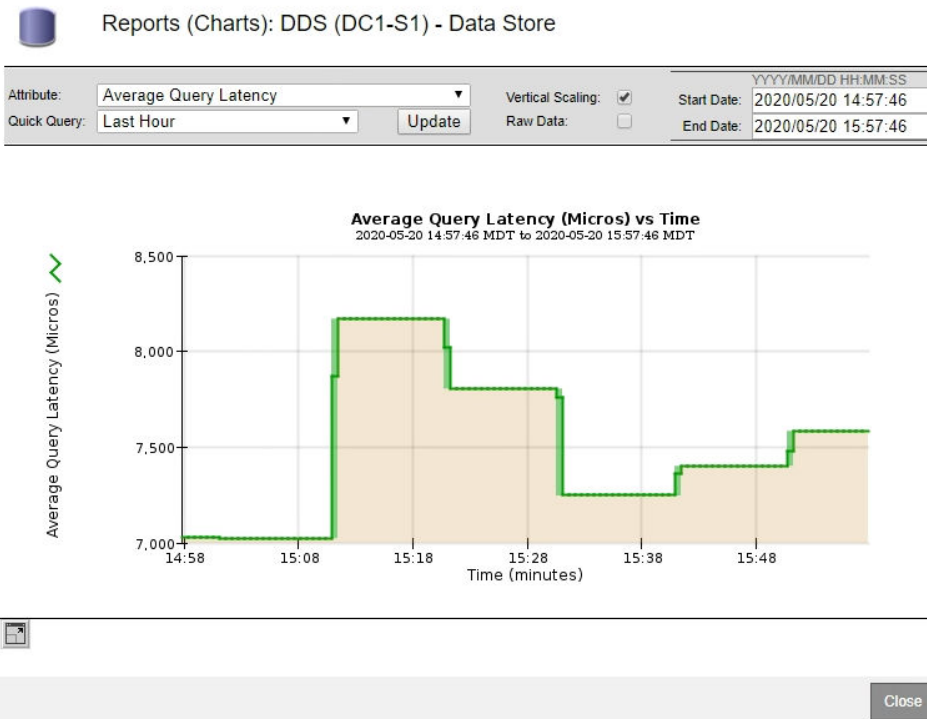


4. If a Grafana chart is not available, you can often display a chart for a specific attribute or metric. From the table on the Nodes page, click the chart icon  to the right of the attribute name.

Note: Charts are not available for all metrics and attributes.

For example, from the Objects tab for a Storage Node, you can click the chart icon to see the average latency for a metadata query over time.

Queries		
Average Latency	14.43 milliseconds	
Queries - Successful	19,786	
Queries - Failed (timed-out)	0	
Queries - Failed (consistency level unmet)	0	



- To display charts for attributes that are not shown on the Node page, select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
- Select *grid node* > *component or service* > **Overview** > **Main**.

Overview | Alarms | Reports | Configuration

Main

Overview: SSM (DC1-ADM1) - Resources
 Updated: 2018-05-07 16:29:52 MDT

Computational Resources


Service Restarts:	1	
Service Runtime:	6 days	
Service Uptime:	6 days	
Service CPU Seconds:	10666 s	
Service Load:	0.266 %	

Memory

Installed Memory:	8.38 GB	
Available Memory:	2.9 GB	

Processors

Processor Number	Vendor	Type	Cache
1	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
2	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
3	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
4	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
5	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
6	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
7	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
8	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB

7. Click 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.

Generating 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 browser.
- You must have specific access permissions.

Steps

1. Select **Support**. Then, in the Tools section of the menu, select **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, deselect the **Vertical Scaling** check box.
5. To show values at full precision, select the **Raw Data** check box, or to round values to a maximum of three decimal places (for example, for attributes reported as percentages), deselect the **Raw Data** check box.
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/DD HH: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 chart 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.
9. If you want to print the chart, right-click and select **Print**, and modify any necessary printer settings and click **Print**.

Types of 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

Generating 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 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**. Then, in the Tools section of the menu, select **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), unselect the **Raw Data** check box.
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/DD HH: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.
9. If you want to print the report, right-click and select **Print**, and modify any necessary printer settings and click **Print**.


Exporting 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**. Then, in the Tools section of the menu, select **Grid Topology**.
2. Create a text report.
3. Click **Export** .

The screenshot shows the 'Reports (Text): SSM (170-176) - Events' window. The 'Attribute' is set to 'Send to Relay Rate'. The 'Start Date' is 2010/07/19 08:42:09 and the 'End Date' is 2010/07/20 08:42:09. The table below shows the following data:

Time Received	Sample Time	Value
2010-07-20 08:40:46	2010-07-20 08:40:46	0.274981485 Messages/s
2010-07-20 08:38:46	2010-07-20 08:38:46	0.274989 Messages/s
2010-07-20 08:36:46	2010-07-20 08:36:46	0.283317543 Messages/s
2010-07-20 08:34:46	2010-07-20 08:34:46	0.274982493 Messages/s
2010-07-20 08:32:46	2010-07-20 08:32:46	0.291646426 Messages/s

Navigation: Previous « 1 2 3 4 5 » Next

The Export Text Report window opens displaying the report.

```

Grid ID: 000000
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.

Monitoring 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 may require investigation.

Choices

- You can download S3tester from github:

Related information

[Understanding audit messages](#)

[s3tester: S3 Performance Benchmarking](#)

Monitoring 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 browser.

About this task

There are two verification processes that 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.

Note: The **Unidentified corrupt object detected** alert is triggered if the system detects a corrupt object that cannot be corrected automatically.












- **Foreground verification** can be triggered by a user to more quickly verify the existence (although not the correctness) of object data.

Foreground verification allows you to verify the existence of replicated and erasure-coded object data on a specific Storage Node, checking that each object that is expected to be present is there. You can run foreground verification on all or some of a Storage Node's object stores to help determine if there are integrity problems with a storage device. Large numbers of missing objects might indicate that there is an issue with storage.

To review results from background and foreground verifications, such as corrupt or missing objects, you can look at the Nodes page for a Storage Node. You should investigate any instances of corrupt or missing object data immediately, to determine the root cause.







Steps

1. Select **Nodes**.
2. Select *Storage Node* > **Objects**.
3. To check the verification results:
 - To check replicated object data verification, look at the attributes in the Verification section.

Verification		
Status	No Errors	
Rate Setting	Adaptive	
Percent Complete	0.00%	
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	

Note: Click an attribute's name in the table to display help text.

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

Erasure Coding Verification		
Status	Idle	
Next Scheduled	2019-03-01 14:20:29 MST	
Fragments Verified	0	
Data Verified	0 bytes	
Corrupt Copies	0	
Corrupt Fragments	0	
Missing Fragments	0	

Note: Click an attribute's name in the table to display help text.

Related concepts

[Verifying object integrity](#) on page 159

The StorageGRID system verifies the integrity of object data on Storage Nodes, checking for both corrupt and missing objects.

Monitoring 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.

The SMTT (Total events) alarm can be repeatedly triggered by issues such as network problems, power outages or upgrades. This section has information on 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.

Reviewing events from the Nodes page

The Nodes page lists the system events for each grid node.

1. Select **Nodes**.
2. Select *grid node* > **Events**.
3. At the top of the page, determine if an event is shown for **Last Event**, which describes the last event detected by the grid node.
The event is relayed verbatim from the grid node and includes any log messages with a severity level of ERROR or CRITICAL.
4. Review the table to see if the Count for any event or error is not zero.
5. After resolving issues, click **Reset event counts** to return the counts to zero.

Reviewing events from the Grid Topology page

The Grid Topology page also lists the system events for each grid node.

1. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
2. Select *site* > *grid node* > **SSM** > **Events** > **Overview** > **Main**.

Related concepts

[Log files reference](#) on page 246

The following sections list the logs 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.

Related tasks

[Resetting event counts](#) on page 134

After resolving system events, you can reset event counts to zero.

Reviewing previous events

You can generate a list of previous event messages to help isolate issues that occurred in the past.

Steps

1. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
2. Select *site* > *grid node* > **SSM** > **Events** > **Reports**.
3. Select **Text**.
The **Last Event** attribute is not shown in the Charts view.
4. Change **Attribute** to **Last Event**.
5. Optionally, select a time period for **Quick Query**.
6. Click **Update**.

The screenshot shows the 'Reports' tab in the StorageGRID interface. The main heading is 'Reports (Text): SSM (170-41) - Events'. Below this is a control panel with the following elements:

- Attribute: Last Event (dropdown)
- Quick Query: Last 5 Minutes (dropdown)
- Update button
- Results Per Page: 20 (dropdown)
- Raw Data:
- Start Date: 2009/04/15 15:19:53 (YYYY/MM/DD HH:MM:SS)
- End Date: 2009/04/15 15:24:53 (YYYY/MM/DD HH:MM:SS)

The main content area is titled 'Text Results for Last Event' and shows a time range from 2009-04-15 15:19:53 PDT to 2009-04-15 15:24:53 PDT. A pagination indicator shows '1 - 2 of 2'. Below this is a table with the following data:

Time Received	Sample Time	Value
2009-04-15 15:24:22	2009-04-15 15:24:22	hdc: task_no_data_intr: status=0x51 { DriveReady SeekComplete Error }
2009-04-15 15:24:11	2009-04-15 15:23:39	hdc: task_no_data_intr: status=0x51 { DriveReady SeekComplete Error }

Related concepts

[Using charts and reports](#) on page 121

You can use charts and reports to monitor the state of the StorageGRID system and troubleshoot problems. The types of charts and reports available in the Grid Manager include pie charts (on the Dashboard only), graphs, and text reports.

Resetting event counts

After resolving system events, you can reset event counts to zero.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Grid Topology Page Configuration permission.

Steps

1. Select **Nodes** > *Grid Node* > **Events**.
2. Make sure that any event with a count greater than 0 has been resolved.
3. Click **Reset event counts**.

Events ?	
Last Event	No Events
Description	Count
Abnormal Software Events	0
Account Service Events	0
Cassandra Heap Out Of Memory Errors	0
Cassandra unhandled exceptions	0
Chunk Service Events	0
Custom Events	0
Data-Mover Service Events	0
File System Errors	0
Forced Termination Events	0
Hotfix Installation Failure Events	0
I/O Errors	0
IDE Errors	0
Identity Service Events	0
Kernel Errors	0
Kernel Memory Allocation Failure	0
Keystone Service Events	0
Network Receive Errors	0
Network Transmit Errors	0
Node Errors	0
Out Of Memory Errors	0
Replicated State Machine Service Events	0
SCSI Errors	0
Stat Service Events	0
Storage Hardware Events	0
System Time Events	0

Reset event counts

Creating 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.

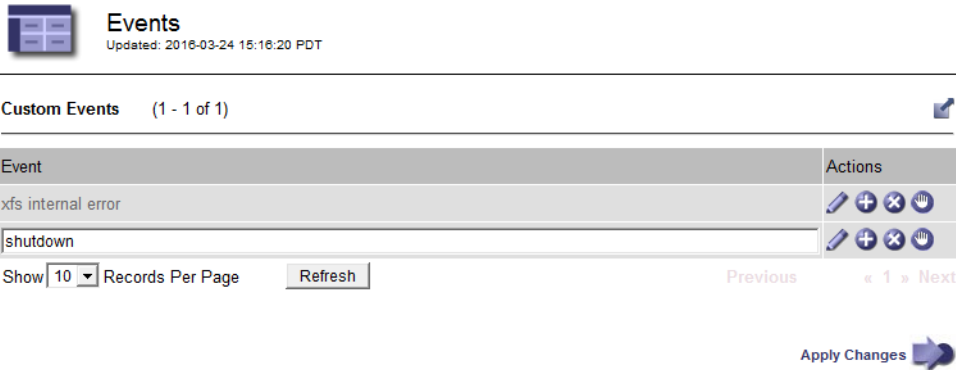
- After a custom event is created, every occurrence of it is monitored. You can view a cumulative Count value for all custom events on the **Nodes > grid node > Events** page.
- To create a custom event based on keywords in the `/var/log/messages` or `/var/log/syslog` files, the logs in those files must be:

- Generated by the kernel
- Generated by daemon or user program at the error or critical level

Note: Not all entries in the `/var/log/messages` or `/var/log/syslog` files will be matched unless they satisfy the requirements stated above.

Steps

1. Select **Configuration**. Then, in the Monitoring section of the menu, select **Events**.
2. Click **Edit**  (or **Insert**  if this is not the first event).
3. Enter a custom event string, for example, shutdown



























4. Click **Apply Changes**.
5. Select **Nodes**. Then, select `grid node > Events`.
6. 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.

Overview Hardware Network Storage **Events**

Events ⓘ

Last Event No Events

Description	Count	
Abnormal Software Events	0	
Account Service Events	0	
Cassandra Heap Out Of Memory Errors	0	
Cassandra unhandled exceptions	0	
Custom Events	0	
File System Errors	0	
Forced Termination Events	0	
Hotfix Installation Failure Events	0	
I/O Errors	0	
IDE Errors	0	
Identity Service Events	0	
Kernel Errors	0	
Kernel Memory Allocation Failure	0	
Keystone Service Events	0	
Network Receive Errors	0	
Network Transmit Errors	0	
Node Errors	0	
Out Of Memory Errors	0	
Replicated State Machine Service Events	0	
SCSI Errors	0	
Stat Service Events	0	
Storage Hardware Events	0	
System Time Events	0	

[Reset event counts](#) 

Resetting the count of custom events to zero

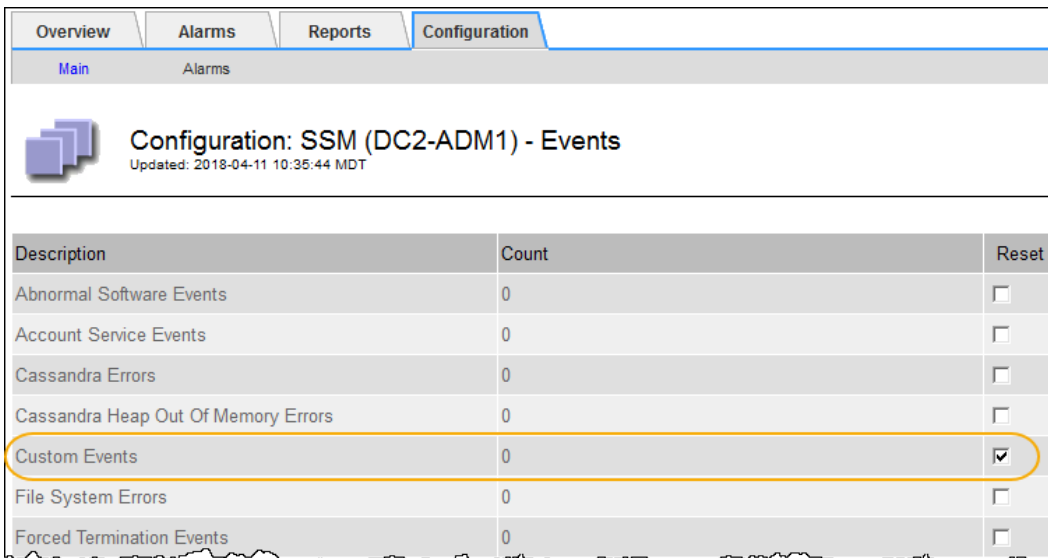
If you want to reset the counter only for custom events, you must use the Grid Topology page in the Support menu.

About this task

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**. Then, in the Tools section of the menu, select **Grid Topology**.
2. Select *grid node* > **SSM** > **Events** > **Configuration** > **Main**.
3. Select the **Reset** check box for Custom Events.



The screenshot shows the 'Configuration' tab for 'SSM (DC2-ADM1) - Events'. The page has a navigation bar with 'Overview', 'Alarms', 'Reports', and 'Configuration'. Below the navigation bar, there are sub-tabs for 'Main' and 'Alarms'. The main content area shows a table with the following data:

Description	Count	Reset
Abnormal Software Events	0	<input type="checkbox"/>
Account Service Events	0	<input type="checkbox"/>
Cassandra Errors	0	<input type="checkbox"/>
Cassandra Heap Out Of Memory Errors	0	<input type="checkbox"/>
Custom Events	0	<input checked="" type="checkbox"/>
File System Errors	0	<input type="checkbox"/>
Forced Termination Events	0	<input type="checkbox"/>

4. Click **Apply Changes**.

Reviewing audit messages

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

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

- System audit messages are related to the auditing system itself, grid node states, system-wide task activity, and service backup operations.
- Object storage audit messages are related to the storage and management of objects within StorageGRID, including object storage and retrievals, grid-node to grid-node transfers, and verifications.
- Client read and write audit messages are logged when an S3 or Swift client application makes a request to create, modify, or retrieve an object.
- Management audit messages log user requests to the Management API.

Each Admin Node stores audit messages in text files. The audit share contains the active file (`audit.log`) as well as compressed audit logs from previous days.

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

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

Related information

[Understanding audit messages](#)

[Administering StorageGRID](#)

Collecting 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 using a supported 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.

Because application log files can be very large, the destination directory where you download the archived log files must have at least 1 GB of free space.

Steps

1. Select **Support**. Then, in the Tools section of the menu, select **Logs**.

Logs

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

StorageGRID Webscale Deployment

- ▲ ▲ Data Center 1
 - DC1-ADM1
 - ▲ DC1-ARC1
 - DC1-G1
 - DC1-S1
 - DC1-S2
 - DC1-S3
- ▲ Data Center 2
 - DC2-ADM1
 - DC2-S1
 - DC2-S2
 - DC2-S3
- ▲ Data Center 3
 - DC3-S1
 - DC3-S2
 - DC3-S3

Log Start Time: 2018-04-18 01 : 38 PM MDT

Log End Time: 2018-04-18 05 : 38 PM MDT

Notes

Provisioning Passphrase

Collect Logs

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. Optionally type 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.
5. Enter the provisioning passphrase for your StorageGRID system in the **Provisioning Passphrase** text box.

6. Click Collect Logs.

When you submit a new request, the previous collection of log files is deleted.

Logs

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

Log collection is in progress.

Last Collected

Log Start Time 2017-05-17 05:01:00 PDT

Log End Time 2017-05-18 09:01:00 PDT

Notes
Issues began approximately 7am on the 17th, then multiple alarms propagated throughout the grid.

23% Collecting logs: 10 of 13 nodes remaining

Download Delete

Name	Status
DC1-ADM1	Complete
DC1-G1	Error: No route to host - connect(2) for "10.96.104.212" port 22
DC1-S1	Collecting
DC1-S2	Collecting
DC1-S3	Collecting
DC2-S1	Collecting
DC2-S2	Collecting
DC2-S3	Collecting

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.

7. Click 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 click **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.

Related concepts

[Log files reference](#) on page 246

The following sections list the logs 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.

Triggering an AutoSupport message

To assist technical support in troubleshooting issues with your StorageGRID system, you can manually trigger the sending of an AutoSupport message.

Before you begin

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

Steps

1. Select **Support**. Then, in the Tools section of the menu, select **AutoSupport**.
2. From the AutoSupport menu, select **User-triggered**.
3. Click **Send**.



Last Attempt	Successful
Last Successful Time	2016-02-22 13:50:53 PST

Send

StorageGRID attempts to send an AutoSupport message to technical support. If the attempt is successful, the Last Attempt attribute updates to Successful. If there is a problem, the Last Attempt attribute updates to Failed, and StorageGRID does not try again.

If a failure occurs and SMTP is the selected protocol, verify that the StorageGRID system's email server is correctly configured and that your email server is running. The following error message might appear on the AutoSupport page: `AutoSupport messages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.`

Related tasks

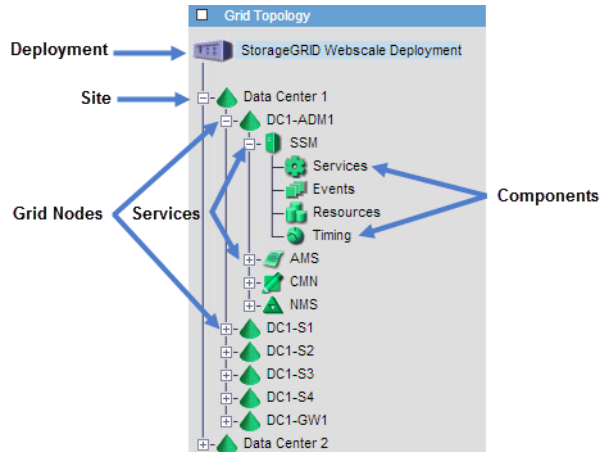
[Configuring email server settings for alarms \(legacy system\)](#) on page 103



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 cannot receive email.

Viewing 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**. Then, in the Tools section of the menu, select **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.

Reviewing 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 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.



Attention: 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.

Steps

1. As directed by technical support, select **Support**. Then, in the Tools section of the menu, select **Metrics**.
The Metrics page appears.

Metrics

Access charts and metrics to help troubleshoot issues.

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

Prometheus

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

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

- [https://\[IP\]/metrics/graph](https://[IP]/metrics/graph)

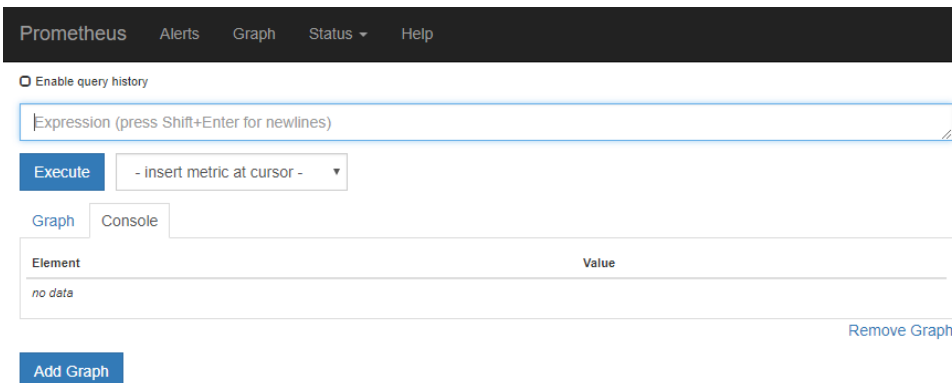
Grafana

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

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

- | | |
|---|---|
| ADE | Node |
| Account Service Overview | Node (Internal Use) |
| Alertmanager | Platform Services Commits |
| Audit Overview | Platform Services Overview |
| Cassandra Cluster Overview | Platform Services Processing |
| Cassandra Network Overview | Renamed Metrics |
| Cassandra Node Overview | Replicated Read Path Overview |
| Cloud Storage Pool Overview | S3 - Node |
| EC Read - Node | S3 Overview |
| EC Read - Overview | Site |
| Grid | Streaming EC - ADE |
| ILM Metrics | Streaming EC - Chunk Service |
| Identity Service Overview | Support Metrics |
| Ingests | |

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.



Note: 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.



Related reference

[Commonly used Prometheus metrics](#) on page 217

The Prometheus service on Admin Nodes collects time series metrics from the services on all nodes. While Prometheus collects more than a thousand metrics, a relatively small number are required to monitor the most critical StorageGRID operations.

Running diagnostics


When troubleshooting an issue, you can work with technical support to run diagnostics on your StorageGRID system and review the results.



Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must 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**. Then, in the Tools section of the menu, select **Diagnostics**.

The Diagnostics page appears and lists the results for each diagnostic check. In the example, all diagnostics have a Normal status.

Diagnostics

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

- ✓ **Normal:** All values are within the normal range.
- ⚠ **Attention:** One or more of the values are outside of the normal range.
- ✖ **Caution:** One or more of the values are significantly outside of the normal range.

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

[Run Diagnostics](#)

✓ Cassandra request timeouts	▼
✓ Cassandra requests unable to achieve consistency	▼
✓ CPU utilization	▼
✓ Dirty page ratio	▼
✓ Disk read latency	▼
✓ Disk write latency	▼
✓ Grid options	▼
✓ Load balancer - request duration	▼
✓ Load balancer - upstream connection problems	▼
✓ Load balancer - upstream retries	▼
✓ Network MTU values	▼
✓ Node uptime	▼
✓ Storage used - object data	▼
✓ TCP connection tracking count	▼
✓ TCP retransmission rate	▼

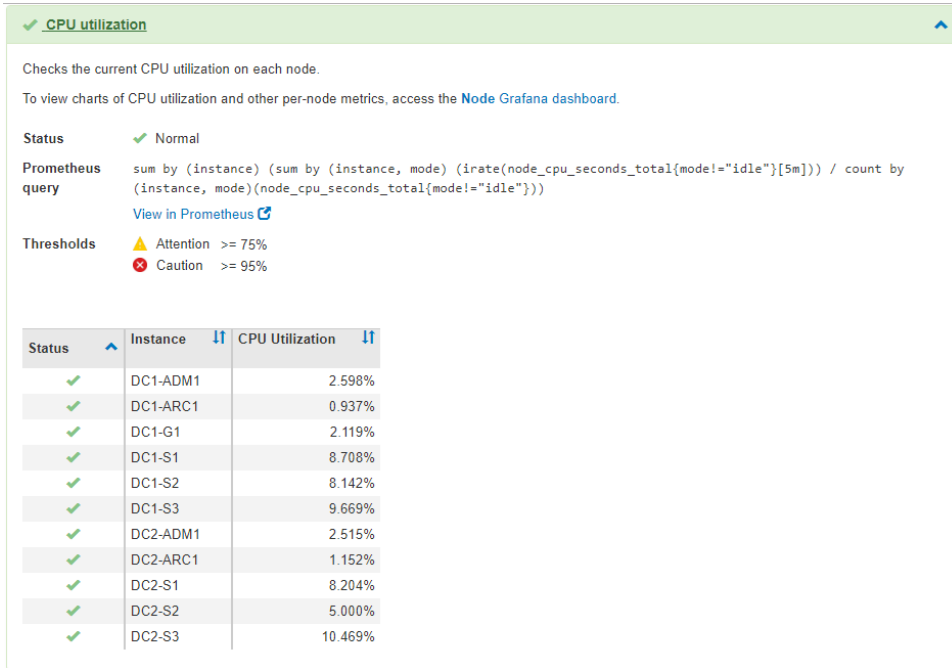
2. To learn more about a specific diagnostic, click anywhere in the row. Details about the diagnostic and its current results appear. The following details are listed:

- **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 are not used for all diagnostics.)

Note: You cannot 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.



- Optional: To see Grafana charts related to this diagnostic, click the **Grafana dashboard** link.

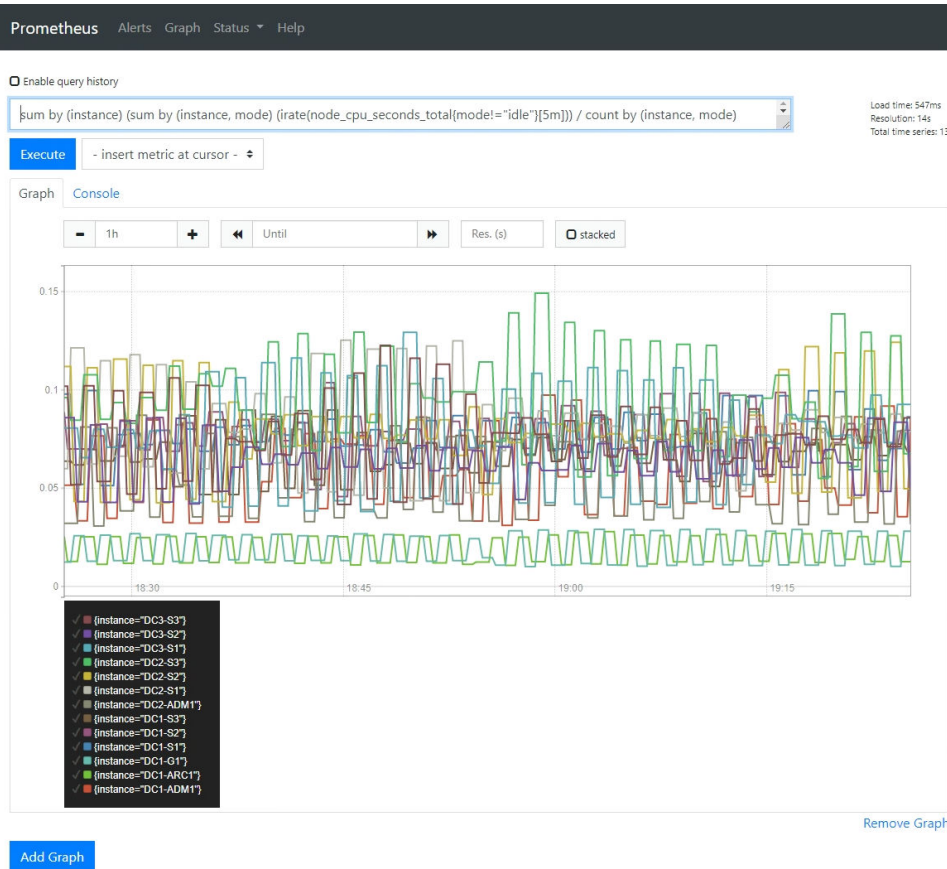
This link is not displayed for all diagnostics.

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.

Note: You can also access the pre-constructed Grafana dashboards from the Grafana section of the **Support > Metrics** page.



- 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.



Related tasks

[Reviewing support metrics](#) on page 142

When troubleshooting an issue, you can work with technical support to review detailed metrics and charts for your StorageGRID system.

Related reference

[Commonly used Prometheus metrics](#) on page 217

The Prometheus service on Admin Nodes collects time series metrics from the services on all nodes. While Prometheus collects more than a thousand metrics, a relatively small number are required to monitor the most critical StorageGRID operations.

Creating 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 are not 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.

To view the metrics API operations, including the complete list of the metrics that are available, go to the Grid Manager and select **Help > API Documentation > metrics**.

metrics Operations on metrics ▼

GET	<code>/grid/metric-labels/{label}/values</code> Lists the values for a metric label	🔒
GET	<code>/grid/metric-names</code> Lists all available metric names	🔒
GET	<code>/grid/metric-query</code> Performs an instant metric query at a single point in time	🔒
GET	<code>/grid/metric-query-range</code> Performs a metric query over a range of time	🔒

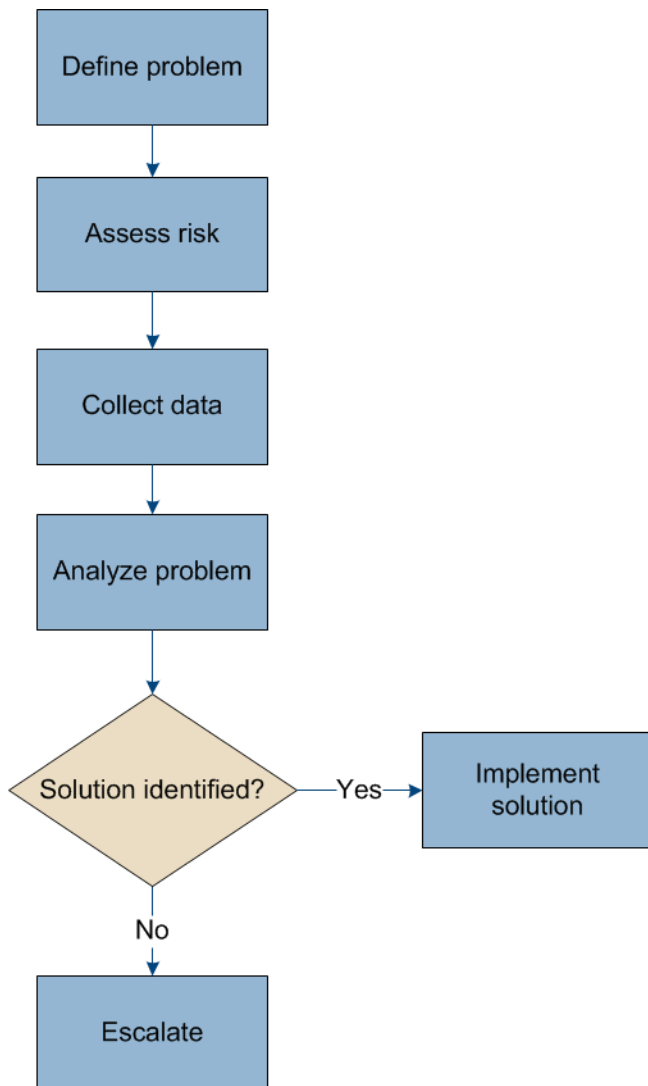
The details of how to implement a custom monitoring application is beyond the scope of this guide.

Troubleshooting a StorageGRID system

If you encounter a problem when using a StorageGRID system, refer to the tips and guidelines in this section for help in determining and resolving the issue.

Overview of problem determination

If you encounter a problem when administering a StorageGRID system, you can use the process outlined in this figure to identify and analyze the issue. In many cases, you can resolve problems on your own; however, you might need to escalate some issues to technical support.



Defining the problem

The first step to solving a problem is to define the problem clearly.

This table provides examples of the types of information that you might collect to define a problem:

Question	Sample response
What is the StorageGRID system doing or not doing? What are its symptoms?	Client applications are reporting that objects cannot be ingested into StorageGRID.
When did the problem start?	Object ingest was first denied at about 14:50 on January 8, 2020.
How did you first notice the problem?	Notified by client application. Also received alert email notifications.
Does the problem happen consistently, or only sometimes?	Problem is ongoing.
If the problem happens regularly, what steps cause it to occur	Problem happens every time a client tries to ingest an object.
If the problem happens intermittently, when does it occur? Record the times of each incident that you are aware of.	Problem is not intermittent.
Have you seen this problem before? How often have you had this problem in the past?	This is the first time I have seen this issue.

Assessing the risk and impact on the system

After you have defined the problem, assess its risk and impact on the StorageGRID system. For example, the presence of critical alerts does not necessarily mean that the system is not delivering core services.

This table summarizes the impact the example problem is having on system operations:

Question	Sample response
Can the StorageGRID system ingest content?	No.
Can client applications retrieve content?	Some objects can be retrieved and others cannot.
Is data at risk?	No.
Is the ability to conduct business severely affected?	Yes, because client applications cannot store objects to the StorageGRID system and data cannot be retrieved consistently.

Collecting data

After you have defined the problem and have assessed its risk and impact, collect data for analysis. The type of data that is most useful to collect depends upon the nature of the problem.

Type of data to collect	Why collect this data	Instructions
Create timeline of recent changes	Changes to your StorageGRID system, its configuration, or its environment can cause new behavior.	<ul style="list-style-type: none"> Creating a timeline of recent changes on page 152

Type of data to collect	Why collect this data	Instructions
Review alerts and alarms	<p>Alerts and alarms can help you quickly determine the root cause of a problem by providing important clues as to the underlying issues that might be causing it.</p> <p>Review the list of current alerts and alarms to see if StorageGRID has identified the root cause of a problem for you.</p> <p>Review alerts and alarms triggered in the past for additional insights.</p>	<ul style="list-style-type: none"> • Viewing current alerts on page 39 • Viewing legacy alarms on page 46 • Viewing resolved alerts on page 41 • Reviewing historical alarms and alarm frequency (legacy system) on page 96
Monitor events	<p>Events include any system error or fault events for a node, including errors such as network errors. Monitor events to learn more about issues or to help with troubleshooting.</p>	<ul style="list-style-type: none"> • Viewing the Events tab on page 15 • Monitoring events on page 132
Identify trends, using chart and text reports	<p>Trends can provide valuable clues about when issues first appeared, and can help you understand how quickly things are changing.</p>	<ul style="list-style-type: none"> • Using charts and reports on page 121
Establish baselines	<p>Collect information about the normal levels of various operational values. These baseline values, and deviations from these baselines, can provide valuable clues.</p>	<ul style="list-style-type: none"> • Establishing baselines on page 153
Perform ingest and retrieval tests	<p>To troubleshoot performance issues with ingest and retrieval, use a workstation to store and retrieve objects. Compare results against those seen when using the client application.</p>	<ul style="list-style-type: none"> • Monitoring PUT and GET performance on page 130
Review audit messages	<p>Review audit messages to follow StorageGRID operations in detail. The details in audit messages can be useful for troubleshooting many types of issues, including performance issues.</p>	<ul style="list-style-type: none"> • Reviewing audit messages on page 138
Check object locations and storage integrity	<p>If you are having storage problems, verify that objects are being placed where you expect. Check the integrity of object data on a Storage Node.</p>	<ul style="list-style-type: none"> • Monitoring object verification operations on page 131 • Confirming object data locations on page 155 • Verifying object integrity on page 159
Collect data for technical support	<p>Technical support might ask you to collect data or review specific information to help troubleshoot issues.</p>	<ul style="list-style-type: none"> • Collecting log files and system data on page 139 • Triggering an AutoSupport message on page 141 • Reviewing support metrics on page 142

Creating a timeline of recent changes

When a problem occurs, you should consider what has changed recently and when those changes occurred.

- Changes to your StorageGRID system, its configuration, or its environment can cause new behavior.
- A timeline of changes can help you identify which changes might be responsible for an issue, and how each change might have affected its development.

Create a table of recent changes to your system that includes information about when each change occurred and any relevant details about the change, such information about what else was happening while the change was in progress:

Time of change	Type of change	Details
For example: <ul style="list-style-type: none"> • When did you start the node recovery? • When did the software upgrade complete? • Did you interrupt the process? 	What happened? What did you do?	Document any relevant details about the change. For example: <ul style="list-style-type: none"> • Details of the network changes. • Which hotfix was installed. • How client workloads changed. Make sure to note if more than one change was happening at the same time. For example, was this change made while an upgrade was in progress? Examples of significant recent changes on page 152

Examples of significant recent changes

Here are some examples of potentially significant changes:

- Was the StorageGRID system recently installed, expanded, or recovered?
- Has the system been upgraded recently? Was a hotfix applied?
- Has any hardware been repaired or changed recently?
- Has the ILM policy been updated?
- Has the client workload changed?
- Has the client application or its behavior changed?
- Have you changed load balancers, or added or removed a high availability group of Admin Nodes or Gateway Nodes?
- Have any tasks been started that might take a long time to complete? Examples include:
 - Recovery of a failed Storage Node
 - Storage Node decommissioning
- Have any changes been made to user authentication, such as adding a tenant or changing LDAP configuration?
- Is data migration taking place?
- Were platform services recently enabled or changed?
- Was compliance enabled recently?
- Have Cloud Storage Pools been added or removed?
- Have any changes been made to storage compression or encryption?
- Have there been any changes to the network infrastructure? For example, VLANs, routers, or DNS.
- Have any changes been made to NTP sources?

- Have any changes been made to the Grid, Admin, or Client Network interfaces?
- Have any configuration changes been made to the Archive Node?
- Have any other changes been made to the StorageGRID system or its environment?

Establishing baselines

You can establish baselines for your system by recording the normal levels of various operational values. In the future, you can compare current values to these baselines to help detect and resolve abnormal values.

Property	Value	How to obtain
Average storage consumption	_____ GB consumed/day _____ % consumed/day	Go to the Grid Manager. On the Nodes page, select the entire grid or a site and go to the Storage tab. On the Storage Used - Object Data chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate how much storage is consumed each day You can collect this information for the entire system or for a specific data center.
Average metadata consumption	_____ GB consumed/day _____ % consumed/day	Go to the Grid Manager. On the Nodes page, select the entire grid or a site and go to the Storage tab. On the Storage Used - Object Metadata chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate how much metadata storage is consumed each day You can collect this information for the entire system or for a specific data center.
Rate of S3/Swift operations	_____ operations/second	Go to the Dashboard in the Grid Manager. In the Protocol Operations section, view the values for S3 rate and the Swift rate. To see ingest and retrieval rates and counts for a specific site or node, select Nodes > site or Storage Node > Objects . Hover your cursor over the Ingest and Retrieve chart for S3 or Swift.
Failed S3/Swift operations	_____	Select Support . Then, in the Tools section of the menu, select Grid Topology . On the Overview tab in the API Operations section, view the value for S3 Operations - Failed or Swift Operations - Failed.
ILM evaluation rate	_____ objects/ second	From the Nodes page, select grid > ILM . On the ILM Queue chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate a baseline value for Evaluation rate for your system.
ILM scan rate	_____ objects/ second	Select Nodes > grid > ILM . On the ILM Queue chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate a baseline value for Scan rate for your system.
Objects queued from client operations	_____ objects/ second	Select Nodes > grid > ILM . On the ILM Queue chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate a baseline value for Objects queued (from client operations) for your system.

Property	Value	How to obtain
Average query latency	_____ milliseconds	Select Nodes > <i>Storage Node</i> > Objects . In the Queries table, view the value for Average Latency.

Analyzing data


Use the information that you collect to determine the cause of the problem and potential solutions.

The analysis is problem-dependent, but in general:

- Locate points of failure and bottlenecks using the alarms.
- Reconstruct the problem history using the alarm history and charts.
- Use charts to find anomalies and compare the problem situation with normal operation.

Escalation information checklist

If you cannot resolve the problem on your own, contact technical support. Before contacting technical support, gather the information listed in the following table to facilitate problem resolution.

	Item	Notes
	Problem statement	What are the problem symptoms? When did the problem start? Does it happen consistently or intermittently? If intermittently, what times has it occurred? <i>Defining the problem</i> on page 149
	Impact assessment	What is the severity of the problem? What is the impact to the client application? <ul style="list-style-type: none"> • Has the client connected successfully before? • Can the client ingest, retrieve, and delete data?
	StorageGRID System ID	Select Maintenance . Then, in the System section of the menu, select License . The StorageGRID System ID is shown as part of the current license.
	Software version	Click Help > About to see the StorageGRID version.
	Customization	Summarize how your StorageGRID system is configured. For example, list the following: <ul style="list-style-type: none"> • Does the grid use storage compression, storage encryption, or compliance? • Does ILM make replicated or erasure coded objects? Does ILM ensure site redundancy? Do ILM rules use the Strict, Balanced, or Dual Commit ingest behaviors?

✓	Item	Notes
	Log files and system data	<p>Collect log files and system data for your system. Select Support. Then, in the Tools section of the menu, select Logs.</p> <p>You can collect logs for the entire grid, or for selected nodes.</p> <p>If you are collecting logs only for selected nodes, be sure to include at least one Storage Node that has the ADC service. (The first three Storage Nodes at a site include the ADC service.)</p> <p>Collecting log files and system data on page 139</p>
	Baseline information	<p>Collect baseline information regarding ingest operations, retrieval operations, and storage consumption.</p> <p>Establishing baselines on page 153</p>
	Timeline of recent changes	<p>Create a timeline that summarizes any recent changes to the system or its environment.</p> <p>Creating a timeline of recent changes on page 152</p>
	History of efforts to diagnose the issue	<p>If you have taken steps to diagnose or troubleshoot the issue yourself, make sure to record the steps you took and the outcome.</p>

Related concepts

[Collecting data](#) on page 150

After you have defined the problem and have assessed its risk and impact, collect data for analysis. The type of data that is most useful to collect depends upon the nature of the problem.

[Analyzing data](#) on page 154

Use the information that you collect to determine the cause of the problem and potential solutions.

Related reference

[Defining the problem](#) on page 149

The first step to solving a problem is to define the problem clearly.

[Assessing the risk and impact on the system](#) on page 150

After you have defined the problem, assess its risk and impact on the StorageGRID system. For example, the presence of critical alerts does not necessarily mean that the system is not delivering core services.

Related information

[Administering StorageGRID](#)

Troubleshooting object and storage issues

There are several tasks you can perform to help determine the source of object and storage issues.

Confirming object data locations

Depending on the problem, you might want to confirm where object data is being stored. For example, you might want to verify that the ILM policy is performing as expected and object data is being stored where intended.

Before you begin

- You must have an object identifier, which can be one of:

- **UUID:** The object's Universally Unique Identifier. Enter the UUID in all uppercase.
- **CBID:** The object's unique identifier within StorageGRID. You can obtain an object's CBID from the audit log. Enter the CBID in all uppercase.
- **S3 bucket and object key:** When an object is ingested through the S3 interface, the client application uses a bucket and object key combination to store and identify the object.
- **Swift container and object name:** When an object is ingested through the Swift interface, the client application uses a container and object name combination to store and identify the object.

Steps

1. Select **ILM > Object Metadata Lookup**.
2. Type the object's identifier in the **Identifier** field.

You can enter a UUID, CBID, S3 bucket/object-key, or Swift container/object-name.

Object Metadata Lookup

Enter the identifier for any object stored in the grid to view its metadata.

Identifier	<input type="text" value="source/testobject"/>	<input type="button" value="Look Up"/>
------------	--	--

3. Click **Look Up**.

The object metadata lookup results appear. This page lists the following types of information:

- System metadata, including the object ID (UUID), the object name, the name of the container, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
- Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object.
- For replicated object copies, the current storage location of each copy.
- For erasure-coded object copies, the current storage location of each fragment.
- For object copies in a Cloud Storage Pool, the location of the object, including the name of the external bucket and the object's unique identifier.
- For segmented objects and multipart objects, a list of object segments including segment identifiers and data sizes. For objects with more than 100 segments, only the first 100 segments are shown.
- All object metadata in the unprocessed, internal storage format. This raw metadata includes internal system metadata that is not guaranteed to persist from release to release.

The following example shows the object metadata lookup results for an S3 test object that is stored as two replicated copies.

System Metadata

Object ID	A12E96FF-B13F-4905-9E9E-45373F6E7DA8
Name	testobject
Container	source
Account	t-1582139188
Size	5.24 MB
Creation Time	2020-02-19 12:15:59 PST
Modified Time	2020-02-19 12:15:59 PST

Replicated Copies

Node	Disk Path
99-97	/var/local/rangedb/2/p/06/0B/00nM8H\$ TFbnQQ} CV2E
99-99	/var/local/rangedb/1/p/12/0A/00nM8H\$ TFboW28 CXG%

Raw Metadata

```
{
  "TYPE": "CTNT",
  "CHND": "A12E96FF-B13F-4905-9E9E-45373F6E7DA8",
  "NAME": "testobject",
  "CBID": "0x8823DE7EC7C10416",
  "PHND": "FEA0AE51-534A-11EA-9FCD-31FF00C36D56",
  "PPTH": "source",
  "META": {
    "BASE": {
      "PAWS": "2",
```

Related information

[Administering StorageGRID](#)

[Implementing S3 client applications](#)

[Implementing Swift client applications](#)

Object store (storage volume) failures

The underlying storage on a Storage Node is divided into object stores. These object stores are physical partitions that act as mount points for StorageGRID system’s storage. Object stores are also known as storage volumes.

You can view object store information for each Storage Node. Object stores are shown at the bottom of the **Nodes > Storage Node > Storage** page.

Disk Devices				
Name	World Wide Name	I/O Load	Read Rate	Write Rate
croot(8:1,sda1)	N/A	1.62%	0 bytes/s	177 KB/s
cvloc(8:2,sda2)	N/A	17.28%	0 bytes/s	2 MB/s
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	11 KB/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	0 bytes/s
sds(8:48,sdd)	N/A	0.00%	0 bytes/s	0 bytes/s

Volumes					
Mount Point	Device	Status	Size	Available	Write Cache Status
/	croot	Online	21.00 GB	14.25 GB	Unknown
/var/local	cvloc	Online	85.86 GB	84.39 GB	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/2	sds	Online	107.32 GB	107.18 GB	Enabled

Object Stores						
ID	Size	Available	Replicated Data	EC Data	Object Data (%)	Health
0000	107.32 GB	96.45 GB	994.37 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

To see more details about each Storage Node, follow these steps:

1. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
2. Select *site* > *Storage Node* > **LDR** > **Storage** > **Overview** > **Main**.



Overview: LDR (DC1-S1) - Storage

Updated: 2020-01-29 15:03:39 PST

Storage State - Desired:	Online	
Storage State - Current:	Online	
Storage Status:	No Errors	

Utilization

Total Space:	322 GB	
Total Usable Space:	311 GB	
Total Usable Space (Percent):	96.534 %	
Total Data:	994 KB	
Total Data (Percent):	0 %	

Replication

Block Reads:	0	
Block Writes:	0	
Objects Retrieved:	0	
Objects Committed:	0	
Objects Deleted:	0	
Delete Service State:	Enabled	

Object Store Volumes

ID	Total	Available	Replicated Data	EC Data	Stored (%)	Health
0000	107 GB	96.4 GB	994 KB	0 B	0.001 %	No Errors
0001	107 GB	107 GB	0 B	0 B	0 %	No Errors
0002	107 GB	107 GB	0 B	0 B	0 %	No Errors

Depending on the nature of the failure, faults with a storage volume might be reflected in an alarm on the storage status or on the health of an object store. If a storage volume fails, you should repair the failed storage volume to restore the Storage Node to full functionality as soon as possible. If necessary, you can go to the **Configuration** tab and place the Storage Node in a read-only state so that the StorageGRID system can use it for data retrieval while you prepare for a full recovery of the server.

Related information

[Recovery and maintenance](#)

Verifying object integrity

The StorageGRID system verifies the integrity of object data on Storage Nodes, checking for both corrupt and missing objects.

There are two verification processes: background verification and foreground verification. They work together to ensure data integrity. Background verification runs automatically, and continuously checks the correctness of object data. Foreground verification can be triggered by a user, to more quickly verify the existence (although not the correctness) of objects.

Related concepts

[What background verification is](#) on page 160

The background verification process automatically and continuously checks Storage Nodes for corrupt copies of object data, and automatically attempts to repair any issues that it finds.

[What foreground verification is](#) on page 162

Foreground verification is a user-initiated process that checks if all expected object data exists on a Storage Node. Foreground verification is used to verify the integrity of a storage device.

What background verification is

The background verification process automatically and continuously checks Storage Nodes for corrupt copies of object data, and automatically attempts to repair any issues that it finds.

Background verification checks the integrity of replicated objects and erasure-coded objects, as follows:

- **Replicated objects:** If the background verification process finds a replicated object that is corrupt, the corrupt copy is removed from its location and quarantined elsewhere on the Storage Node. Then, a new uncorrupted copy is generated and placed to satisfy the active ILM policy. The new copy might not be placed on the Storage Node that was used for the original copy.

Note: Corrupt object data is quarantined rather than deleted from the system, so that it can still be accessed. For more information on accessing quarantined object data, contact technical support.

- **Erasure-coded objects:** If the background verification process detects that a fragment of an erasure-coded object is corrupt, StorageGRID automatically attempts to rebuild the missing fragment in place on the same Storage Node, using the remaining data and parity fragments. If the corrupted fragment cannot be rebuilt, the Corrupt Copies Detected (ECOR) attribute is incremented by one, and an attempt is made to retrieve another copy of the object. If retrieval is successful, an ILM evaluation is performed to create a replacement copy of the erasure-coded object.

The background verification process checks objects on Storage Nodes only. It does not check objects on Archive Nodes or in a Cloud Storage Pool. Objects must be older than four days to qualify for background verification.

Background verification runs at a continuous rate that is designed not to interfere with ordinary system activities. Background verification cannot be stopped. However you can increase the background verification rate to more quickly verify the contents of a Storage Node if you suspect a problem.

Alerts and alarms (legacy) related to background verification

If the system detects a corrupt object that it cannot correct automatically (because the corruption prevents the object from being identified), the **Unidentified corrupt object detected** alert is triggered.

If background verification cannot replace a corrupted object because it cannot locate another copy, the **Objects lost** alert and the LOST (Lost Objects) legacy alarm are triggered.

Configuring the background verification rate

You can change the rate at which background verification checks replicated object data on a Storage Node if you have concerns about data integrity.

Before you begin

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

About this task

You can change the Verification Rate for background verification on a Storage Node:

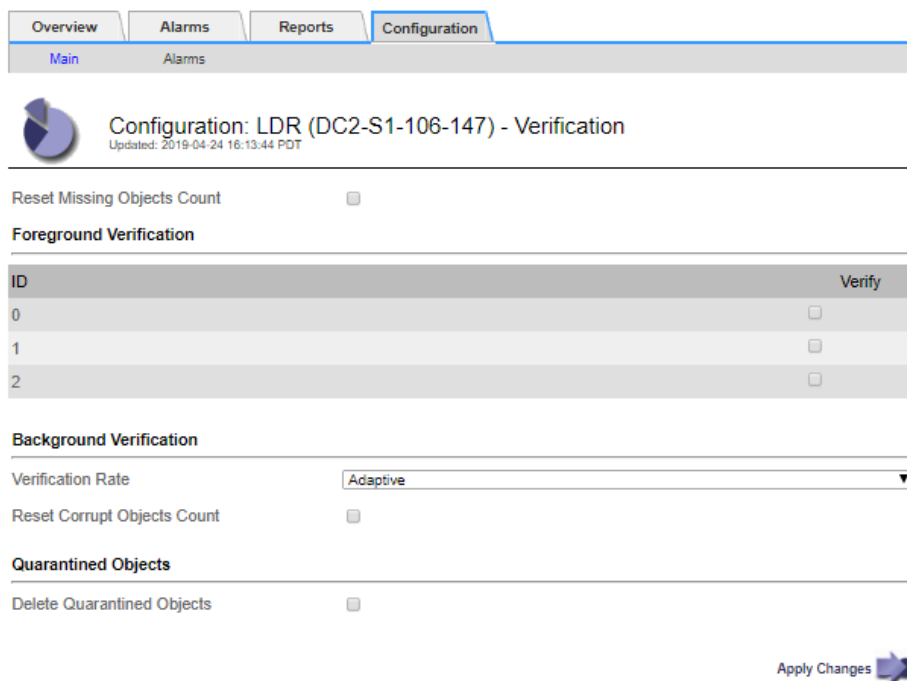
- Adaptive: Default setting. The task is designed to verify at a maximum of 4 MB/s or 10 objects/s (whichever is exceeded first).

- High: Storage verification proceeds quickly, at a rate that can slow ordinary system activities.

Use the High verification rate only when you suspect that a hardware or software fault might have corrupted object data. After the High priority background verification completes, the Verification Rate automatically resets to Adaptive.

Steps

1. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
2. Select **Storage Node > LDR > Verification**.
3. Click **Configuration > Main**.
4. Go to **LDR > Verification > Configuration > Main**.
5. Under Background Verification, select **Verification Rate > High** or **Verification Rate > Adaptive**.



Note: Setting the Verification Rate to High triggers the VPRI (Verification Rate) legacy alarm at the Notice level.

6. Click **Apply Changes**.
7. Monitor the results of background verification for replicated objects.
 - a. Go to **Nodes > Storage Node > Objects**.
 - b. In the Verification section, monitor the values for **Corrupt Objects** and **Corrupt Objects Unidentified**.

If background verification finds corrupt replicated object data, the **Corrupt Objects** metric is incremented, and StorageGRID attempts to extract the object identifier from the data, as follows:

- If the object identifier can be extracted, StorageGRID automatically creates a new copy of the object data. The new copy can be made anywhere in the StorageGRID system that satisfies the active ILM policy.

- If the object identifier cannot be extracted (because it has been corrupted), the **Corrupt Objects Unidentified** metric is incremented, and the **Unidentified corrupt object detected** alert is triggered.
 - c. If corrupt replicated object data is found, contact technical support to determine the root cause of the corruption.
8. Monitor the results of background verification for erasure-coded objects.
- If background verification finds corrupt fragments of erasure-coded object data, the Corrupt Fragments Detected attribute is incremented. StorageGRID recovers by rebuilding the corrupt fragment in place on the same Storage Node.
- a. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
 - b. Select **Storage Node > LDR > Erasure Coding > . > . .**
 - c. In the Verification Results table, monitor the Corrupt Fragments Detected (ECCD) attribute.
9. After corrupt objects have been automatically restored by the StorageGRID system, reset the count of corrupt objects.
- a. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
 - b. Select **Storage Node > LDR > Verification > Configuration**.
 - c. Select **Reset Corrupt Object Count**.
 - d. Click **Apply Changes**.
10. If you are confident that quarantined objects are not required, you can delete them.
- Note:** If the **Objects lost** alert or the LOST (Lost Objects) legacy alarm was triggered, technical support might want to access quarantined objects to help debug the underlying issue or to attempt data recovery.
- a. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
 - b. Select **Storage Node > LDR > Verification > Configuration**.
 - c. Select **Delete Quarantined Objects**.
 - d. Click **Apply Changes**.

What foreground verification is

Foreground verification is a user-initiated process that checks if all expected object data exists on a Storage Node. Foreground verification is used to verify the integrity of a storage device.

Foreground verification is a faster alternative to background verification that checks the existence, but not the integrity, of object data on a Storage Node. If foreground verification finds that many items are missing, there might be an issue with all or part of a storage device associated with the Storage Node.

Foreground verification checks both replicated object data and erasure-coded object data, as follows:

- **Replicated objects:** If a copy of replicated object data is found to be missing, StorageGRID automatically attempts to replace the copy from copies stored elsewhere in the system. The Storage Node runs an existing copy through an ILM evaluation, which will determine that the current ILM policy is no longer being met for this object because the missing copy no longer exists at the expected location. A new copy is generated and placed to satisfy the system's active ILM policy. This new copy might not be placed in the same location that the missing copy was stored.
- **Erasure-coded objects:** If a fragment of an erasure-coded object is found to be missing, StorageGRID automatically attempts to rebuild the missing fragment in place on the same Storage Node using the remaining fragments. If the missing fragment cannot be rebuilt (because too many fragments have been lost), the Corrupt Copies Detected (ECOR) attribute is

incremented by one. ILM then attempts to find another copy of the object, which it can use to generate a new erasure-coded copy.

If foreground verification identifies an issue with erasure coding on a storage volume, the foreground verification task pauses with an error message that identifies the affected volume. You must perform a recovery procedure for any affected storage volumes.

If no other copies of a missing replicated object or a corrupted erasure-coded object can be found in the grid, the **Objects lost** alert and the LOST (Lost Objects) legacy alarm are triggered.

Running foreground verification

Foreground verification enables you to verify the existence of data on a Storage Node. Missing object data might indicate that an issue exists with the underlying storage device.

Before you begin

- You have ensured that the following grid tasks are not running:
 - Grid Expansion: Add Server (GEXP), when adding a Storage Node
 - Storage Node Decommissioning (LDCM) on the same Storage Node

If these grid tasks are running, wait for them to complete or release their lock.

- You have ensured that the storage is online. (Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**. Then, select **Storage Node > LDR > Storage > Overview > Main**. Ensure that **Storage State - Current** is Online.)
- You have ensured that the following recovery procedures are not running on the same Storage Node:
 - Recovery of a failed storage volume
 - Recovery of a Storage Node with a failed system drive

Foreground verification does not provide useful information while recovery procedures are in progress.

About this task

Foreground verification checks for both missing replicated object data and missing erasure-coded object data:

- If foreground verification finds large amounts of missing object data, there is likely an issue with the Storage Node's storage that needs to be investigated and addressed.
- If foreground verification finds a serious storage error associated with erasure-coded data, it will notify you. You must perform storage volume recovery to repair the error.

You can configure foreground verification to check all of a Storage Node's object stores or only specific object stores.

If foreground verification finds missing object data, the StorageGRID system attempts to replace it. If a replacement copy cannot be made, the LOST (Lost Objects) alarm might be triggered.

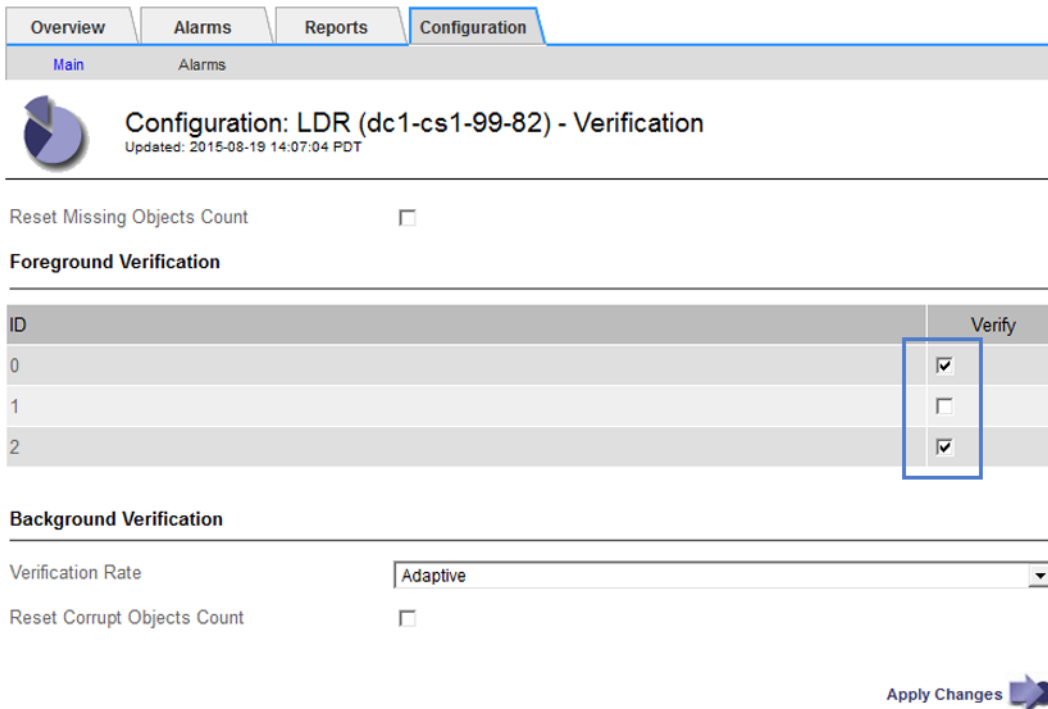
Foreground verification generates an LDR Foreground Verification grid task that, depending on the number of objects stored on a Storage Node, can take days or weeks to complete. It is possible to select multiple Storage Nodes at the same time; however, these grid tasks are not run simultaneously. Instead, they are queued and run one after the other until completion. When foreground verification is in progress on a Storage Node, you cannot start another foreground verification task on that same Storage Node even though the option to verify additional volumes might appear to be available for the Storage Node.

If a Storage Node other than the one where foreground verification is being run goes offline, the grid task continues to run until the **% Complete** attribute reaches 99.99 percent. The **% Complete** attribute then falls back to 50 percent and waits for the Storage Node to return to online status.

When the Storage Node's state returns to online, the LDR Foreground Verification grid task continues until it completes.

Steps

1. Select **Storage Node > LDR > Verification**.
2. Click **Configuration > Main**.
3. Under **Foreground Verification**, select the check box for each storage volume ID you want to verify.



4. Click **Apply Changes**.

Wait until the page auto-refreshes and reloads before you leave the page. Once refreshed, object stores become unavailable for selection on that Storage Node.

An LDR Foreground Verification grid task is generated and runs until it completes, pauses, or is aborted.

5. Monitor missing objects or missing fragments:
 - a. Select **Storage Node > LDR > Verification**.
 - b. On the Overview tab under **Verification Results**, note the value of **Missing Objects Detected**.

Note: The same value is reported as **Lost Objects** on the Nodes page. Go to **Nodes > Storage Node**, and select the **Objects** tab.

If the number of **Missing Objects Detected** is large (if there are a hundreds of missing objects), there is likely an issue with the Storage Node's storage. Contact technical support.

- c. Select **Storage Node > LDR > Erasure Coding**.
- d. On the Overview tab under **Verification Results**, note the value of **Missing Fragments Detected**.

If the number of **Missing Fragments Detected** is large (if there are a hundreds of missing fragments), there is likely an issue with the Storage Node's storage. Contact technical support.

If foreground verification does not detect a significant number of missing replicated object copies or a significant number of missing fragments, then the storage is operating normally.

6. Monitor the completion of the foreground verification grid task:
 - a. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**. Then select *site* > *Admin Node* > **CMN** > **Grid Task** > **Overview** > **Main**.
 - b. Verify that the foreground verification grid task is progressing without errors.

Note: A notice-level alarm is triggered on grid task status (SCAS) if the foreground verification grid task pauses.
 - c. If the grid task pauses with a `critical storage error`, recover the affected volume and then run foreground verification on the remaining volumes to check for additional errors.



Attention: If the foreground verification grid task pauses with the message `Encountered a critical storage error in volume volID`, you must perform the procedure for recovering a failed storage volume. See the recovery and maintenance instructions.

After you finish

If you still have concerns about data integrity, go to **LDR** > **Verification** > **Configuration** > **Main** and increase the background Verification Rate. Background verification checks the correctness of all stored object data and repairs any issues that it finds. Finding and repairing potential issues as quickly as possible reduces the risk of data loss.

Related information

[Recovery and maintenance](#)

Troubleshooting lost and missing object data

Objects can be retrieved for several reasons, including read requests from a client application, background verifications of replicated object data, ILM re-evaluations, and the restoration of object data during the recovery of a Storage Node.

The StorageGRID system uses location information in an object's metadata to determine from which location to retrieve the object. If a copy of the object is not found in the expected location, the system attempts to retrieve another copy of the object from elsewhere in the system, assuming that the ILM policy contains a rule to make two or more copies of the object.

If this retrieval is successful, the StorageGRID system replaces the missing copy of the object. Otherwise, the **Objects lost** alert and the legacy LOST (Lost Objects) alarm are triggered, as follows:

- For replicated copies, if another copy cannot be retrieved, the object is considered lost, and the alert and alarm are triggered.
- For erasure coded copies, if a copy cannot be retrieved from the expected location, the Corrupt Copies Detected (ECOR) attribute is incremented by one before an attempt is made to retrieve a copy from another location. If no other copy is found, the alert and alarm are triggered.

You should investigate all **Objects lost** alerts immediately to determine the root cause of the loss and to determine if the object might still exist in an offline, or otherwise currently unavailable, Storage Node or Archive Node.

In the case where object data without copies is lost, there is no recovery solution. However, you must reset the Lost Object counter to prevent known lost objects from masking any new lost objects.

Related tasks

[Investigating lost objects](#) on page 166

When the **Objects lost** alert and the legacy LOST (Lost Objects) alarm are triggered, you must investigate immediately. Collect information about the affected objects and contact technical support.

[Resetting lost and missing object counts](#) on page 173

After investigating the StorageGRID system and verifying that all recorded lost objects are permanently lost or that it is a false alarm, you can reset the value of the Lost Objects attribute to zero.

Investigating lost objects

When the **Objects lost** alert and the legacy LOST (Lost Objects) alarm are triggered, you must investigate immediately. Collect information about the affected objects and contact technical support.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have the `Passwords.txt` file.

About this task

The **Objects lost** alert and the LOST alarm indicate that StorageGRID believes that there are no copies of an object in the grid. Data might have been permanently lost.

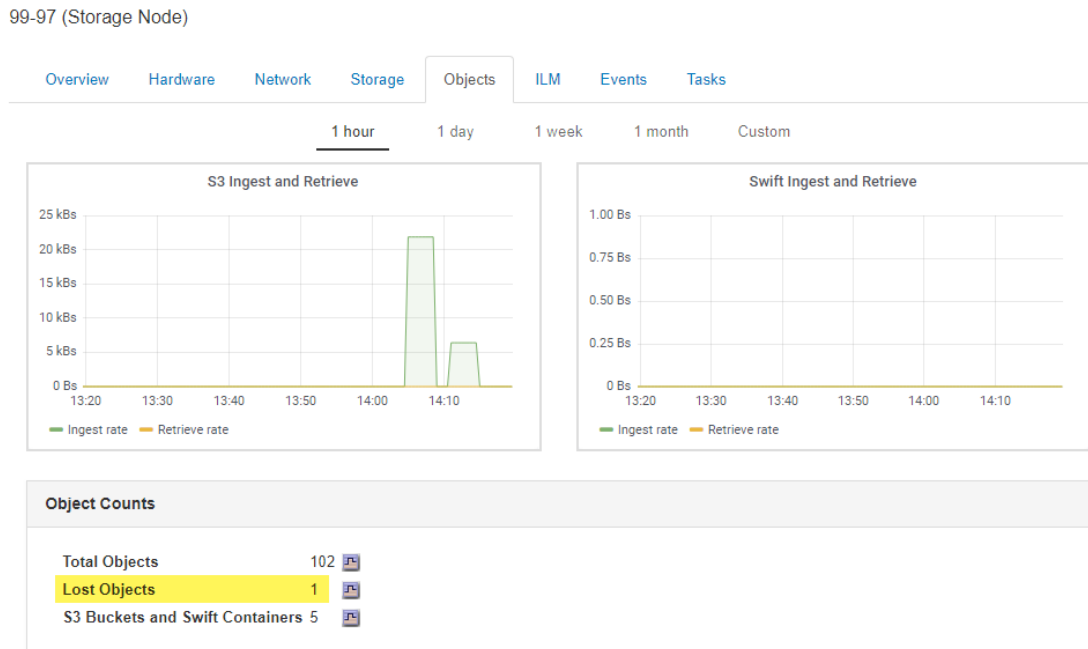
Investigate lost object alarms or alerts immediately. You might need to take action to prevent further data loss. In some cases, you might be able to restore a lost object if you take prompt action.

The number of Lost Objects can be seen in the Grid Manager.

Steps

1. Select **Nodes**.
2. Select **Storage Node > Objects**.
3. Review the number of Lost Objects shown in the Object Counts table.

This number indicates the total number of objects this grid node detects as missing from the entire StorageGRID system. The value is the sum of the Lost Objects counters of the Data Store component within the LDR and DDS services.



4. From an Admin Node, access the audit log to determine the unique identifier (UUID) of the object that triggered the **Objects lost** alert and the LOST alarm:

a. Log in to the grid node:

- i. Enter the following command: `ssh admin@grid_node_IP`
- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

b. Change to the directory where the audit logs are located. Enter:

```
cd /var/local/audit/export/
```

c. Use `grep` to extract the Object Lost (OLST) audit messages. Enter:

```
grep OLST audit_file_name
```

d. Note the UUID value included in the message.

```
>Admin: # grep OLST audit.log
2020-02-12T19:18:54.780426 [AUDT:[CBID(UI64):0x38186FE53E3C49A5]
[UUID(CSTR):"926026C4-00A4-449B-AC72-BCCA72DD1311"]
[PATH(CSTR):"source/cats"][NOID(UI32):12288733][VOLI(UI64):3222345986][RSLT(FC32):NONE]
[AVER(UI32):10]
[ATIM(UI64):1581535134780426][ATYP(FC32):OLST][ANID(UI32):12448208][AMID(FC32):ILMX]
[ATID(UI64):7729403978647354233]]
```

5. Use the `ObjectByUUID` command to find the object by its identifier (UUID), and then determine if data is at risk.

a. Telnet to localhost 1402 to access the LDR console.

b. Enter: `/proc/OBRP/ObjectByUUID UUID_value`

In this first example, the object with UUID `926026C4-00A4-449B-AC72-BCCA72DD1311` has two locations listed.

```
ade 12448208: /proc/OBRP > ObjectByUUID 926026C4-00A4-449B-AC72-BCCA72DD1311
{
  "TYPE(Object Type)": "Data object",
```

```

"CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
"NAME": "cats",
"CBID": "0x38186FE53E3C49A5",
"PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-ACBB00BB82DD",
"PPTH(Parent path)": "source",
"META": {
  "BASE(Protocol metadata)": {
    "PAWS(S3 protocol version)": "2",
    "ACCT(S3 account ID)": "44084621669730638018",
    "*ctp(HTTP content MIME type)": "binary/octet-stream"
  },
  "BYCB(System metadata)": {
    "CSIZ(Plaintext object size)": "5242880",
    "SHSH(Supplementary Plaintext hash)": "MD5D 0xBAC2A2617C1DFF7E959A76731E6EAF5E",
    "BSIZ(Content block size)": "5252084",
    "CVER(Content block version)": "196612",
    "CTME(Object store begin timestamp)": "2020-02-12T19:16:10.983000",
    "MTME(Object store modified timestamp)": "2020-02-12T19:16:10.983000",
    "ITME": "1581534970983000"
  },
  "CMSM": {
    "LATM(Object last access time)": "2020-02-12T19:16:10.983000"
  },
  "AWS3": {
    "LOCC": "us-east-1"
  }
},
"CLCO(Locations)": [
  {
    "Location Type": "CLDI(Location online)",
    "NOID(Node ID)": "12448208",
    "VOLI(Volume ID)": "3222345473",
    "Object File Path": "/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila#3udu",
    "LTIM(Location timestamp)": "2020-02-12T19:36:17.880569"
  },
  {
    "Location Type": "CLDI(Location online)",
    "NOID(Node ID)": "12288733",
    "VOLI(Volume ID)": "3222345984",
    "Object File Path": "/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb#3s;L",
    "LTIM(Location timestamp)": "2020-02-12T19:36:17.934425"
  }
]
}

```

In the second example, the object with UUID 926026C4-00A4-449B-AC72-BCCA72DD1311 has no locations listed.

```

ade 12448208: / > /proc/OBRP/ObjectByUUID 926026C4-00A4-449B-AC72-BCCA72DD1311
{
  "TYPE(Object Type)": "Data object",
  "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
  "NAME": "cats",
  "CBID": "0x38186FE53E3C49A5",
  "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-ACBB00BB82DD",
  "PPTH(Parent path)": "source",
  "META": {
    "BASE(Protocol metadata)": {
      "PAWS(S3 protocol version)": "2",
      "ACCT(S3 account ID)": "44084621669730638018",
      "*ctp(HTTP content MIME type)": "binary/octet-stream"
    },
    "BYCB(System metadata)": {
      "CSIZ(Plaintext object size)": "5242880",
      "SHSH(Supplementary Plaintext hash)": "MD5D 0xBAC2A2617C1DFF7E959A76731E6EAF5E",
      "BSIZ(Content block size)": "5252084",
      "CVER(Content block version)": "196612",
      "CTME(Object store begin timestamp)": "2020-02-12T19:16:10.983000",
      "MTME(Object store modified timestamp)": "2020-02-12T19:16:10.983000",
      "ITME": "1581534970983000"
    },
    "CMSM": {
      "LATM(Object last access time)": "2020-02-12T19:16:10.983000"
    },
    "AWS3": {

```



```

        "LOCC": "us-east-1"
    }
}
}
    
```

c. Review the output of `/proc/OBRP/ObjectByUUID`, and take the appropriate action:

Metadata	Conclusion
No object found ("ERROR":")	If the object is not found, the message "ERROR:" is returned. If the object is not found, it is safe to ignore the alarm. The lack of an object indicates that the object was intentionally deleted.
Locations > 0	If there are locations listed in the output, the Lost Objects alarm might be a false positive. Confirm that the objects exist. Use the Node ID and filepath listed in the output to confirm that the object file is in the listed location. (The procedure for finding potentially lost objects explains how to use the Node ID to find the correct Storage Node.) Searching for and restoring potentially lost objects on page 169 If the objects exist, you can reset the count of Lost Objects to clear the alarm and the alert.
Locations = 0	If there are no locations listed in the output, the object is potentially missing. You can try to find and restore the object yourself, or you can contact technical support. Searching for and restoring potentially lost objects on page 169 Technical support might ask you to determine if there is a storage recovery procedure in progress. That is, has a <code>repair-data</code> command been issued on any Storage Node and is the recovery still in progress? See "Restoring object data to a storage volume" in the recovery and maintenance instructions.

Related information

- [Recovery and maintenance](#)
- [Understanding audit messages](#)

Searching for and restoring potentially lost objects

It might be possible to find and restore objects that have triggered a Lost Objects (LOST) alarm and a **Object lost** alert and that you have identified as potentially lost.

Before you begin

- You must have the UUID of any lost object, as identified in "Investigating lost objects."
- You must have the `Passwords.txt` file.

About this task

You can follow this procedure to look for replicated copies of the lost object elsewhere in the grid. In most cases, the lost object will not be found. However, in some cases, you might be able to find and restore a lost replicated object if you take prompt action.



Attention: Contact technical support for assistance with this procedure.

Steps

- From an Admin Node, search the audit logs for possible object locations:
 - Log in to the grid node:
 - Enter the following command: `ssh admin@grid_node_IP`

- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

- b. Change to the directory where the audit logs are located:

```
cd /var/local/audit/export/
```

- c. Use `grep` to extract the audit messages associated with the potentially lost object and send them to an output file. Enter:

```
grep uuid-value audit_file_name > output_file_name
```

For example:

```
Admin: # grep 926026C4-00A4-449B-AC72-BCCA72DD1311 audit.log >
messages_about_lost_object.txt
```

- d. Use `grep` to extract the Location Lost (LLST) audit messages from this output file. Enter:

```
grep LLST output_file_name
```

For example:

```
Admin: # grep LLST messages_about_lost_objects.txt
```

An LLST audit message looks like this sample message.

```
[AUDT:[NOID(UI32):12448208][CBIL(UI64):0x38186FE53E3C49A5]
[UUID(CSTR):"926026C4-00A4-449B-AC72-BCCA72DD1311"]][LTYP(FC32):CLDI]
[PCLD(CSTR):"/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6"]
[TSRC(FC32):SYST][RSLT(FC32):NONE][AVER(UI32):10][ATIM(UI64):
1581535134379225][ATYP(FC32):LLST][ANID(UI32):12448208][AMID(FC32):CLSM]
[ATID(UI64):7086871083190743409]]
```

- e. Find the PCLD field and the NOID field in the LLST message.

If present, the value of PCLD is the complete path on disk to the missing replicated object copy. The value of NOID is the node id of the LDR where a copy of the object might be found.



Attention: If you do not find LLST messages in your output file, or if the LLST message does not contain a PCLD field, audit messages for the object might have been recorded before the grid was upgraded to StorageGRID 11.4. Contact technical support for assistance with searching for an object location using the object's CBID.

If you find an object location, you might be able to restore the object.

- f. Find the Storage Node for this LDR node ID.

There are two ways to use the node ID to find the Storage Node:

- In the Grid Manager, select **Support**. Then, in the Tools section of the menu, select **Grid Topology**. Then select **Data Center > Storage Node > LDR**. The LDR node ID is in the Node Information table. Review the information for each Storage Node until you find the one that hosts this LDR.
- Download and unzip the Recovery Package for the grid. There is a `\docs` directory in the SAID package. If you open the `index.html` file, the Servers Summary shows all node IDs for all grid nodes.

- 2. Determine if the object exists on the Storage Node indicated in the audit message:

- a. Log in to the grid node:

- i. Enter the following command: `ssh admin@grid_node_IP`
- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

- b. Determine if the file path for the object exists.

For the file path of the object, use the value of `PCLD` from the LLST audit message.

For example, enter:

```
ls '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6'
```

Note: Always enclose the object file path in single quotes in commands to escape any special characters.

- If the object path is not found, the object is lost and cannot be restored using this procedure. Contact technical support.
- If the object path is found, continue with step 3 on page 171. You can attempt to restore the found object back to StorageGRID.

3. If the object path was found, attempt to restore the object to StorageGRID:

- a. From the same Storage Node, change the ownership of the object file so that it can be managed by StorageGRID. Enter:

```
chown ldr-user:bycast 'file_path_of_object'
```

- b. Telnet to localhost 1402 to access the LDR console. Enter:

```
telnet 0 1402
```

- c. Enter:

```
cd /proc/STOR
```

- d. Enter:

```
Object_Found 'file_path_of_object'
```

For example, enter:

```
Object_Found '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6'
```

Issuing the `Object_Found` command notifies the grid of the object's location. It also triggers the active ILM policy, which makes additional copies as specified in the policy.

Note: If the Storage Node where you found the object is offline, you can copy the object to any Storage Node that is online. Place the object in any `/var/local/rangedb` directory of the online Storage Node. Then, issue the `Object_Found` command using that file path to the object.

- If the object cannot be restored, the `Object_Found` command fails. Contact technical support.
- If the object was successfully restored to StorageGRID, a success message appears. For example:

```
ade 12448208: /proc/STOR > Object_Found '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6'  
ade 12448208: /proc/STOR > Object found succeeded.  
First packet of file was valid. Extracted key: 38186FE53E3C49A5  
Renamed '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6' to '/var/local/rangedb/1/p/  
17/11/00rH0%DkRt78Ila#3udu'
```

Continue with step 4 on page 171.

4. If the object was successfully restored to StorageGRID, verify that new locations were created.

- a. Enter:

```
cd /proc/OBRP
```

- b. Enter: **ObjectByUUID *UUID_value***

The following example shows that there are two locations for the object with UUID 926026C4-00A4-449B-AC72-BCCA72DD1311.

```
ade 12448208: /proc/OBRP > ObjectByUUID 926026C4-00A4-449B-AC72-BCCA72DD1311
{
  "TYPE(Object Type)": "Data object",
  "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
  "NAME": "cats",
  "CBID": "0x38186FE53E3C49A5",
  "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-ACBB00BB82DD",
  "PPTH(Parent path)": "source",
  "META": {
    "BASE(Protocol metadata)": {
      "PAWS(S3 protocol version)": "2",
      "ACCT(S3 account ID)": "44084621669730638018",
      "*ctp(HTTP content MIME type)": "binary/octet-stream"
    },
    "BYCB(System metadata)": {
      "CSIZ(Plaintext object size)": "5242880",
      "SHSH(Supplementary Plaintext hash)": "MD5D 0xBAC2A2617C1DF7E959A76731E6EAF5E",
      "BSIZ(Content block size)": "5252084",
      "CVER(Content block version)": "196612",
      "CTME(Object store begin timestamp)": "2020-02-12T19:16:10.983000",
      "MTME(Object store modified timestamp)": "2020-02-12T19:16:10.983000",
      "ITME": "1581534970983000"
    },
    "CMSM": {
      "LATM(Object last access time)": "2020-02-12T19:16:10.983000"
    },
    "AWS3": {
      "LOCC": "us-east-1"
    }
  },
  "CLCO(Locations)": [
    {
      "Location Type": "CLDI(Location online)",
      "NOID(Node ID)": "12448208",
      "VOLI(Volume ID)": "3222345473",
      "Object File Path": "/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila#3udu",
      "LTIM(Location timestamp)": "2020-02-12T19:36:17.880569"
    },
    {
      "Location Type": "CLDI(Location online)",
      "NOID(Node ID)": "12288733",
      "VOLI(Volume ID)": "3222345984",
      "Object File Path": "/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb#3s;L",
      "LTIM(Location timestamp)": "2020-02-12T19:36:17.934425"
    }
  ]
}
```

- c. Sign out of the LDR console. Enter:

```
exit
```

5. From an Admin Node, search the audit logs for the ORLM audit message for this object to confirm that information lifecycle management (ILM) has placed copies as required.

- a. Log in to the grid node:

- i. Enter the following command: `ssh admin@grid_node_IP`
- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

- b. Change to the directory where the audit logs are located:

```
cd /var/local/audit/export/
```

- c. Use grep to extract the audit messages associated with the object to an output file. Enter:

```
grep uuid-value audit_file_name > output_file_name
```

For example:

```
Admin: # grep 926026C4-00A4-449B-AC72-BCCA72DD1311 audit.log >
messages_about_restored_object.txt
```

- d. Use grep to extract the Object Rules Met (ORLM) audit messages from this output file.

Enter:

```
grep ORLM output_file_name
```

For example:

```
Admin: # grep ORLM messages_about_restored_object.txt
```

An ORLM audit message looks like this sample message.

```
[AUDT:[CBID(UI64):0x38186FE53E3C49A5][RULE(CSTR):"Make 2 Copies"]
[STAT(FC32):DONE][CSIZ(UI64):0][UUID(CSTR):"926026C4-00A4-449B-AC72-BCCA72DD1311"]
[LOCS(CSTR):"CLDI 12828634 2148730112, CLDI 12745543 2147552014"]
[RSLT(FC32):SUCS][AVER(UI32):10][ATYP(FC32):ORLM][ATIM(UI64):1563398230669]
[ATID(UI64):15494889725796157557][ANID(UI32):13100453][AMID(FC32):BCMS]]
```

- e. Find the LOCS field in the audit message.

If present, the value of CLDI in LOCS is the node ID and the volume ID where an object copy has been created. This message shows that the ILM has been applied and that two object copies have been created in two locations in the grid.

6. Reset the count of lost objects in the Grid Manager.

Related tasks

[Investigating lost objects](#) on page 166

When the **Objects lost** alert and the legacy LOST (Lost Objects) alarm are triggered, you must investigate immediately. Collect information about the affected objects and contact technical support.

[Confirming object data locations](#) on page 155

Depending on the problem, you might want to confirm where object data is being stored. For example, you might want to verify that the ILM policy is performing as expected and object data is being stored where intended.

[Resetting lost and missing object counts](#) on page 173

After investigating the StorageGRID system and verifying that all recorded lost objects are permanently lost or that it is a false alarm, you can reset the value of the Lost Objects attribute to zero.

Related information

[Understanding audit messages](#)

Resetting lost and missing object counts

After investigating the StorageGRID system and verifying that all recorded lost objects are permanently lost or that it is a false alarm, you can reset the value of the Lost Objects attribute to zero.

Before you begin

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

About this task

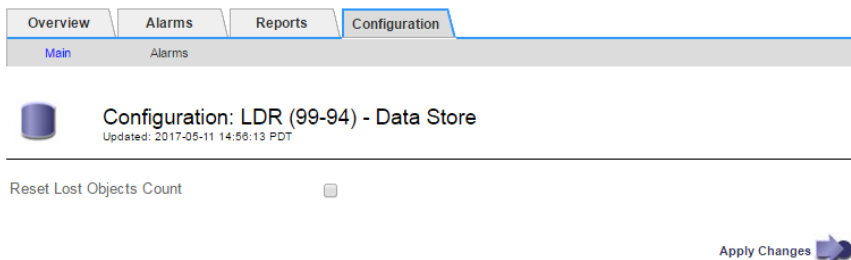
You can reset the Lost Objects counter from either of the following pages:

- **Support > Grid Topology > site > Storage Node > LDR > Data Store > Overview > Main**
- **Support > Grid Topology > site > Storage Node > DDS > Data Store > Overview > Main**

These instructions show resetting the counter from the **LDR > Data Store** page.

Steps

1. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
2. Select **Site > Storage Node > LDR > Data Store > Configuration** for the Storage Node that has the **Objects lost** alert or the LOST alarm.
3. Select **Reset Lost Objects Count**.



4. Click **Apply Changes**.
The Lost Objects attribute is reset to 0 and the **Objects lost** alert and the LOST alarm clear, which can take a few minutes.
5. Optionally, reset other related attribute values that might have been incremented in the process of identifying the lost object.
 - a. Select **Site > Storage Node > LDR > Erasure Coding > Configuration**.
 - b. Select **Reset Reads Failure Count** and **Reset Corrupt Copies Detected Count**.
 - c. Click **Apply Changes**.
 - d. Select **Site > Storage Node > LDR > Verification > Configuration**.
 - e. Select **Reset Missing Objects Count** and **Reset Corrupt Objects Count**.
 - f. If you are confident that quarantined objects are not required, you can select **Delete Quarantined Objects**.

Quarantined objects are created when background verification identifies a corrupt replicated object copy. In most cases StorageGRID automatically replaces the corrupt object, and it is safe to delete the quarantined objects. However, if the **Objects lost** alert or the LOST alarm is triggered, technical support might want to access the quarantined objects.

- g. Click **Apply Changes**.

It can take a few moments for the attributes to reset after you click **Apply Changes**.

Related information

[Administering StorageGRID](#)

Troubleshooting the Low object data storage alert

The **Low object data storage** alert monitors how much space is available for storing object data on each Storage Node. It is related to the Storage Status (SSTS) legacy alarm, but it is not exactly equivalent.

Before you begin

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

- You must have specific access permissions.

About this task

The **Low object data storage** is triggered when the total amount of replicated and erasure coded object data on a Storage Node meets one of the conditions configured in the alert rule.

By default, a major alert is triggered when this condition evaluates as true:

```
(storagegrid_storage_utilization_data_bytes/  
(storagegrid_storage_utilization_data_bytes +  
storagegrid_storage_utilization_usable_space_bytes)) >=0.90
```

In this condition:

- `storagegrid_storage_utilization_data_bytes` is an estimate of the total size of replicated and erasure coded object data for a Storage Node.
- `storagegrid_storage_utilization_usable_space_bytes` is the total amount of object storage space remaining for a Storage Node.

If a major or minor **Low object data storage** alert is triggered, you should perform an expansion procedure as soon as possible.

Steps

1. Select **Alerts > Current**.
The Alerts page appears.
2. From the table of alerts, expand the **Low object data storage** alert group, if required, and select the alert you want to view.


Note: Select the alert, not the heading for a group of alerts.

3. Review the details in the dialog box, and note the following:

- Time triggered
- The name of the site and node
- The current values of the metrics for this alert

4. Select **Nodes > Storage_Node or Site > Storage**.
5. Hover your cursor over the **Storage Used - Object Data** graph.

A pop-up displays Used (%), Used, and Total capacities. The Used value is the `storagegrid_storage_utilization_data_bytes` metric.



6. Select the time controls above the graph to view storage use over different time periods.
Looking at storage use over time can help you understand how much storage was used before and after the alert was triggered and can help you estimate how long it might take for the node's remaining space to become full.

7. As soon as possible, perform an expansion procedure to add storage capacity.

You can add storage volumes (LUNs) to existing Storage Nodes, or you can add new Storage Nodes.

Note: To manage a full Storage Node, see the instructions for administering StorageGRID.

Related tasks

[Troubleshooting the Storage Status \(SSTS\) alarm](#) on page 177

The Storage Status (SSTS) alarm is triggered if a Storage Node has insufficient free space remaining for object storage.

Related information

[Expanding a StorageGRID system](#)

[Administering StorageGRID](#)

Troubleshooting the Total Usable Space (Percent) (SAVP) alarm


If the Total Usable Space (SAVP) alarm appears, you can investigate the cause.

Before you begin

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


Steps

1. Select **Support**. Then, use the options in the Alarms (legacy) section of the menu.
2. Notice the SAVP alarm.
3. Click the attribute name to display more information.



4. Determine when the Storage Node is likely to reach capacity:
 - a. Select **Nodes > Storage Node > Storage**.
 - b. Choose a range (1 hour, 1 day, 1 week, 1 month, 1 year, or Custom).
 - c. Hover your cursor over the **Storage Used - Object Data** graph. Then, slide your cursor to the right to determine how quickly storage is being used over the chosen range of time.

Note: The graph shows how much storage is being used for object data. You can use the Storage Used - Object Metadata graph to determine how much storage is being used for object metadata.



- If an inadequate amount of space remains on this Storage Node and other Storage Nodes, add storage to the system by adding Storage Nodes in an expansion procedure.

Related reference

[Alarms reference \(legacy system\)](#) on page 222

Related information

[Expanding a StorageGRID system](#)

Troubleshooting the Storage Status (SSTS) alarm

The Storage Status (SSTS) alarm is triggered if a Storage Node has insufficient free space remaining for object storage.

Before you begin

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

About this task

The SSTS (Storage Status) alarm is triggered at the Notice level when the amount of free space on every volume in a Storage Node falls below the value of the Storage Volume Soft Read Only Watermark (**Configuration > Storage Options > Overview**).



Storage Options Overview
 Updated: 2019-10-09 13:09:30 MDT

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark	30 GB
Storage Volume Soft Read-Only Watermark	10 GB
Storage Volume Hard Read-Only Watermark	5 GB
Metadata Reserved Space	3,000 GB

For example, suppose the Storage Volume Soft Read-Only Watermark is set to 10 GB, which is its default value. The SSTS alarm is triggered if less than 10 GB of usable space remains on each

storage volume in the Storage Node. If any of the volumes has 10 GB or more of available space, the alarm is not triggered.

If an SSTS alarm has been triggered, you can follow these steps to better understand the issue.

Steps

1. Select **Support**. Then, in the Alarms (legacy) section of the menu, select **Current Alarms**.
2. From the Service column, select the data center, node, and service that are associated with the SSTS alarm.

The Grid Topology page appears. The Alarms tab shows the active alarms for the node and service you selected.

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Notice	SSTS (Storage Status)	Insufficient Free Space	2019-10-09 12:42:51 MDT	Insufficient Free Space	Insufficient Free Space		<input type="checkbox"/>
Notice	SAVP (Total Usable Space (Percent))	Under 10 %	2019-10-09 12:43:21 MDT	7.95 %	7.95 %		<input type="checkbox"/>
Normal	SHLH (Health)						<input type="checkbox"/>

In this example, both the SSTS (Storage Status) and SAVP (Total Usable Space (Percent)) alarms have been triggered at the Notice level.

Note: Typically, both the SSTS alarm and the SAVP alarm are triggered at about the same time; however, whether both alarms are triggered depends on the the watermark setting in GB and the SAVP alarm setting in percent.

3. To determine how much usable space is actually available, select **LDR > Storage > Overview**, and find the Total Usable Space (STAS) attribute.

Overview: LDR (DC1-S1-101-193) - Storage
 Updated: 2019-10-09 12:51:07 MDT

Storage State - Desired: Online
 Storage State - Current: Read-only
 Storage Status: Insufficient Free Space

Utilization

Total Space:	164 GB
Total Usable Space:	19.6 GB
Total Usable Space (Percent):	11.937 %
Total Data:	139 GB
Total Data (Percent):	84.567 %

Replication

Block Reads:	0
Block Writes:	2,279,881
Objects Retrieved:	0
Objects Committed:	88,882
Objects Deleted:	16
Delete Service State:	Enabled

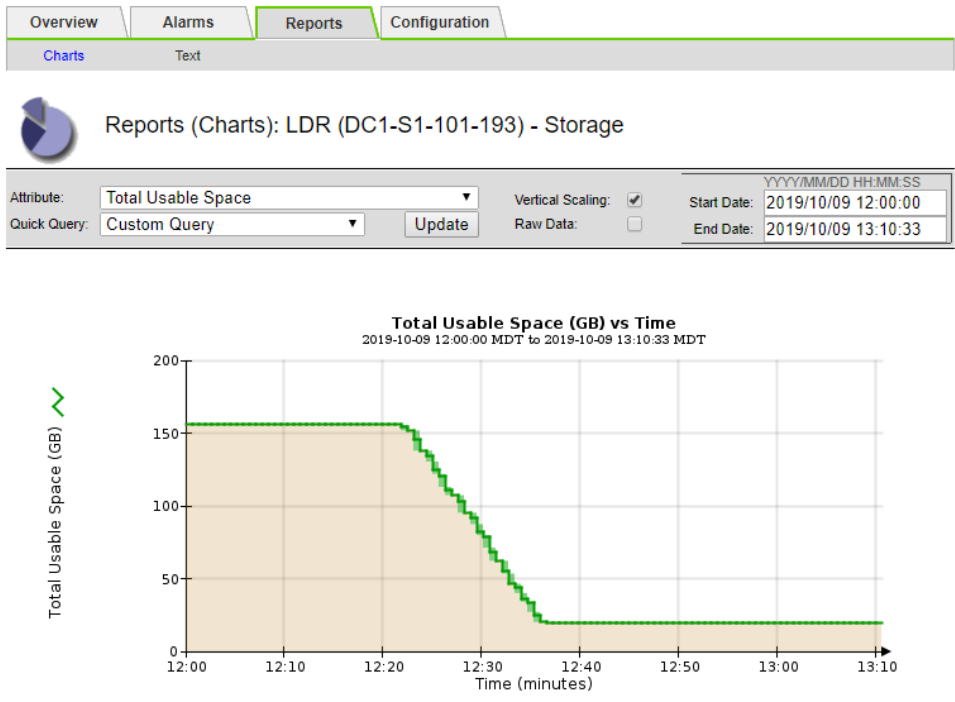
Object Store Volumes

ID	Total	Available	Replicated Data	EC Data	Stored (%)	Health
0000	54.7 GB	2.93 GB	46.2 GB	0 B	84.486 %	No Errors
0001	54.7 GB	8.32 GB	46.3 GB	0 B	84.644 %	No Errors
0002	54.7 GB	8.36 GB	46.3 GB	0 B	84.57 %	No Errors

In this example, only 19.6 GB of the 164 GB of space on this Storage Node remains available. Note that the total value is the sum of the **Available** values for the three object store volumes. The SSTS alarm was triggered because each of the three storage volumes had less than 10 GB of available space.

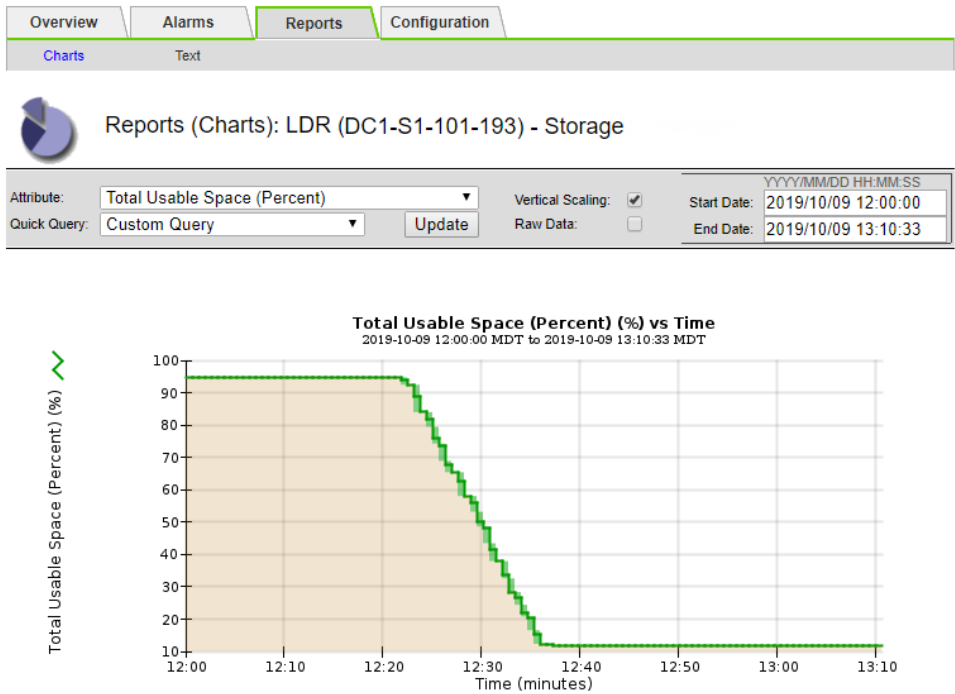
- To understand how storage has been used over time, select the **Reports** tab, and plot Total Usable Space over the last few hours.

In this example, Total Usable Space dropped from roughly 155 GB at 12:00 to 20 GB at 12:35, which corresponds to the time at which the SSTS alarm was triggered.



- To understand how storage is being used as a percent of the total, plot Total Usable Space (Percent) over the last few hours.

In this example, the total usable space dropped from 95% to just over 10% at approximately the same time.



- As required, add storage capacity by expanding the StorageGRID system.

For procedures on how to manage a full Storage Node, see the instructions for administering StorageGRID.

Related information

[Expanding a StorageGRID system](#)

[Administering StorageGRID](#)

Troubleshooting delivery of platform services messages (SMTT alarm)

The Total Events (SMTT) alarm is triggered in the Grid Manager if a platform service message is delivered to an destination that cannot accept the data.

About this task

For example, an S3 multipart upload can succeed even though the associated replication or notification message cannot be delivered to the configured endpoint. Or, a message for CloudMirror replication can fail to be delivered if the metadata is too long.

The SMTT alarm contains a Last Event message that says, `Failed to publish notifications for bucket-name object key` for the last object whose notification failed.

For additional information about troubleshooting platform services, see the instructions for administering StorageGRID. You might need to access the tenant from the Tenant Manager to debug a platform service error.

Steps

1. To view the alarm, select **Nodes** > *site* > *grid node* > **Events**.
2. View Last Event at the top of the table.
Event messages are also listed in `/var/local/log/bycast-err.log`.
3. Follow the guidance provided in the SMTT alarm contents to correct the issue.
4. Click **Reset event counts**.
5. Notify the tenant of the objects whose platform services messages have not been delivered.
6. Instruct the tenant to trigger the failed replication or notification by updating the object's metadata or tags.

Related concepts

[Log files reference](#) on page 246

The following sections list the logs 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.

Related tasks

[Resetting event counts](#) on page 134

After resolving system events, you can reset event counts to zero.

Related information

[Administering StorageGRID](#)

[Using tenant accounts](#)

Troubleshooting the S3 multipart part too small alert

The **S3 multipart part too small** alert is triggered if an S3 client attempts to complete a multipart upload with parts that do not meet Amazon S3 size limits.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have the `Passwords.txt` file.

About this task

StorageGRID enforces Amazon S3 size limits for multipart parts. S3 clients must follow these guidelines:

- Each part in a multipart upload must be between 5 MiB (5,242,880 bytes) and 5 GiB (5,368,709,120 bytes).
- The last part can be smaller than 5 MiB (5,242,880 bytes).
- In general, part sizes should be as large as possible. For example, use part sizes of 5 GiB for a 100 GiB object. Since each part is considered a unique object, using large part sizes reduces StorageGRID metadata overhead.
- For objects smaller than 5 GiB, consider using non-multipart upload instead.

If the **S3 multipart part too small** alert is triggered, you must inform S3 client users to modify their request settings. To give clients time to adjust their multipart upload settings, you can run a script to temporarily disable the enforcement of minimum part size.

Steps

1. Use the tenant ID shown in the alert details to identify the tenant account.
2. Use the client IP address shown in the alert details to identify the specific S3 client.
3. Inform all client users that each part in a multipart upload must be between 5 MiB and 5 GiB. Only the last part can be smaller than 5 MiB (5,242,880 bytes).
4. Optionally, temporarily disable the minimum size restriction for StorageGRID:
 - a. Log in to the primary Admin Node:

```
ssh admin@grid_node_IP
```
 - b. Enter the password from `Passwords.txt`.
 - c. Switch to root:

```
su-
```
 - d. Run this script:

```
'/usr/local/mgmt-api/app/scripts/s3-multipart-part-size-enforcement.rb --disable'
```

(Use `--enable` to undo this change.)

Note: This script will be removed in StorageGRID 11.5. After you upgrade to StorageGRID 11.5, all S3 clients must use part sizes between 5 MiB and 5 GiB.
5. To determine if an S3 client has successfully updated their multipart upload requests.
 - a. Ask the client to upload a test object.
 - b. Go to **ILM > Object Metadata Lookup**, and look up the object.
 - c. Confirm that the size of each part (*segment*) is greater than 5 MiB (5,242,880 bytes).
6. If you disabled the size restriction, look for S3 multipart upload requests in which the parts are too small (requests that previously would have failed):
 - a. Access the `/var/local/log/bycast-err.log` file.
 - b. Grep the log file.
 - Grep for the tenant ID to find errors related to a specific S3 tenant account.
 - Grep for `EntityTooSmall` to find requests that would have failed if minimum size checking were enabled.

Troubleshooting metadata issues

There are several tasks you can perform to help determine the source of metadata problems.

Troubleshooting the Low metadata storage alert

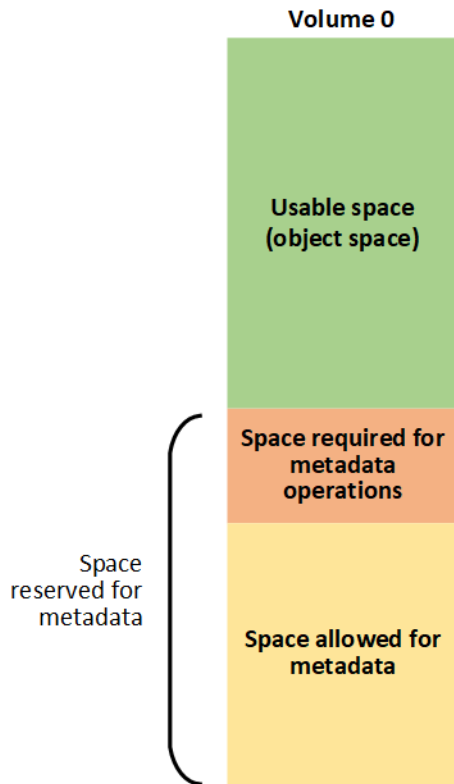
If the **Low metadata storage** alert or the legacy CDLP alarm is triggered, you must add new Storage Nodes.

Before you begin

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

About this task

StorageGRID reserves a certain amount of space on each Storage Node for metadata. This space is known as the Metadata Reserved Space (CAWM). The Metadata Reserved Space is subdivided into the space available for object metadata (the Metadata Allowed Space) and the space required for essential database operations, such as compaction and repair.



If object metadata consumes more than 100% of the Metadata Allowed Space, database operations cannot run efficiently and errors will occur.

StorageGRID uses the following Prometheus metric to measure how full the Metadata Allowed Space is:

```
storagegrid_storage_utilization_metadata_bytes/  
storagegrid_storage_utilization_metadata_allowed_bytes
```

When this Prometheus expression reaches certain thresholds, the **Low metadata storage** alert is triggered.

Note: This Prometheus expression is equivalent to the legacy CDLP (Metadata Used Space (Percent)) attribute, which triggers the legacy CDLP alarm at the same thresholds.

- **Minor:** Object metadata is using 70% or more of the Metadata Allowed Space. You should add new Storage Nodes as soon as possible.
- **Major:** Object metadata is using 90% or more of the Metadata Allowed Space. You must add new Storage Nodes immediately.



Attention: When the legacy CDLP alarm is triggered at the major level, a warning appears on the Dashboard. If this warning appears, you must add new Storage Nodes immediately. You must never allow object metadata to use more than 100% of the allowed space.

- **Critical:** Object metadata is using 100% or more of the Metadata Allowed Space and is starting to consume the space required for essential database operations. You must stop the ingest of new objects, and you must add new Storage Nodes immediately.

When you add the new nodes, StorageGRID evenly distributes object metadata across all Storage Nodes at each site to increase the overall metadata capacity of the grid. No other user action is required. The **Low metadata storage** alert (and the legacy CDLP alarm) are cleared.

Note: Because StorageGRID keeps all object metadata at every site, the metadata capacity of the entire grid is limited by the metadata capacity of the smallest site. If you need to add metadata capacity to one site, you should also expand any other sites by adding the same number of Storage Nodes.

In the following example, object metadata is using more than 100% of the Metadata Allowed Space. This is a critical situation, which will result in inefficient database operation and errors.

The following Storage Nodes are using more than 90% of the space allowed for object metadata:

Node	% Used	Used	Allowed
DC1-S2-227	104.51%	6.73 GB	6.44 GB
DC1-S3-228	104.36%	6.72 GB	6.44 GB
DC2-S2-233	104.20%	6.71 GB	6.44 GB
DC1-S1-226	104.20%	6.71 GB	6.44 GB
DC2-S3-234	103.43%	6.66 GB	6.44 GB

Undesirable results can occur if object metadata uses more than 100% of the allowed space. You must add new Storage Nodes immediately or contact support.



Attention: If the first storage volume is smaller than the Metadata Reserved Space (for example, in a non-production environment), the calculation for the **Low metadata storage** alert (and the legacy CDLP alarm) might be inaccurate.

Steps

1. To view details about the **Low metadata storage** alert:
 - a. Select **Alerts > Current**.
 - b. From the table of alerts, expand the **Low metadata storage** alert group, if required, and select the specific alert you want to view.
 - c. Review the details in the alert dialog box.
2. To view details about the legacy CDLP alarm:
 - a. Select **Support**. Then, in the Alarms (legacy) section of the menu, select **Current Alarms**.
 - b. Review the details for the alarm in the table.

CDLP (Metadata Used Space (Percent)) alarms of all severities (minor, major, and critical) are displayed on this page.

Current Alarms

Last Refreshed: 2018-05-04 10:56:02 MDT

Show Acknowledged Alarms (1 - 18 of 18)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value
Minor	CDLP (Metadata Used Space (Percent))	Data Center 1/SGA-Lab11/DDS	The metadata store is more than 70% full. You should add new Storage Nodes as soon as possible.	2018-05-04 10:53:51 MDT	89.319 %	89.319 %

Note: You can also see the CDLP alarm for a Storage Node by selecting **Support > Grid Topology > Storage Node > DDS > Data Store > Alarms**.

- If a major or critical **Low metadata storage** alert (or legacy CDLP alarm) has been triggered, perform an expansion to add Storage Nodes immediately.

When you add the new nodes, the system automatically rebalances object metadata across all Storage Nodes at each site, and the alerts and alarms clear.

Troubleshooting the Services: Status - Cassandra (SVST) alarm

The Services: Status - Cassandra (SVST) alarm indicates that you might need to rebuild the Cassandra database for a Storage Node. Cassandra is used as the metadata store for StorageGRID.

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have the `Passwords.txt` file.

About this task

If Cassandra is stopped for more than 15 days (for example, the Storage Node is powered off), Cassandra will not start when the node is brought back online. You must rebuild the Cassandra database for the affected DDS service.



Attention: If two or more of the Cassandra database services are down for more than 15 days, contact technical support, and do not proceed with the steps below.

Steps

- Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
- Select **site > Storage Node > SSM > Services > Alarms > Main** to display alarms.

This example shows that the SVST alarm was triggered.

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge	Time
Minor	SVST (Services: Status - Cassandra)	Not Running	2014-08-14 14:56:26 PDT	Not Running	Not Running	<input type="checkbox"/>	

The SSM Services Main page also indicates that Cassandra is not running.

Service	Version	Status	Threads	Load	Memory
Account Service	10.4.0-20161224.0333.803cd91	Running	7	0.002 %	12 MB
Administrative Domain Controller (ADC)	10.4.0-20170329.0039.8800cae	Running	52	0.14 %	63.1 MB
Cassandra	4.6.12-1.byc.0-20170308.0109.ba3598a	Not Running	0	0 %	0 B
Content Management System (CMS)	10.4.0-20170220.1846.1a76aed	Running	18	0.055 %	20.6 MB
Distributed Data Store (DDS)	10.4.0-20170329.0039.8800cae	Running	104	1.301 %	76 MB
Identity Service	10.4.0-20170203.2038.a457d45	Running	6	0 %	8.75 MB
Keystone Service	10.4.0-20170104.1815.6e52138	Running	5	0 %	7.77 MB
Local Distribution Router (LDR)	10.4.0-20170329.0039.8800cae	Running	109	0.218 %	96.6 MB
Server Manager	10.4.0-20170306.2303.9649faf	Running	4	3.58 %	19.1 MB

3. Try restarting Cassandra from the Storage Node:

a. Log in to the grid node:

- i. Enter the following command: `ssh admin@grid_node_IP`
- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

b. Enter:

```
/etc/init.d/cassandra status
```

c. If Cassandra is not running, restart it:

```
/etc/init.d/cassandra restart
```

4. If Cassandra does not restart, determine how long Cassandra has been down. If Cassandra has been down for longer than 15 days, you must rebuild the Cassandra database.



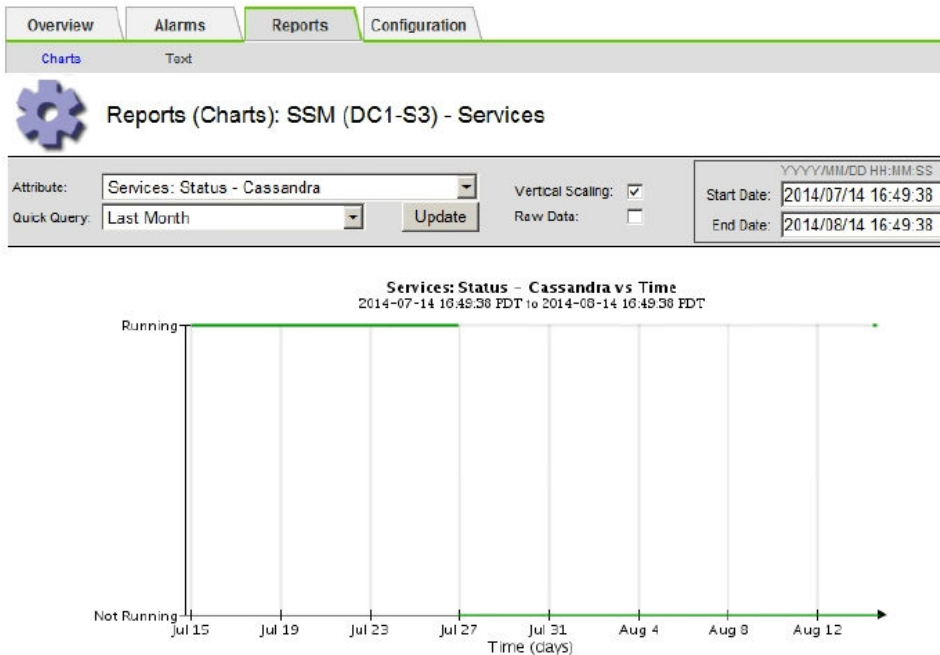
Attention: If two or more of the Cassandra database services are down, contact technical support, and do not proceed with the steps below.

You can determine how long Cassandra has been down by charting it or by reviewing the `servermanager.log` file.

5. To chart Cassandra:

- a. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**. Then select `site > Storage Node > SSM > Services > Reports > Charts`.
- b. Select **Attribute > Service: Status - Cassandra**.
- c. For **Start Date**, enter a date that is at least 16 days before the current date. For **End Date**, enter the current date.
- d. Click **Update**.
- e. If the chart shows Cassandra as being down for more than 15 days, rebuild the Cassandra database.

The following chart example shows that Cassandra has been down for at least 17 days.



6. To review the `servermanager.log` file on the Storage Node:

a. Log in to the grid node:

- i. Enter the following command: `ssh admin@grid_node_IP`
- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

b. Enter:

```
cat /var/local/log/servermanager.log
```

The contents of the `servermanager.log` file are displayed.

If Cassandra has been down for longer than 15 days, the following message is displayed in the `servermanager.log` file:

```
"2014-08-14 21:01:35 +0000 | cassandra | cassandra not
started because it has been offline for longer than
its 15 day grace period - rebuild cassandra
```

c. Make sure the timestamp of this message is the time when you attempted restarting Cassandra as instructed in step 3.

There can be more than one entry for Cassandra; you must locate the most recent entry.

d. If Cassandra has been down for longer than 15 days, you must rebuild the Cassandra database.

For instructions, see "Recovering from a single Storage Node down more than 15 days" in the recovery and maintenance instructions.

e. Contact technical support if alarms do not clear after Cassandra is rebuilt.

Related information

[Recovery and maintenance](#)

Troubleshooting Cassandra Out of Memory errors (SMTT alarm)

A Total Events (SMTT) alarm is triggered when the Cassandra database has an out-of-memory error. If this error occurs, contact technical support to work through the issue.

About this task

If an out-of-memory error occurs for the Cassandra database, a heap dump is created, a Total Events (SMTT) alarm is triggered, and the Cassandra Heap Out Of Memory Errors count is incremented by one.

Steps

1. To view the event, select **Nodes** > *grid node* > **Events**.
2. Verify that the Cassandra Heap Out Of Memory Errors count is 1 or greater.
3. Go to `/var/local/core/`, compress the `Cassandra.hprof` file, and send it to technical support.
4. Make a backup of the `Cassandra.hprof` file, and delete it from the `/var/local/core/` directory.

This file can be as large as 24 GB, so you should remove it to free up space.

5. Once the issue is resolved, click **Reset event counts**.

Note: To reset event counts, you must have the Grid Topology Page Configuration permission.

Related tasks

[Resetting event counts](#) on page 134

After resolving system events, you can reset event counts to zero.

Troubleshooting certificate errors

If you see a security or certificate issue when you try to connect to StorageGRID using a web browser or an S3 or Swift client, you should check the certificate.

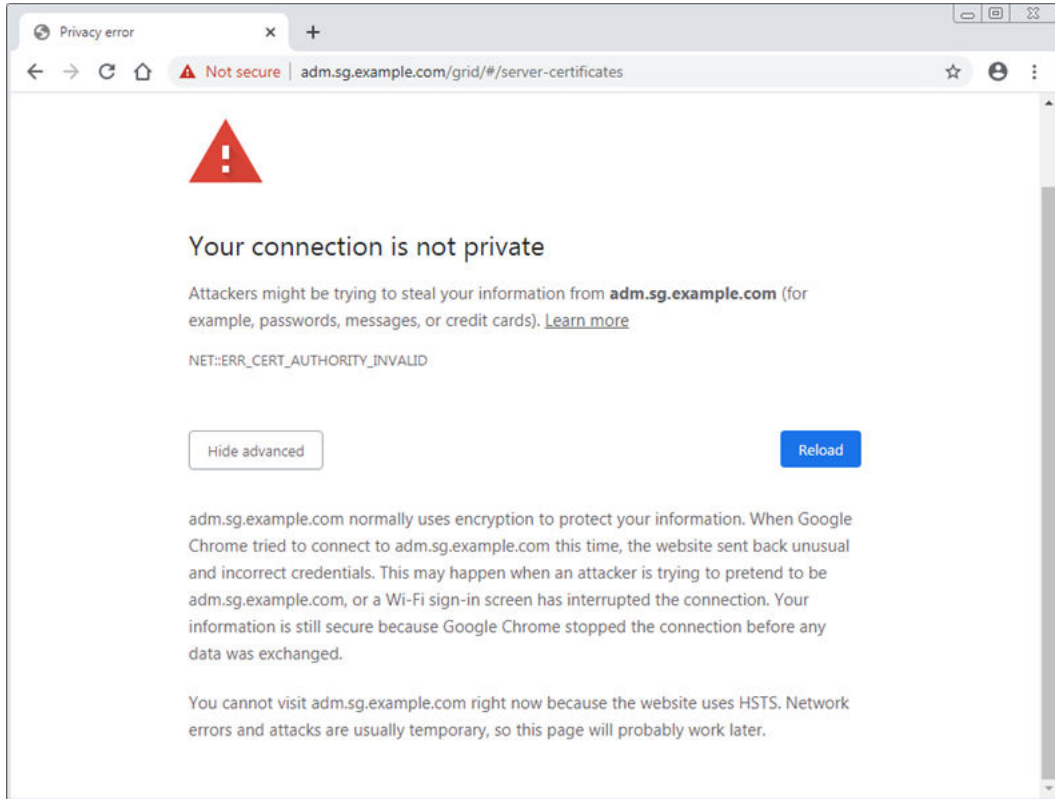
About this task

Certificate errors can cause problems when you try to connect to StorageGRID using the Grid Manager, Grid Management API, the Tenant Manager, or the Tenant Management API, or when you try to connect with an S3 or Swift client.

If you are accessing the Grid Manager or Tenant Manager using a domain name instead of an IP address, the browser shows a certificate error without an option to bypass if either of the following occurs:

- Your custom management interface server certificate expires.
- You revert from a custom management interface server certificate to the default server certificate.

The following example shows a certificate error when the custom management interface server certificate expired:



Note: To ensure that operations are not disrupted by a failed server certificate, the **Expiration of server certificate for Management Interface** alert and the legacy Management Interface Certificate Expiry (MCEP) alarm are both triggered when the server certificate is about to expire. As required, you can view the number of days until the current service certificate expires by selecting **Support > Grid Topology > primary Admin Node > CMN > Resources**.

Steps

1. Check if the certificate used for the connection has expired.
2. Check the validity period of the certificate.
Some web browsers and S3 or Swift clients do not accept certificates with a validity period greater than 825 days.
3. Ensure that the Subject Alternative Name (SAN) of the certificate is populated, and that the SAN matches the IP address or host name of the node that you are connecting to.
4. If you are attempting to connect to StorageGRID using a domain name and a certificate error appears, follow these steps:
 - a. Enter the IP address of the Admin Node instead of the domain name to bypass the connection error and access the Grid Manager.
 - b. From the Grid Manager, select **Configuration > Server Certificate** to install a new custom certificate or continue with the default certificate.
 - c. In the instructions for administering StorageGRID, see the steps for configuring a custom server certificate for the Grid Manager and the Tenant Manager.

Related information

[Administering StorageGRID](#)

Troubleshooting Admin Node and user interface issues

There are several tasks you can perform to help determine the source of issues related to Admin Nodes and the StorageGRID user interface.

Troubleshooting sign-on errors

If you experience an error when you are signing in to a StorageGRID Admin Node, your system might have an issue with the identity federation configuration, a networking or hardware problem, an issue with Admin Node services, or an issue with the Cassandra database on connected Storage Nodes.

Before you begin

- You must have the `Passwords.txt` file.
- You must have specific access permissions.

About this task

Use these troubleshooting guidelines if you see any of the following error messages when attempting to sign in to an Admin Node:

- `Your credentials for this account were invalid. Please try again.`
- `Waiting for services to start...`
- `Internal server error. The server encountered an error and could not complete your request. Please try again. If the problem persists, contact Technical Support.`
- `Unable to communicate with server. Reloading page...`

Steps

1. Wait 10 minutes, and try signing in again.
If the error is not resolved automatically, go to the next step.
2. If your StorageGRID system has more than one Admin Node, try signing in to the Grid Manager from another Admin Node.
 - If you are able to sign in, you can use the **Dashboard, Nodes, Alerts, and Support > Grid Topology** options to help determine the cause of the error.
 - If you have only one Admin Node or you still cannot sign in, go to the next step.
3. Determine if the node's hardware is offline.
4. If single sign-on (SSO) is enabled for your StorageGRID system, refer to the steps for configuring single sign-on, in the instructions for administering StorageGRID.
You might need to temporarily disable and re-enable SSO for a single Admin Node to resolve any issues.
Note: If SSO is enabled, you cannot sign on using a restricted port. You must use port 443.
5. Determine if the account you are using belongs to a federated user.
If the federated user account is not working, try signing in to the Grid Manager as a local user, such as root.
 - If the local user can sign in:
 - a. Review any displayed alarms.
 - b. Select **Configuration > Identity Federation**.
 - c. Click **Test Connection** to validate your connection settings for the LDAP server.
 - d. If the test fails, resolve any configuration errors.

- If the local user cannot sign in and you are confident that the credentials are correct, go to the next step.

6. Use Secure Shell (ssh) to log in to the Admin Node:
 - a. Enter the following command: `ssh admin@Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.

When you are logged in as root, the prompt changes from `$` to `#`.

7. View the status of all services running on the grid node:

storagegrid-status

Make sure the `nms`, `mi`, `nginx`, and `mgmt api` services are all running.

The output is updated immediately if the status of a service changes.

```
$ storagegrid-status
Host Name                99-211
IP Address               10.96.99.211
Operating System Kernel 4.19.0                 Verified
Operating System Environment Debian 10.1             Verified
StorageGRID Webscale Release 11.4.0                 Verified
Networking                Verified
Storage Subsystem         Verified
Database Engine           5.5.9999+default      Running
Network Monitoring       11.4.0                 Running
Time Synchronization     1:4.2.8p10+dfsg      Running
ams                       11.4.0                 Running
cmn                       11.4.0                 Running
nms                       11.4.0                 Running
ssm                       11.4.0                 Running
mi                        11.4.0                 Running
dynip                    11.4.0                 Running
nginx                     1.10.3                 Running
tomcat                    9.0.27                 Running
grafana                   6.4.3                 Running
mgmt api                  11.4.0                 Running
prometheus                11.4.0                 Running
persistence               11.4.0                 Running
ade exporter              11.4.0                 Running
alertmanager              11.4.0                 Running
attrDownPurge             11.4.0                 Running
attrDownSamp1             11.4.0                 Running
attrDownSamp2             11.4.0                 Running
node exporter              0.17.0+ds              Running
sg snmp agent             11.4.0                 Running
```

8. Confirm that the Apache web server is running:

```
# service apache2 status
```

9. Use Lumberjack to collect logs:

```
# /usr/local/sbin/lumberjack.rb
```

If the failed authentication happened in the past, you can use the `-start` and `-end` Lumberjack script options to specify the appropriate time range. Use `lumberjack -h` for details on these options.

The output to the terminal indicates where the log archive has been copied.

10. Review the following logs:

- `/var/local/log/bycast.log`
- `/var/local/log/bycast-err.log`
- `/var/local/log/nms.log`
- `**/*commands.txt`

11. If you could not identify any issues with the Admin Node, issue either of the following commands to determine the IP addresses of the three Storage Nodes that run the ADC service at your site. Typically, these are the first three Storage Nodes that were installed at the site.

```
# cat /etc/hosts
```

```
# vi /var/local/gpt-data/specs/grid.xml
```

Admin Nodes use the ADC service during the authentication process.

12. From the Admin Node, log in to each of the ADC Storage Nodes, using the IP addresses you identified.
- a. Enter the following command: `ssh admin@grid_node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.
- When you are logged in as root, the prompt changes from `$` to `#`.

13. View the status of all services running on the grid node:

```
storagegrid-status
```

Make sure the `idnt`, `acct`, `nginx`, and `cassandra` services are all running.

14. Repeat steps 9 on page 191 and 10 on page 191 to review the logs on the Storage Nodes.
15. If you are unable to resolve the issue, contact technical support.
Provide the logs you collected to technical support.

Related concepts

[Log files reference](#) on page 246

The following sections list the logs 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.

Related information

[Administering StorageGRID](#)

Troubleshooting user interface issues

You might see issues with the Grid Manager or the Tenant Manager after upgrading to a new version of StorageGRID software.

Web interface does not respond as expected

The Grid Manager or the Tenant Manager might not respond as expected after StorageGRID software is upgraded. For example, after you use the Grid Manager to acknowledge an alarm and click **Apply Changes**, the change might not be saved.

If you experience issues with the web interface:

- Make sure you are using a supported browser.
- Clear your web browser cache.

Clearing the cache removes outdated resources used by the previous version of StorageGRID software, and permits the user interface to operate correctly again. For instructions, see the documentation for your web browser.

Related reference

[Web browser requirements](#) on page 5

You must use a supported web browser.

Related information

[Administering StorageGRID](#)

Checking the status of an unavailable Admin Node

If the StorageGRID system includes multiple Admin Nodes, you can use another Admin Node to check the status of an unavailable Admin Node.

Before you begin

You must have specific access permissions.

Steps

1. From an available Admin Node, sign in to the Grid Manager using a supported browser.
2. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
3. Select **Site > unavailable Admin Node > SSM > Services > Overview > Main**.
4. Look for services that have a status of Not Running and that might also be displayed in blue.

The screenshot shows the 'Overview: SSM (MM-10-224-4-81-ADM1) - Services' page. The 'Management Application Program Interface (mgmt-api)' service is highlighted in blue and has a status of 'Not Running'.

Service	Version	Status	Threads	Load	Memory
Audit Management System (AMS)	10.4.0-20170113.2207.3ec2cd0	Running	52	0.043 %	35.7 MB
CIFS Filesharing (nmbd)	2:4.2.14+dfsg-0+deb8u2	Running	1	0 %	5.5 MB
CIFS Filesharing (smbd)	2:4.2.14+dfsg-0+deb8u2	Running	1	0 %	14.5 MB
CIFS Filesharing (winbindd)	2:4.2.14+dfsg-0+deb8u2	Not Running	0	0 %	0 B
Configuration Management Node (CMN)	10.4.0-20170113.2207.3ec2cd0	Running	52	0.055 %	41.3 MB
Database Engine	5.5.53-0+deb8u1	Running	47	0.354 %	1.33 GB
Grid Deployment Utility Server	10.4.0-20170112.2125.c4253bb	Running	3	0 %	32.8 MB
Management Application Program Interface (mgmt-api)	10.4.0-20170113.2136.07c4997	Not Running	0	0 %	0 B
NFS Filesharing	10.4.0-20161224.0333.803cd91	Not Running	0	0 %	0 B
NMS Data Cleanup	10.4.0-20161224.0333.803cd91	Running	22	0.008 %	52.4 MB
NMS Data Downsampler 1	10.4.0-20161224.0333.803cd91	Running	22	0.049 %	195 MB
NMS Data Downsampler 2	10.4.0-20161224.0333.803cd91	Running	22	0.009 %	157 MB
NMS Processing Engine	10.4.0-20161224.0333.803cd91	Running	40	0.132 %	200 MB

5. Determine if alarms have been triggered.
6. Take the appropriate actions to resolve the issue.

Related information

[Administering StorageGRID](#)

Troubleshooting network, hardware, and platform issues

There are several tasks you can perform to help determine the source of issues related to StorageGRID network, hardware, and platform issues.

Choices

- [Troubleshooting "422: Unprocessable Entity" errors](#) on page 194
- [Troubleshooting the Grid Network MTU mismatch alert](#) on page 195
- [Troubleshooting the Network Receive Error \(NRER\) alarm](#) on page 196
- [Troubleshooting time synchronization errors](#) on page 197
- [Linux: Network connectivity issues](#) on page 198
- [Linux: Node status is "orphaned"](#) on page 198
- [Linux: Troubleshooting IPv6 support](#) on page 199

Troubleshooting "422: Unprocessable Entity" errors

The error 422: Unprocessable Entity can occur in a number of circumstances. Check the error message to determine what caused your issue.

About this task

If you see one of the listed error messages, take the recommended action.

Error message	Root cause and corrective action
<pre>422: Unprocessable Entity Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration. Unable to authenticate, please verify your username and password: LDAP Result Code 8 "Strong Auth Required": 00002028: LdapErr: DSID-0C090256, comment: The server requires binds to turn on integrity checking if SSL\TLS are not already active on the connection, data 0, v3839</pre>	<p>This message might occur if you select the Do not use TLS option for Transport Layer Security (TLS) when configuring identity federation using Windows Active Directory (AD).</p> <p>Using the Do not use TLS option is not supported for use with AD servers that enforce LDAP signing. You must select either the Use STARTTLS option or the Use LDAPS option for TLS.</p>
<pre>422: Unprocessable Entity Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration.Unable to begin TLS, verify your certificate and TLS configuration: LDAP Result Code 200 "Network Error": TLS handshake failed (EOF)</pre>	<p>This message appears if you try to use an unsupported cipher to make a Transport Layer Security (TLS) connection from StorageGRID to an external system used for identify federation or Cloud Storage Pools.</p> <p>Check the ciphers that are offered by the external system. The system must use one of the ciphers supported by StorageGRID for outgoing TLS connections, as shown in the instructions for administering StorageGRID.</p>

Related information

[Administering StorageGRID](#)

Troubleshooting the Grid Network MTU mismatch alert

The **Grid Network MTU mismatch** alert is triggered when the maximum transmission unit (MTU) setting for the Grid Network interface (eth0) differs significantly across nodes in the grid.

About this task

The differences in MTU settings could indicate that some, but not all, eth0 networks are configured for jumbo frames. An MTU size mismatch of greater than 1000 might cause network performance problems.

Steps

1. List the MTU settings for eth0 on all nodes.

- Use the query provided in the Grid Manager.
- Navigate to `<primary Admin Node IP address>/metrics/graph` and enter the following query:

```
node_network_mtu_bytes{interface='eth0'}
```

2. Modify the MTU settings as necessary to ensure they are the same for the Grid Network interface (eth0) on all nodes.

- For appliance nodes, see the installation and maintenance instructions for your appliance.
- For Linux- and VMware-based nodes, use the following command:

```
/usr/sbin/change-mtu.py [-h] [-n node] mtu network [network...]
```

Example: `change-mtu.py -n node 1500 grid admin`

Note: On Linux-based nodes, if the desired MTU value for the network in the container exceeds the value already configured on the host interface, you must first configure the host interface to have the desired MTU value, and then use the `change-mtu.py` script to change the MTU value of the network in the container.

Use the following arguments for modifying the MTU on Linux- or VMware-based nodes.

Positional arguments	Description
mtu	The MTU to set. Must be in the range 1280 to 9216.
network	The networks to apply the MTU to. Include one or more of the following network types: <ul style="list-style-type: none"> ◦ grid ◦ admin ◦ client
Optional arguments	Description
-h, - help	Show the help message and exit.
-n node, --node node	The node. The default is the local node.

Related information

[SG100 and SG1000 appliance installation and maintenance](#)

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)

Troubleshooting the Network Receive Error (NRER) alarm

Network Receive Error (NRER) alarms can be caused by connectivity issues between StorageGRID and your network hardware. In some cases, NRER errors can clear without manual intervention. If the errors do not clear, take the recommended actions.

About this task

NRER alarms can be caused by the following issues with networking hardware that connects to StorageGRID:

- Forward error correction (FEC) is required and not in use
- Switch port and NIC MTU mismatch
- High link error rates
- NIC ring buffer overrun

Steps

1. Follow the troubleshooting steps for all potential causes of the NRER alarm given your network configuration.

- If the error is caused by FEC mismatch, perform the following steps:

Note: These steps are applicable only for NRER errors caused by FEC mismatch on StorageGRID appliances.

- a. Check the FEC status of the port in the switch attached to your StorageGRID appliance.
- b. Check the physical integrity of the cables from the appliance to the switch.
- c. If you want to change FEC settings to try to resolve the NRER alarm, first ensure that the appliance is configured for **Auto** mode on the Link Configuration page of the StorageGRID Appliance Installer (see the installation and maintenance instructions for your appliance). Then, change the FEC settings on the switch ports. The StorageGRID appliance ports will adjust their FEC settings to match, if possible.

(You cannot configure FEC settings on StorageGRID appliances. Instead, the appliances attempt to discover and mirror the FEC settings on the switch ports they are connected to. If the links are forced to 25-GbE or 100-GbE network speeds, the switch and NIC might fail to negotiate a common FEC setting. Without a common FEC setting, the network will fall back to "no-FEC" mode. When FEC is not enabled, the connections are more susceptible to errors caused by electrical noise.)

Note: StorageGRID appliances support Firecode (FC) and Reed Solomon (RS) FEC, as well as no FEC.

- If the error is caused by a switch port and NIC MTU mismatch, check that the MTU size configured on the node is the same as the MTU setting for the switch port.

The MTU size configured on the node might be smaller than the setting on the switch port the node is connected to. If a StorageGRID node receives an Ethernet frame larger than its MTU, which is possible with this configuration, the NRER alarm might be reported. If you believe this is what is happening, either change the MTU of the switch port to match the StorageGRID network interface MTU, or change the MTU of the StorageGRID network interface to match the switch port, depending on your end-to-end MTU goals or requirements.



Attention: For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The **Grid Network MTU mismatch** alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The

MTU values do not have to be the same for all network types.

Note: To change the MTU setting, see the installation and maintenance guide for your appliance.

- If the error is caused by high link error rates, perform the following steps:
 - a. Enable FEC, if not already enabled.
 - b. Verify that your network cabling is of good quality and is not damaged or improperly connected.
 - c. If the cables do not appear to be the problem, contact technical support.

Note: You might notice high error rates in an environment with high electrical noise.

- If the error is a NIC ring buffer overrun, contact technical support.
The ring buffer can be overrun when the StorageGRID system is overloaded and unable to process network events in a timely manner.

2. After you resolve the underlying problem, reset the error counter.
 - a. Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**.
 - b. Select *site* > *grid node* > **SSM** > **Resources** > **Configuration** > **Main**.
 - c. Select **Reset Receive Error Count** and click **Apply Changes**.

Related tasks

[Troubleshooting the Grid Network MTU mismatch alert](#) on page 195

The **Grid Network MTU mismatch** alert is triggered when the maximum transmission unit (MTU) setting for the Grid Network interface (eth0) differs significantly across nodes in the grid.

Related reference

[Alarms reference \(legacy system\)](#) on page 222

Related information

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)

[SG100 and SG1000 appliance installation and maintenance](#)

Troubleshooting time synchronization errors

You might see issues with time synchronization in your grid.

If you encounter time synchronization problems, verify that you have specified at least four external NTP sources, each providing a Stratum 3 or better reference, and that all external NTP sources are operating normally and are accessible by your StorageGRID nodes.

Note: When specifying the external NTP source for a production-level StorageGRID installation, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

Related information

[Recovery and maintenance](#)

Linux: Network connectivity issues

You might see issues with network connectivity for StorageGRID grid nodes hosted on Linux hosts.

Promiscuous mode

If you are deploying StorageGRID nodes on Linux hosts and you observe that the node containers are unable to communicate over the network even though the host networking is working, verify that *promiscuous mode* is enabled in the hypervisor. Promiscuous mode is disabled by default for VMware hypervisors.

Note: This information only applies to containers deployed on Linux hosts; VMware virtual machine nodes do not require promiscuous mode.

Linux: Node status is "orphaned"

A Linux node in an orphaned state usually indicates that either the storagegrid service or the StorageGRID node daemon controlling the node's container died unexpectedly.

About this task

If a Linux node reports that it is in an orphaned state, you should:

- Check logs for errors and messages.
- Attempt to start the node again.
- If necessary, use Docker commands to stop the existing node container.
- Restart the node.

Steps

1. Check logs for both the service daemon and the orphaned node for obvious errors or messages about exiting unexpectedly.
2. Log in to the host as root or using an account with sudo permission.
3. Attempt to start the node again by running the following command:

```
$ sudo storagegrid node start node-name
```

```
$ sudo storagegrid node start DC1-S1-172-16-1-172
```

If the node is orphaned, the response is

```
Not starting ORPHANED node DC1-S1-172-16-1-172
```

4. From Linux, stop the Docker container and any controlling storagegrid-node processes:

```
sudo docker stop --time seconds container-name
```

For *seconds*, enter the number of seconds you want to wait for the container to stop (typically 15 minutes or less).

```
sudo docker stop --time 900 storagegrid-DC1-S1-172-16-1-172
```

5. Restart the node:

```
storagegrid node start node-name
```

```
storagegrid node start DC1-S1-172-16-1-172
```

Linux: Troubleshooting IPv6 support

You might need to enable IPv6 support in the kernel if you have installed StorageGRID nodes on Linux hosts and you notice that IPv6 addresses have not been assigned to the node containers as expected.

About this task

You can see the IPv6 address that has been assigned to a grid node in the following locations in the Grid Manager:

- Select **Nodes**, and select the node. Then, click **Show more** next to **IP Addresses** on the Overview tab.

DC1-S1 (Storage Node)

The screenshot shows the Grid Manager interface for a Storage Node named DC1-S1. The 'Node Information' section is expanded to show the 'IP Addresses' section. Below this, a table lists the network interfaces and their assigned IP addresses. The IPv6 address for the 'eth0' interface is highlighted with an orange circle.

Interface	IP Address
eth0	10.96.106.102
eth0	fe80::250:56ff:fea7:5c83

- Select **Support**. Then, in the Tools section of the menu, select **Grid Topology**. Then, select **node > SSM > Resources**. If an IPv6 address has been assigned, it is listed below the IPv4 address in the **Network Addresses** section.

If the IPv6 address is not shown and the node is installed on a Linux host, follow these steps to enable IPv6 support in the kernel.

Steps

1. Log in to the host as root or using an account with sudo permission.
2. Run the following command:

```
sysctl net.ipv6.conf.all.disable_ipv6
```

```
root@SG:~ # sysctl net.ipv6.conf.all.disable_ipv6
```

The result should be 0.

```
net.ipv6.conf.all.disable_ipv6 = 0
```

Note: If the result is not 0, see the documentation for your operating system for changing sysctl settings. Then, change the value to 0 before continuing.

3. Enter the StorageGRID node container:

```
storagegrid node enter node-name
```
4. Run the following command:

```
sysctl net.ipv6.conf.all.disable_ipv6
```

```
root@DC1-S1:~ # sysctl net.ipv6.conf.all.disable_ipv6
```

The result should be 1.

```
net.ipv6.conf.all.disable_ipv6 = 1
```

Note: If the result is not 1, this procedure does not apply. Contact technical support.

5. Exit the container:

```
exit
```

```
root@DC1-S1:~ # exit
```

6. As root, edit the following file: `/var/lib/storagegrid/settings/sysctl.d/net.conf`.

```
sudo vi /var/lib/storagegrid/settings/sysctl.d/net.conf
```

7. Locate the following two lines and remove the comment tags. Then, save and close the file.

```
net.ipv6.conf.all.disable_ipv6 = 0
```

```
net.ipv6.conf.default.disable_ipv6 = 0
```

8. Run these commands to restart the StorageGRID container:

```
storagegrid node stop node-name
```

```
storagegrid node start node-name
```


Alerts reference

The following table lists all default StorageGRID alerts. As required, you can create custom alert rules to fit your system management approach.

See information about the commonly used Prometheus metrics to learn about the metrics used in some of these alerts.

Alert name	Description and recommended actions
Appliance battery expired	<p>The battery in the appliance’s storage controller has expired.</p> <ol style="list-style-type: none"> 1. Replace the battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance 2. If this alert persists, contact technical support.
Appliance battery failed	<p>The battery in the appliance’s storage controller has failed.</p> <ol style="list-style-type: none"> 1. Replace the battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance 2. If this alert persists, contact technical support.
Appliance battery has insufficient learned capacity	<p>The battery in the appliance’s storage controller has insufficient learned capacity.</p> <ol style="list-style-type: none"> 1. Replace the battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance 2. If this alert persists, contact technical support.
Appliance battery near expiration	<p>The battery in the appliance’s storage controller is nearing expiration.</p> <ol style="list-style-type: none"> 1. Replace the battery soon. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance 2. If this alert persists, contact technical support.

Alert name	Description and recommended actions
Appliance battery removed	<p>The battery in the appliance's storage controller is missing.</p> <ol style="list-style-type: none"> 1. Install a battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance 2. If this alert persists, contact technical support.
Appliance battery too hot	<p>The battery in the appliance's storage controller is overheated.</p> <ol style="list-style-type: none"> 1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert. 2. Investigate possible reasons for the temperature increase, such as a fan or HVAC failure. 3. If this alert persists, contact technical support.
Appliance cache backup device failed	<p>A persistent cache backup device has failed.</p> <ol style="list-style-type: none"> 1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert. 2. Contact technical support.
Appliance cache backup device insufficient capacity	<p>There is insufficient cache backup device capacity. Contact technical support.</p>
Appliance cache backup device write-protected	<p>A cache backup device is write-protected. Contact technical support.</p>
Appliance cache memory size mismatch	<p>The two controllers in the appliance have different cache sizes. Contact technical support.</p>
Appliance compute controller chassis temperature too high	<p>The temperature of the compute controller in a StorageGRID appliance has exceeded a nominal threshold.</p> <ol style="list-style-type: none"> 1. Check the hardware components for overheating conditions, and follow the recommended actions: <ul style="list-style-type: none"> • If you have an SG100, SG1000, or SG6000, use the BMC. • If you have an SG5600 or SG5700, use SANtricity System Manager. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance • SG100 and SG1000 appliance installation and maintenance

Alert name	Description and recommended actions
Appliance compute controller CPU temperature too high	<p>The temperature of the CPU in the compute controller in a StorageGRID appliance has exceeded a nominal threshold.</p> <ol style="list-style-type: none"> 1. Check the hardware components for overheating conditions, and follow the recommended actions: <ul style="list-style-type: none"> • If you have an SG100, SG1000, or SG6000, use the BMC. • If you have an SG5600 or SG5700, use SANtricity System Manager. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance • SG100 and SG1000 appliance installation and maintenance
Appliance compute controller needs attention	<p>A hardware fault has been detected in the compute controller of a StorageGRID appliance.</p> <ol style="list-style-type: none"> 1. Check the hardware components for errors, and follow the recommended actions: <ul style="list-style-type: none"> • If you have an SG100, SG1000, or SG6000, use the BMC. • If you have an SG5600 or SG5700, use SANtricity System Manager. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance • SG100 and SG1000 appliance installation and maintenance
Appliance compute controller power supply A has a problem	<p>Power supply A in the compute controller has a problem. This alert might indicate that the power supply has failed or that it has a problem providing power.</p> <ol style="list-style-type: none"> 1. Check the hardware components for errors, and follow the recommended actions: <ul style="list-style-type: none"> • If you have an SG100, SG1000, or SG6000, use the BMC. • If you have an SG5600 or SG5700, use SANtricity System Manager. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance • SG100 and SG1000 appliance installation and maintenance

Alert name	Description and recommended actions
Appliance compute controller power supply B has a problem	<p>Power supply B in the compute controller has a problem.</p> <p>This alert might indicate that the power supply has failed or that it has a problem providing power.</p> <ol style="list-style-type: none"> 1. Check the hardware components for errors, and follow the recommended actions: <ul style="list-style-type: none"> • If you have an SG100, SG1000, or SG6000, use the BMC. • If you have an SG5600 or SG5700, use SANtricity System Manager. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance • SG100 and SG1000 appliance installation and maintenance
Appliance compute hardware monitor service stalled	<p>The service that monitors storage hardware status has stopped reporting data.</p> <ol style="list-style-type: none"> 1. Check the status of the eos-system-status service in the base-os. 2. If the service is in a stopped or error state, restart the service. 3. If this alert persists, contact technical support.
Appliance flash cache drives non-optimal	<p>The drives used for the SSD cache are non-optimal.</p> <ol style="list-style-type: none"> 1. Replace the SSD cache drives. See the appliance installation and maintenance instructions. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance 2. If this alert persists, contact technical support.
Appliance interconnect/battery canister removed	<p>The interconnect/battery canister is missing.</p> <ol style="list-style-type: none"> 1. Replace the battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance 2. If this alert persists, contact technical support.

Alert name	Description and recommended actions
Appliance overall power supply degraded	<p>The power of a StorageGRID appliance has deviated from the recommended operating voltage.</p> <ol style="list-style-type: none"> 1. Check the status of power supply A and B to determine which power supply is operating abnormally, and follow the recommended actions: <ul style="list-style-type: none"> • If you have an SG100, SG1000, or SG6000, use the BMC. • If you have an SG5600 or SG5700, use SANtricity System Manager. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance • SG100 and SG1000 appliance installation and maintenance
Appliance storage connectivity degraded	<p>There is a problem with one or more connections between the compute controller and storage controller.</p> <ol style="list-style-type: none"> 1. Go to the appliance to check the port indicator lights. 2. If a port's lights are off, confirm the cable is properly connected. As needed, replace the cable. 3. Wait up to five minutes. <p style="margin-left: 40px;">Note: If a second cable needs to be replaced, do not unplug it for at least 5 minutes. Otherwise, the root volume might become read-only, which requires a hardware restart.</p> 4. From the Grid Manager, select Nodes. Then, select the Hardware tab of the node that had the problem. Verify that the alert condition has resolved.
Appliance storage controller A failure	<p>Storage controller A in a StorageGRID appliance has failed.</p> <ol style="list-style-type: none"> 1. Use SANtricity System Manager to check hardware components, and follow the recommended actions. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance
Appliance storage controller B failure	<p>Storage controller B in a StorageGRID appliance has failed.</p> <ol style="list-style-type: none"> 1. Use SANtricity System Manager to check hardware components, and follow the recommended actions. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance


Alert name	Description and recommended actions
Appliance storage controller drive failure	<p>One or more drives in a StorageGRID appliance has failed or is not optimal.</p> <ol style="list-style-type: none"> 1. Use SANtricity System Manager to check hardware components, and follow the recommended actions. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance
Appliance storage controller hardware issue	<p>SANtricity software is reporting "Needs attention" for a component in a StorageGRID appliance.</p> <ol style="list-style-type: none"> 1. Use SANtricity System Manager to check hardware components, and follow the recommended actions. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance
Appliance storage controller power supply A failure	<p>Power supply A in a StorageGRID appliance has deviated from the recommended operating voltage.</p> <ol style="list-style-type: none"> 1. Use SANtricity System Manager to check hardware components, and follow the recommended actions. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance
Appliance storage controller power supply B failure	<p>Power supply B in a StorageGRID appliance has deviated from the recommended operating voltage.</p> <ol style="list-style-type: none"> 1. Use SANtricity System Manager to check hardware components, and follow the recommended actions. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance
Appliance storage hardware monitor service stalled	<p>The service that monitors storage hardware status has stopped reporting data.</p> <ol style="list-style-type: none"> 1. Check the status of the eos-system-status service in the base-os. 2. If the service is in a stopped or error state, restart the service. 3. If this alert persists, contact technical support.

Alert name	Description and recommended actions
Appliance storage shelves power supply degraded	<p>The status of one of the components in the storage shelf for a storage appliance is degraded.</p> <ol style="list-style-type: none"> 1. Use SANtricity System Manager to check hardware components, and follow the recommended actions. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance
Appliance temperature exceeded	<p>The nominal or maximum temperature for the appliance's storage controller has been exceeded.</p> <ol style="list-style-type: none"> 1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert. 2. Investigate possible reasons for the temperature increase, such as a fan or HVAC failure. 3. If this alert persists, contact technical support.
Appliance temperature sensor removed	<p>A temperature sensor has been removed. Contact technical support.</p>
Cassandra communication error	<p>The nodes that run the Cassandra service are having trouble communicating with each other.</p> <p>This alert indicates that something is interfering with node-to-node communications. There might be a network issue or the Cassandra service might be down on one or more Storage Nodes.</p> <ol style="list-style-type: none"> 1. Determine if there is another alert affecting one or more Storage Nodes. This alert might be resolved when you resolve the other alert. 2. Check for a network issue that might be affecting one or more Storage Nodes. 3. Select Support > Grid Topology. 4. For each Storage Node in your system, select SSM > Services. Ensure that the status of the Cassandra service is "Running." 5. If Cassandra is not running, follow the steps for starting or restarting a service in the recovery and maintenance instructions. 6. If all instances of the Cassandra service are now running and the alert is not resolved, contact technical support. <p>Recovery and maintenance</p>
Cassandra repair metrics out of date	<p>The metrics that describe Cassandra repair jobs are out of date. If this condition persists for more than 48 hours, client queries, such as bucket listings, might show deleted data.</p> <ol style="list-style-type: none"> 1. Reboot the node. From the Grid Manager, go to Nodes, select the node, and select the Tasks tab. 2. If this alert persists, contact technical support.

Alert name	Description and recommended actions
Cassandra repair progress slow	<p>The progress of Cassandra database repairs is slow.</p> <p>When database repairs are slow, Cassandra data consistency operations are impeded. If this condition persists for more than 48 hours, client queries, such as bucket listings, might show deleted data.</p> <ol style="list-style-type: none"> 1. Monitor this alert for up to 2 days to see if the issue resolves on its own. 2. If database repairs continue to proceed slowly, contact technical support.
Cassandra repair service not available	<p>The Cassandra repair service is not available.</p> <p>The Cassandra repair service exists on all Storage Nodes and provides critical repair functions for the Cassandra database. If this condition persists for more than 48 hours, client queries, such as bucket listings, might show deleted data.</p> <ol style="list-style-type: none"> 1. Select Support > Grid Topology. 2. For each Storage Node in your system, select SSM > Services. Ensure that the status of the Cassandra Reaper service is "Running." 3. If Cassandra Reaper is not running, follow the steps for starting or restarting a service in the recovery and maintenance instructions. 4. If all instances of the Cassandra Reaper service are now running and the alert is not resolved, contact technical support. <p><i>Recovery and maintenance</i></p>
Cloud Storage Pool connectivity error	<p>The health check for Cloud Storage Pools detected one or more new errors.</p> <ol style="list-style-type: none"> 1. Go to the Cloud Storage Pools section of the Storage Pools page. 2. Look at the Last Error column to determine which Cloud Storage Pool has an error. 3. See the instructions for administering StorageGRID. <p><i>Administering StorageGRID</i></p>
DHCP lease expired	<p>The DHCP lease on a network interface has expired.</p> <p>If the DHCP lease has expired, follow the recommended actions:</p> <ol style="list-style-type: none"> 1. Ensure there is connectivity between this node and the DHCP server on the affected interface. 2. Ensure there are IP addresses available to assign in the affected subnet on the DHCP server. 3. Ensure there is a permanent reservation for the IP address configured in the DHCP server. Or, use the StorageGRID Change IP tool to assign a static IP address outside of the DHCP address pool. See the recovery and maintenance instructions. <p><i>Recovery and maintenance</i></p>
DHCP lease expiring soon	<p>The DHCP lease on a network interface is expiring soon.</p> <p>To prevent the DHCP lease from expiring, follow the recommended actions:</p> <ol style="list-style-type: none"> 1. Ensure there is connectivity between this node and the DHCP server on the affected interface. 2. Ensure there are IP addresses available to assign in the affected subnet on the DHCP server. 3. Ensure there is a permanent reservation for the IP address configured in the DHCP server. Or, use the StorageGRID Change IP tool to assign a static IP address outside of the DHCP address pool. See the recovery and maintenance instructions. <p><i>Recovery and maintenance</i></p>

Alert name	Description and recommended actions
DHCP server unavailable	<p>The DHCP server is unavailable.</p> <p>The StorageGRID node is unable to contact your DHCP server. The DHCP lease for the node's IP address cannot be validated.</p> <ol style="list-style-type: none"> 1. Ensure there is connectivity between this node and the DHCP server on the affected interface. 2. Ensure there are IP addresses available to assign in the affected subnet on the DHCP server. 3. Ensure there is a permanent reservation for the IP address configured in the DHCP server. Or, use the StorageGRID Change IP tool to assign a static IP address outside of the DHCP address pool. See the recovery and maintenance instructions. <p><i>Recovery and maintenance</i></p>
Email notification failure	<p>The email notification for an alert could not be sent.</p> <p>This alert is triggered when an alert email notification fails or a test email (sent from the Alerts > Email Setup page) cannot be delivered.</p> <ol style="list-style-type: none"> 1. Sign in to Grid Manager from the Admin Node listed in the Site/Node column of the alert. 2. Go to the Alerts > Email Setup page, check the settings, and change them if required. 3. Click Send Test Email, and check the inbox of a test recipient for the email. A new instance of this alert might be triggered if the test email cannot be sent. 4. If the test email could not be sent, confirm your email server is online. 5. If the server is working, go to Support > Logs, and collect the log for the Admin Node. Specify a time period that is 15 minutes before and after the time of the alert. 6. Extract the downloaded archive, and review the contents of <code>prometheus.log (/GID<gid><time_stamp>/<site_node>/<time_stamp>/metrics/prometheus.log)</code>. 7. If you are unable to resolve the problem, contact technical support.
Expiration of load balancer endpoint certificate	<p>One or more load balancer endpoint certificates are about to expire.</p> <ol style="list-style-type: none"> 1. Go to Configuration > Load Balancer Endpoints. 2. Select an endpoint that has a certificate that will expire soon. 3. Select Edit endpoint to upload or generate a new certificate. 4. Repeat these steps for each endpoint that has an expired certificate or one that will expire soon. <p>For more information about managing load balancer endpoints, see the instructions for administering StorageGRID.</p> <p><i>Administering StorageGRID</i></p>
Expiration of server certificate for Management Interface	<p>The server certificate used for the management interface is about to expire.</p> <ol style="list-style-type: none"> 1. Go to Configuration > Server Certificates. 2. In the Management Interface Server Certificate section, upload a new certificate. <p><i>Administering StorageGRID</i></p>

Alert name	Description and recommended actions
Expiration of server certificate for Storage API Endpoints	<p>The server certificate used for accessing storage API endpoints is about to expire.</p> <ol style="list-style-type: none"> 1. Go to Configuration > Server Certificates. 2. In the Object Storage API Service Endpoints Server Certificate section, upload a new certificate. <p><i>Administering StorageGRID</i></p>
Grid Network MTU mismatch	<p>The maximum transmission unit (MTU) setting for the Grid Network interface (eth0) differs significantly across nodes in the grid.</p> <p>The differences in MTU settings could indicate that some, but not all, eth0 networks are configured for jumbo frames. An MTU size mismatch of greater than 1000 might cause network performance problems.</p> <p><i>Troubleshooting the Grid Network MTU mismatch alert</i> on page 195</p>
High Java heap use	<p>A high percentage of Java heap space is being used.</p> <p>If the Java heap becomes full, metadata services can become unavailable and client requests can fail.</p> <ol style="list-style-type: none"> 1. Review the ILM activity on the Dashboard. This alert might resolve on its own when the ILM workload decreases. 2. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert. 3. If this alert persists, contact technical support.
High latency for metadata queries	<p>The average time for Cassandra metadata queries is too long.</p> <p>An increase in query latency can be caused by a hardware change, such as replacing a disk, or a workload change, such as a sudden increase in ingests.</p> <ol style="list-style-type: none"> 1. Determine if there were any hardware or workload changes around the time the query latency increased. 2. If you are unable to resolve the problem, contact technical support.
Identity federation synchronization failure	<p>Unable to synchronize federated groups and users from the identity source.</p> <ol style="list-style-type: none"> 1. Confirm that the configured LDAP server is online and available. 2. Review the settings on the Identity Federation page. Confirm that all values are current. See "Configuring a federated identity source" in the instructions for administering StorageGRID. 3. Click Test Connection to validate the settings for the LDAP server. 4. If you cannot resolve the issue, contact technical support. <p><i>Administering StorageGRID</i></p>

Alert name	Description and recommended actions
ILM placement unachievable	<p>A placement instruction in an ILM rule cannot be achieved for certain objects. This alert indicates that a node required by a placement instruction is unavailable or that an ILM rule is misconfigured. For example, a rule might specify more replicated copies than there are Storage Nodes.</p> <ol style="list-style-type: none"> 1. Ensure that all nodes are online. 2. If all nodes are online, review the placement instructions in all ILM rules that are used the active ILM policy. Confirm that there are valid instructions for all objects. See the instructions for administering StorageGRID. 3. As required, update rule settings and activate a new policy. <p style="text-align: center;">Note: It might take up to 1 day for the alert to clear.</p> <ol style="list-style-type: none"> 4. If the problem persists, contact technical support. <p>Note: This alert might appear during an upgrade and could persist for 1 day after the upgrade is completed successfully. When this alert is triggered by an upgrade, it will clear on its own.</p> <p><i>Administering StorageGRID</i></p>
ILM scan period too long	<p>The time required to scan, evaluate objects, and apply ILM is too long. If the estimated time to complete a full ILM scan of all objects is too long (see Scan Period - Estimated on the Dashboard), the active ILM policy might not be applied to newly ingested objects. Changes to the ILM policy might not be applied to existing objects.</p> <ol style="list-style-type: none"> 1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert. 2. Confirm that all Storage Nodes are online. 3. Temporarily reduce the amount of client traffic. For example, from the Grid Manager, select Configuration > Traffic Classification Policies, and create a policy that limits bandwidth or the number of requests. 4. If disk I/O or CPU are overloaded, try to reduce the load or increase the resource. 5. If necessary, update ILM rules to use synchronous placement (default for rules created after StorageGRID 11.3). 6. If this alert persists, contact technical support. <p><i>Administering StorageGRID</i></p>
ILM scan rate low	<p>The ILM scan rate is set to less than 100 objects/second. This alert indicates that someone has changed the ILM scan rate for your system to less than 100 objects/second (default: 400 objects/second). The active ILM policy might not be applied to newly ingested objects. Subsequent changes to the ILM policy will not be applied to existing objects.</p> <ol style="list-style-type: none"> 1. Determine if a temporary change was made to the ILM scan rate as part of an ongoing support investigation. 2. Contact technical support. <p style="text-align: center;"> Attention: Never change the ILM scan rate without contacting technical support.</p>

Alert name	Description and recommended actions
Large audit queue	<p>The disk queue for audit messages is full.</p> <ol style="list-style-type: none"> 1. Check the load on the system—if there have been a significant number of transactions, the alert should resolve itself over time, and you can ignore the alert. 2. If the alert 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. 3. Reduce the client operation rate or decrease the number of audit messages logged by changing the audit level for Client Writes and Client Reads to Error or Off (Configuration > Audit). <p>Understanding audit messages</p>
Low audit log disk capacity	<p>The space available for audit logs is low.</p> <ol style="list-style-type: none"> 1. Monitor this alert to see if the issue resolves on its own and the disk space becomes available again. 2. Contact technical support if the available space continues to decrease.
Low available node memory	<p>The amount of RAM available on a node is low.</p> <p>Low available RAM could indicate a change in the workload or a memory leak with one or more nodes.</p> <ol style="list-style-type: none"> 1. Monitor this alert to see if the issue resolves on its own. 2. If the available memory falls below the major alert threshold, contact technical support.
Low installed node memory	<p>The amount of installed memory on a node is low.</p> <p>Increase the amount of RAM available to the virtual machine or Linux host. Check the threshold value for the major alert to determine the default minimum requirement for a StorageGRID node. See the installation instructions for your platform:</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux or CentOS installation • Ubuntu or Debian installation • VMware installation
Low metadata storage	<p>The space available for storing object metadata is low.</p> <p>Critical alert</p> <ol style="list-style-type: none"> 1. Stop ingesting objects. 2. Immediately add Storage Nodes in an expansion procedure. <p>Major alert</p> <p>Immediately add Storage Nodes in an expansion procedure.</p> <p>Minor alert</p> <ol style="list-style-type: none"> 1. Monitor the rate at which object metadata space is being used. Select Nodes > Storage Nodes > Storage, and view the Storage Used - Object Metadata graph. 2. Add Storage Nodes in an expansion procedure as soon as possible. <p>Once new Storage Nodes are added, the system automatically rebalances object metadata across all Storage Nodes, and the alarm clears.</p> <p>Troubleshooting the Low metadata storage alert on page 183</p> <p>Expanding a StorageGRID system</p>

Alert name	Description and recommended actions
Low metrics disk capacity	<p>The space available for the metrics database is low.</p> <ol style="list-style-type: none"> 1. Monitor this alert to see if the issue resolves on its own and the disk space becomes available again. 2. Contact technical support if the available space continues to decrease.
Low object data storage	<p>The space available for storing object data is low.</p> <p>Perform an expansion procedure. You can add storage volumes (LUNs) to existing Storage Nodes, or you can add new Storage Nodes.</p> <p>Troubleshooting the Low object data storage alert on page 174</p> <p>Expanding a StorageGRID system</p>
Low root disk capacity	<p>The space available for the root disk is low.</p> <ol style="list-style-type: none"> 1. Monitor this alert to see if the issue resolves on its own and the disk space becomes available again. 2. Contact technical support if the available space continues to decrease.
Low system data capacity	<p>The space available for StorageGRID system data on the <code>/var/local</code> file system is low.</p> <ol style="list-style-type: none"> 1. Monitor this alert to see if the issue resolves on its own and the disk space becomes available again. 2. Contact technical support if the available space continues to decrease.
Node network connectivity error	<p>Errors have occurred while transferring data between nodes.</p> <p>Network connectivity errors might clear without manual intervention. Contact technical support if the errors do not clear.</p> <p>Troubleshooting the Network Receive Error (NRER) alarm on page 196</p>
Node not in sync with NTP server	<p>The node's time is not in sync with the network time protocol (NTP) server.</p> <ol style="list-style-type: none"> 1. Verify that you have specified at least four external NTP servers, each providing a Stratum 3 or better reference. 2. Check that all NTP servers are operating normally. 3. Verify the connections to the NTP servers. Make sure they are not blocked by a firewall.
Node not locked with NTP server	<p>The node is not locked to a network time protocol (NTP) server.</p> <ol style="list-style-type: none"> 1. Verify that you have specified at least four external NTP servers, each providing a Stratum 3 or better reference. 2. Check that all NTP servers are operating normally. 3. Verify the connections to the NTP servers. Make sure they are not blocked by a firewall.
Non appliance node network down	<p>One or more network devices are down or disconnected.</p> <p>This alert indicates that a network interface (eth) for a node installed on a virtual machine or Linux host is not accessible.</p> <p>Contact technical support.</p>

Alert name	Description and recommended actions
Objects lost	<p>One or more objects have been lost from the grid. This alert might indicate that data has been permanently lost and is not retrievable.</p> <ol style="list-style-type: none"> 1. Investigate this alert immediately. You might need to take action to prevent further data loss. You also might be able to restore a lost object if you take prompt action. Troubleshooting lost and missing object data on page 165 2. When the underlying problem is resolved, reset the counter: <ol style="list-style-type: none"> a. Select Support. Then, in the Tools section of the menu, select Grid Topology. b. For the Storage Node that raised the alert, select <i>site</i> > <i>grid node</i> > LDR > Data Store > Configuration > Main. c. Select Reset Lost Objects Count and click Apply Changes.
Platform services unavailable	<p>Too few Storage Nodes with the RSM service are running or available at a site. Make sure that the majority of the Storage Nodes that have the RSM service at the affected site are running and in a non-error state. See "Troubleshooting platform services" in the instructions for administering StorageGRID. Administering StorageGRID</p>
S3 multipart part too small	<p>An S3 client is attempting to upload multipart parts that do not meet Amazon S3 size limits.</p> <ol style="list-style-type: none"> 1. Use the tenant ID shown in the alert details to identify the tenant account. 2. Use the IP address shown in the alert details to identify the specific S3 client. 3. Inform all client users that each part in a multipart upload must be between 5 MiB and 5 GiB. Only the last part can be less than 5 MiB (5,242,880 bytes). 4. Optionally, temporarily disable the minimum size restriction for StorageGRID: <ol style="list-style-type: none"> a. Log in to the primary Admin Node: <pre>ssh admin@grid_node_IP</pre> b. Enter the password from <code>Passwords.txt</code>. c. Switch to root: <pre>su-</pre> d. Run this script: <pre>'/usr/local/mgmt-api/app/scripts/s3-multipart-part-size-enforcement.rb --disable'</pre> <p>(Use <code>--enable</code> to undo this change.)</p> <p>Note: This script will be removed in StorageGRID 11.5. When you upgrade to the next release, all S3 clients must use part sizes between 5 MiB and 5 GiB.</p> <p>Troubleshooting the S3 multipart part too small alert on page 181</p>

Alert name	Description and recommended actions
Services appliance link down on Admin Network port 1	<p>The Admin Network port 1 on the appliance is down or disconnected.</p> <ol style="list-style-type: none"> 1. Check the cable and physical connection to Admin Network port 1. 2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware. 3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select Alert > Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> • SG100 and SG1000 appliance installation and maintenance • Disabling an alert rule on page 80
Services appliance link down on Grid Network (or Admin Network or Client Network)	<p>The appliance interface to the Grid Network (eth0), Admin Network (eth1), or the Client Network (eth2) is down or disconnected.</p> <ol style="list-style-type: none"> 1. Check the cables, SFPs, and physical connections to the StorageGRID network. 2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware. 3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select Alert > Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> • SG100 and SG1000 appliance installation and maintenance • Disabling an alert rule on page 80
Services appliance link down on network port 1, 2, 3, or 4	<p>Network port 1, 2, 3, or 4 on the appliance is down or disconnected.</p> <ol style="list-style-type: none"> 1. Check the cables, SFPs, and physical connections to the StorageGRID network. 2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware. 3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select Alert > Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> • SG100 and SG1000 appliance installation and maintenance • Disabling an alert rule on page 80
Storage appliance link down on Admin Network port 1	<p>The Admin Network port 1 on the appliance is down or disconnected.</p> <ol style="list-style-type: none"> 1. Check the cable and physical connection to Admin Network port 1. 2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware. 3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select Alert > Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance • Disabling an alert rule on page 80

Alert name	Description and recommended actions
Storage appliance link down on Grid Network (or Admin Network or Client Network)	<p>The appliance interface to the Grid Network (eth0), Admin Network (eth1), or Client Network (eth2) is down or disconnected.</p> <ol style="list-style-type: none"> 1. Check the cables, SFPs, and physical connections to the StorageGRID network. 2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware. 3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select Alert > Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance • Disabling an alert rule on page 80
Storage appliance link down on network port 1, 2, 3, or 4	<p>Network port 1, 2, 3, or 4 on the appliance is down or disconnected.</p> <ol style="list-style-type: none"> 1. Check the cables, SFPs, and physical connections to the StorageGRID network. 2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware. 3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select Alert > Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance • Disabling an alert rule on page 80
Storage device inaccessible	<p>A storage device cannot be accessed.</p> <p>This alert indicates that a volume cannot be mounted or accessed because of a problem with an underlying storage device.</p> <ol style="list-style-type: none"> 1. Check the status of all storage devices used for the node: <ul style="list-style-type: none"> • If the node is installed on a virtual machine or Linux host, follow the instructions for your operating system to run hardware diagnostics or perform a filesystem check. <ul style="list-style-type: none"> ◦ Red Hat Enterprise Linux or CentOS installation ◦ Ubuntu or Debian installation ◦ VMware installation • If the node is installed on an SG100, SG1000, or SG6000 appliance, use the BMC. • If the node is installed on a SG5600 or SG5700 appliance, use SANtricity System Manager. 2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware. <ul style="list-style-type: none"> • SG6000 appliance installation and maintenance • SG5700 appliance installation and maintenance • SG5600 appliance installation and maintenance

Alert name	Description and recommended actions
Unable to communicate with node	<p>One or more services are unresponsive, or the node cannot be reached.</p> <p>This alert indicates that a node is disconnected for an unknown reason. 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.</p> <p>Monitor this alert to see if the issue resolves on its own. If the issue persists:</p> <ol style="list-style-type: none"> 1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert. 2. Confirm that all of the services on this node are running. If a service is stopped, try starting it. See the recovery and maintenance instructions. 3. Ensure that the host for the node is powered on. If it is not, start the host. <p>Note: If more than one host is powered off, see the recovery and maintenance instructions.</p> <ol style="list-style-type: none"> 4. Determine if there is a network connectivity issue between this node and the Admin Node. 5. If you cannot resolve the alert, contact technical support. <p><i>Recovery and maintenance</i></p>
Unidentified corrupt object detected	<p>A file was found in replicated object storage that could not be identified as a replicated object.</p> <ol style="list-style-type: none"> 1. Determine if there are any issues with the underlying storage on a Storage Node. For example, run hardware diagnostics or perform a filesystem check. 2. After resolving any storage issues, run foreground verification to determine if objects are missing and to replace them if possible. 3. Monitor this alert. The alert will clear after 24 hours, but will be triggered again if the issue has not been fixed. 4. If you cannot resolve the alert, contact technical support. <p><i>Running foreground verification</i> on page 163</p>

Related reference

Commonly used Prometheus metrics on page 217

The Prometheus service on Admin Nodes collects time series metrics from the services on all nodes. While Prometheus collects more than a thousand metrics, a relatively small number are required to monitor the most critical StorageGRID operations.

Commonly used Prometheus metrics

The Prometheus service on Admin Nodes collects time series metrics from the services on all nodes. While Prometheus collects more than a thousand metrics, a relatively small number are required to monitor the most critical StorageGRID operations.

The following table lists the most commonly used Prometheus metrics and provides a mapping of each metric to the equivalent attribute (used in the alarm system).

You can refer to this list to better understand the conditions in the default alert rules or to construct the conditions for custom alert rules. For a complete list of metrics, select **Help > API Documentation**.

Note: Metrics that include `_private_` in their names are intended for internal use only and are subject to change between StorageGRID releases without notice.

Note: Prometheus metrics are retained for 31 days.

Prometheus metric	Attribute	Description
alertmanager_notifications_failed_total		The total number of failed alert notifications.
node_filesystem_avail_bytes		The amount of filesystem 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	BADD	For the storage controller in an appliance, the number of drives that are not optimal.
storagegrid_appliance_storage_controller_hardware_status		The overall status of the storage controller hardware in an appliance.
storagegrid_content_buckets_and_containers	SBKC	The total number of S3 buckets and Swift containers known by this Storage Node.
storagegrid_content_objects	SDOC	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	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. Troubleshooting lost and missing object data on page 165
storagegrid_http_sessions_incoming_attempted	HAIS	The total number of HTTP sessions that have been attempted to a Storage Node.
storagegrid_http_sessions_incoming_currently_established	HCCS	The number of HTTP sessions that are currently active (open) on the Storage Node.
storagegrid_http_sessions_incoming_failed	HEIS	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	HISC	The total number of HTTP sessions that have completed successfully.
storagegrid_ilm_awaiting_background_objects	BQUZ	The total number of objects on this node awaiting ILM evaluation from the scan.

Prometheus metric	Attribute	Description
storagegrid_ilm_awaiting_client_evaluation_objects_per_second	EVRT	The current rate at which objects are evaluated against the ILM policy on this node.
storagegrid_ilm_awaiting_client_objects	CQUZ	The total number of objects on this node awaiting ILM evaluation from client operations (for example, ingest).
storagegrid_ilm_awaiting_total_objects	QUSZ	The total number of objects awaiting ILM evaluation.
storagegrid_ilm_scan_objects_per_second	SCRT	The rate at which objects owned by this node are scanned and queued for ILM.
storagegrid_ilm_scan_period_estimated_minutes	SCTM	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_certificate_expiry_time		The expiration time of the load balancer endpoint certificate in seconds since the epoch.
storagegrid_metadata_queries_average_latency_milliseconds	CQST	The average time required to run a query against the metadata store through this service.
storagegrid_network_received_bytes	TRXB	The total amount of data received since installation.
storagegrid_network_transmitted_bytes	TTXB	The total amount of data sent since installation.
storagegrid_ntp_chosen_time_source_offset_milliseconds	NTSO	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	SRXB	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	STXB	The total amount of data retrieved by S3 clients from this Storage Node since the attribute was last reset.
storagegrid_s3_operations_failed	SFAL	The total number of failed S3 operations (HTTP status codes 4xx and 5xx), excluding those caused by S3 authorization failure.
storagegrid_s3_operations_successful	SSUC	The total number of successful S3 operations (HTTP status code 2xx).
storagegrid_s3_operations_unauthorized	SUAU	The total number of failed S3 operations that are the result of an authorization failure.
storagegrid_servercertificate_management_interface_certificate_expiry_days		The number of days before the Management Interface certificate expires.
storagegrid_servercertificate_storage_api_endpoints_certificate_expiry_days		The number of days before the Object Storage API certificate expires.
storagegrid_service_cpu_seconds	SUTM	The cumulative amount of time that the CPU has been used by this service since installation.

Prometheus metric	Attribute	Description
storagegrid_service_load	SLOD	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_service_memory_usage_bytes	SMEM	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	BREC	The total amount of data received by this service since installation.
storagegrid_service_network_transmitted_bytes	BTRA	The total amount of data sent by this service.
storagegrid_service_restarts	RSTS	The total number of times the service has been restarted.
storagegrid_service_runtime_seconds	SVRT	The total amount of time that the service has been running since installation.
storagegrid_service_uptime_seconds	SVUT	The total amount of time the service has been running since it was last restarted.
storagegrid_storage_state_current	SSCR	The current state of the storage services. Attribute values are: <ul style="list-style-type: none"> 10 = Offline 15 = Maintenance 20 = Read-only 30 = Online
storagegrid_storage_status	SSTS	The current status of the storage services. Attribute values are: <ul style="list-style-type: none"> 0 = No Errors 10 = In Transition 20 = Insufficient Free Space 30 = Volume(s) Unavailable 40 = Error
storagegrid_storage_utilization_data_bytes	SPSD	An estimate of the total size of replicated and erasure coded object data on the Storage Node.
storagegrid_storage_utilization_metadata_allowed_bytes	CEMS	The total space available on storage volume 0 for object metadata. Metadata Allowed Space (CEMS) is always less than the Metadata Reserved Space (CAWM) because a portion of the reserved metadata space is required for essential database operations, such as compaction and repair.
storagegrid_storage_utilization_metadata_bytes	CADL	The amount of object metadata on storage volume 0, in bytes.
storagegrid_storage_utilization_total_space_bytes	STTS	The total amount of storage space allocated to all object stores.
storagegrid_storage_utilization_usable_space_bytes	STAS	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.

Prometheus metric	Attribute	Description
storagegrid_swift_data_transfers_bytes_ingested	WRXB	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	WTXB	The total amount of data retrieved by Swift clients from this Storage Node since the attribute was last reset.
storagegrid_swift_operations_failed	WFAL	The total number of failed Swift operations (HTTP status codes 4xx and 5xx), excluding those caused by Swift authorization failure.
storagegrid_swift_operations_successful	WSUC	The total number of successful Swift operations (HTTP status code 2xx).
storagegrid_swift_operations_unauthorized	WUAU	The total number of failed Swift operations that are the result of an authorization failure (HTTP status codes 401, 403, 405).

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.

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new 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 cannot report attribute values to the NMS service. Thus, the NMS service can no longer monitor the status of the service, or update attributes for the service. If the problem persists, contact technical support.
ACMS	Available Metadata Services	BARC, BLDR, BCMN	An alarm is triggered when an LDR or ARC service loses connection to a DDS service. If this occurs, ingest or retrieve transactions cannot be processed. If the unavailability of DDS services is only a brief transient issue, transactions can be delayed. Check and restore connections to a DDS service to clear this alarm and return the service to full functionality.
ACTS	Cloud Tiering Service Status	ARC	Only available for Archive Nodes with a Target Type of Cloud Tiering - Simple Storage Service (S3). If the ACTS attribute for the Archive Node is set to Read-Only Enabled or Read-Write Disabled, you must set the attribute to Read-Write Enabled. If a major alarm is triggered due to an authentication failure, verify the credentials associated with destination bucket and update values, if necessary. If a major alarm is triggered due to any other reason, contact technical support.
ADCA	ADC Status	ADC	If an alarm is triggered, select Support . Then, in the Tools section of the menu, select Grid Topology . Then select <i>site</i> > <i>grid node</i> > ADC > Overview > Main and ADC > Alarms > Main to determine the cause of the alarm. If the problem persists, contact technical support.
ADCE	ADC State	ADC	If the value of ADC State is Standby, continue monitoring the service and if the problem persists, contact technical support. If the value of ADC State is Offline, restart the service. If the problem persists, contact technical support.

Code	Name	Service	Recommended action
AITE	Retrieve State	BARC	<p>Only available for Archive Node's with a Target Type of Tivoli Storage Manager (TSM).</p> <p>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.</p> <p>If the value of Archive Retrieve State is Offline, attempt to update the state to Online. Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site</i> > <i>grid node</i> > ARC > Retrieve > Configuration > Main, select Archive Retrieve State > Online, and click Apply Changes.</p> <p>If the problem persists, contact technical support.</p>
AITU	Retrieve Status	BARC	<p>If the value of Retrieve Status is Target Error, check the targeted external archival storage system for errors.</p> <p>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.</p> <p>If the value of Archive Retrieve Status is Unknown Error, contact technical support.</p>
ALIS	Inbound Attribute Sessions	ADC	<p>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.</p> <p>If the problem persists, contact technical support.</p>
ALOS	Outbound Attribute Sessions	ADC	<p>The ADC service has a high number of attribute sessions, and is becoming overloaded. If this alarm is triggered, contact technical support.</p>
ALUR	Unreachable Attribute Repositories	ADC	<p>Check network connectivity with the NMS service to ensure that the service can contact the attribute repository.</p> <p>If this alarm is triggered and network connectivity is good, contact technical support.</p>

Code	Name	Service	Recommended action
AMQS	Audit Messages Queued	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BDDS	<p>If audit messages cannot 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.</p> <p>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:</p> <ul style="list-style-type: none"> • Notice: More than 100,000 messages • Minor: At least 500,000 messages • Major: At least 2,000,000 messages • Critical: At least 5,000,000 messages <p>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.</p> <p>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 "Changing audit message levels" in <i>Understanding audit messages</i>.</p> <p>Understanding audit messages</p>
AOTE	Store State	BARC	<p>Only available for Archive Node's with a Target Type of Tivoli Storage Manager (TSM).</p> <p>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.</p> <p>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.</p>
AOTU	Store Status	BARC	<p>If the value of Store Status is Session Lost check that the external archival storage system is connected and online.</p> <p>If the value of Target Error, check the external archival storage system for errors.</p> <p>If the value of Store Status is Unknown Error, contact technical support.</p>

Code	Name	Service	Recommended action
APMS	Storage Multipath Connectivity	SSM	<p>If the multipath state alarm appears as "Degraded" (select Support. Then, in the Tools section of the menu, select Grid Topology, then select <i>site > grid node > SSM > Events</i>), do the following:</p> <ol style="list-style-type: none"> 1. Plug in or replace the cable that does not display any indicator lights. 2. Wait one to five minutes. Do not unplug the other cable until at least five minutes after you plug in the first one. Unplugging too early can cause the root volume to become read-only, which requires that the hardware be restarted. 3. Return to the SSM > Resources page, and verify that the "Degraded" Multipath status has changed to "Nominal" in the Storage Hardware section.
ARCE	ARC State	ARC	<p>The ARC service has a state of Standby until all ARC components (Replication, Store, Retrieve, Target) have started. It then transitions to Online.</p> <p>If the value of ARC State does not transition from Standby to Online, check the status of the ARC components.</p> <p>If the value of ARC State is Offline, restart the service.</p> <p>If the problem persists, contact technical support.</p>
AROQ	Objects Queued	ARC	<p>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.</p> <p>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.</p>

Code	Name	Service	Recommended action
ARRF	Request Failures	ARC	<p>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.</p> <p>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.</p> <p>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.</p> <p>Once the problem that triggered this alarm is addressed, reset the failures count. Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site > grid node > ARC > Retrieve > Configuration > Main</i>, select Reset Request Failure Count and click Apply Changes.</p>
ARRV	Verification Failures	ARC	<p>To diagnose and correct this problem, contact technical support.</p> <p>Once the problem that triggered this alarm is addressed, reset the failures count. Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site > grid node > ARC > Retrieve > Configuration > Main</i>, select Reset Verification Failure Count and click Apply Changes.</p>
ARVF	Store Failures	ARC	<p>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.</p> <p>Once the problem that triggered this alarm is addressed, reset the failures count. Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site > grid node > ARC > Retrieve > Configuration > Main</i>, select Reset Store Failure Count, and click Apply Changes.</p>
ASXP	Audit Shares	AMS	<p>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.</p> <p>If the problem persists, contact technical support.</p>
AUMA	AMS Status	AMS	<p>If the value of AMS Status is DB Connectivity Error, restart the grid node.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
AUME	AMS State	AMS	<p>If the value of AMS State is Standby, continue monitoring the StorageGRID system. If the problem persists, contact technical support.</p> <p>If the value of AMS State is Offline, restart the service. If the problem persists, contact technical support.</p>
AUXS	Audit Export Status	AMS	<p>If an alarm is triggered, correct the underlying problem, and then restart the AMS service.</p> <p>If the problem persists, contact technical support.</p>
BADD	Storage Controller Failed Drive Count	SSM	<p>This alarm is triggered when one or more drives in a StorageGRID appliance has failed or is not optimal.</p> <p>Replace the drives as required.</p>
BASF	Available Object Identifiers	CMN	<p>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.</p> <p>To allocate more identifiers, contact technical support.</p>
BASS	Identifier Block Allocation Status	CMN	<p>By default, an alarm is triggered when object identifiers cannot be allocated because ADC quorum cannot be reached.</p> <p>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 re-established. 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.</p> <p>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.</p> <p>If the problem persists, contact technical support.</p>
BRDT	Compute Controller Chassis Temperature	SSM	<p>An alarm is triggered if the temperature of the compute controller in a StorageGRID appliance exceeds a nominal threshold.</p> <p>Check hardware components and environmental issues for overheated condition. If necessary, replace the component.</p>

Code	Name	Service	Recommended action
BTOF	Offset	BADC, BLDR, BNMS, BAMS, BCLB, BCMN, BARC	An alarm is triggered if the service time (seconds) differs significantly from the operating system time. Under normal conditions, the service should resynchronize itself. If the service time drifts too far from the operating system time, system operations can be affected. Confirm that the StorageGRID system's time source is correct. If the problem persists, contact technical support.
BTSE	Clock State	BADC, BLDR, BNMS, BAMS, BCLB, BCMN, BARC	An alarm is triggered if the service's time is not synchronized with the time tracked by the operating system. Under normal conditions, the service should resynchronize itself. If the time drifts too far from operating system time, system operations can be affected. Confirm that the StorageGRID system's time source is correct. If the problem persists, contact technical support.
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 . Then, in the Tools section of the menu, select Grid Topology , then select <i>site > grid node > DDS > Resources > Overview > Main</i> . If the problem persists, contact technical support.
CAIH	Number Available Ingest Destinations	CLB	This alarm is deprecated.
CAQH	Number Available Destinations	CLB	This alarm clears when underlying issues of available LDR services are corrected. Ensure that the HTTP component of LDR services are online and running normally. If the problem persists, contact technical support.

Code	Name	Service	Recommended action
CASA	Data Store Status	DDS	<p>An alarm is raised if the Cassandra metadata store becomes unavailable.</p> <p>Check the status of Cassandra:</p> <ol style="list-style-type: none"> 1. At the Storage Node, log in as admin and <code>su</code> to root using the password listed in the <code>Passwords.txt</code> file. 2. Enter: <code>service cassandra status</code> 3. If Cassandra is not running, restart it: <code>service cassandra restart</code> <p>This alarm might also indicate that the metadata store (Cassandra database) for a Storage Node requires rebuilding.</p> <p><i>Troubleshooting the Services: Status - Cassandra (SVST) alarm</i> on page 185</p> <p>If the problem persists, contact technical support.</p>
CASE	Data Store State	DDS	<p>This alarm is triggered during installation or expansion to indicate a new data store is joining the grid.</p>
CCES	Incoming Sessions - Established	CLB	<p>This alarm is triggered if there are 20,000 or more HTTP sessions currently active (open) on the Gateway Node. If a client has too many connections, you might see connection failures. You should reduce the workload.</p>
CCNA	Compute Hardware	SSM	<p>This alarm is triggered if the status of the compute controller hardware in a StorageGRID appliance is Needs Attention.</p>

Code	Name	Service	Recommended action
CDLP	Metadata Used Space (Percent)	DDS	<p>This alarm is triggered when the Metadata Effective Space (CEMS) reaches 70% full (minor alarm), 90% full (major alarm), and 100% full (critical alarm).</p> <p>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 the instructions for expanding a StorageGRID grid.</p> <p>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.</p> <p>Note: Contact technical support if you are unable to add Storage Nodes.</p> <p>Once new Storage Nodes are added, the system automatically rebalances object metadata across all Storage Nodes, and the alarm clears.</p> <p>Troubleshooting the Low metadata storage alert on page 183</p> <p>Expanding a StorageGRID system</p>
CLBA	CLB Status	CLB	<p>If an alarm is triggered, select Support. Then, in the Tools section of the menu, select Grid Topology, then select <i>site</i> > <i>grid node</i> > CLB > Overview > Main and CLB > Alarms > Main to determine the cause of the alarm and to troubleshoot the problem.</p> <p>If the problem persists, contact technical support.</p>
CLBE	CLB State	CLB	<p>If the value of CLB State is Standby, continue monitoring the situation and if the problem persists, contact technical support.</p> <p>If the state is Offline and there are no known server hardware issues (for example, the server is unplugged) or scheduled downtime, restart the service. If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
CMNA	CMN Status	CMN	<p>If the value of CMN Status is Error, select Support. Then, in the Tools section of the menu, select Grid Topology, then select <i>site > grid node > CMN > Overview > Main</i> and CMN > Alarms > Main to determine the cause of the error and to troubleshoot the problem.</p> <p>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).</p> <p>If the problem persists, contact technical support.</p>
CPRC	Remaining Capacity	NMS	<p>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.</p> <p>If an alarm is triggered, contact technical support.</p>
CPSA	Compute Controller Power Supply A	SSM	<p>An alarm is triggered if there is an issue with power supply A in the compute controller for a StorageGRID appliance.</p> <p>If necessary, replace the component.</p>
CPSB	Compute Controller Power Supply B	SSM	<p>An alarm is triggered if there is an issue with power supply B in the compute controller for a StorageGRID appliance.</p> <p>If necessary, replace the component.</p>
CPUT	Compute Controller CPU Temperature	SSM	<p>An alarm is triggered if the temperature of the CPU in the compute controller in a StorageGRID appliance exceeds a nominal threshold.</p> <p>If the Storage Node is a StorageGRID appliance, the StorageGRID system indicates that the controller needs attention.</p> <p>Check hardware components and environment issues for overheated condition. If necessary, replace the component.</p>
DNST	DNS Status	SSM	<p>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.</p>

Code	Name	Service	Recommended action
ECCD	Corrupt Fragments Detected	LDR	<p>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.</p> <p>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 may 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.</p> <p>If the problem persists, contact technical support.</p>
ECST	Verification Status	LDR	<p>This alarm indicates the current status of the background verification process for erasure coded object data on this Storage Node.</p> <p>A major alarm is triggered if there is an error in the background verification process.</p>
FOPN	Open File Descriptors	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BSSM, BDDS	<p>FOPN can become large during peak activity. If it does not diminish during periods of slow activity, contact technical support.</p>
HCCS	Currently Established Incoming Sessions	LDR	<p>This alarm is triggered if there are 10,000 or more HTTP sessions currently active (open) on the Storage Node. If a node has too many connections, you might see connection failures. You should reduce the workload.</p>
HSTE	HTTP State	BLDR	<p>HSTE and HSTU are related to the HTTP protocol for all LDR traffic, including S3, Swift, and other internal StorageGRID traffic. An alarm indicates that one of the following situations has occurred:</p> <ul style="list-style-type: none"> • The HTTP protocol has been taken offline manually. • The Auto-Start HTTP attribute has been disabled. • The LDR service is shutting down. <p>The Auto-Start HTTP attribute is enabled by default. If this setting is changed, HTTP could remain offline after a restart.</p> <p>If necessary, wait for the LDR service to restart.</p> <p>Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>Storage Node</i> > LDR > Configuration. If the HTTP protocol is offline, place it online. Verify that the Auto-Start HTTP attribute is enabled.</p> <p>If the HTTP protocol remains offline, contact technical support.</p>
HSTU	HTTP Status	BLDR	

Code	Name	Service	Recommended action
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 . Then, in the Tools section of the menu, select Grid Topology . Then select <i>site</i> > <i>grid node</i> > LDR > Replication > Configuration > Main .
LATA	Average Latency	NMS	Check for connectivity issues. Check system activity to confirm that there is an increase in system activity. An increase in system activity will result in an increase to attribute data activity. This increased activity will result in a delay to the processing of attribute data. This can be normal system activity and will subside. Check for multiple alarms. An increase in average latency times can be indicated by an excessive number of triggered alarms. If the problem persists, contact technical support.
LDRE	LDR State	LDR	If the value of LDR State is Standby, continue monitoring the situation and if the problem persists, contact technical support. If the value of LDR State is Offline, restart the service. If the problem persists, contact technical support.
LOST	Lost Objects	DDS, LDR	Triggered when the StorageGRID system fails to retrieve a copy of the requested object from anywhere in the system. Before a LOST (Lost Objects) alarm is triggered, the system attempts to retrieve and replace a missing object from elsewhere in the system. Lost objects represent a loss of data. The Lost Objects attribute is incremented whenever the number of locations for an object drops to zero without the DDS service purposely purging the content to satisfy the ILM policy. Investigate LOST (LOST Object) alarms immediately. If the problem persists, contact technical support. Troubleshooting lost and missing object data on page 165
MCEP	Management Interface Certificate Expiry	CMN	Triggered when the certificate used for accessing the management interface is about to expire. <ol style="list-style-type: none"> Go to Configuration > Server Certificates. In the Management Interface Server Certificate section, upload a new certificate. Administering StorageGRID

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. Configuring email server settings for alarms (legacy system) on page 103
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. Configuring email server settings for alarms (legacy system) on page 103
MISS	NMS Interface Engine Status	BNMS	An alarm is triggered if the NMS interface engine on the Admin Node that gathers and generates interface content is disconnected from the system. Check Server Manager to determine if the server individual application is down.
NANG	Network Auto Negotiate Setting	SSM	Check the network adapter configuration. The setting must match preferences of your network routers and switches. An incorrect setting can have a severe impact on system performance.
NDUP	Network Duplex Setting	SSM	Check the network adapter configuration. The setting must match preferences of your network routers and switches. An incorrect setting can have a severe impact on system performance.
NLNK	Network Link Detect	SSM	Check the network cable connections on the port and at the switch. Check the network router, switch, and adapter configurations. Restart the server. If the problem persists, contact technical support.
NRER	Receive Errors	SSM	The following can be causes of NRER alarms: <ul style="list-style-type: none"> • Forward error correction (FEC) mismatch • Switch port and NIC MTU mismatch • High link error rates • NIC ring buffer overrun Troubleshooting the Network Receive Error (NRER) alarm on page 196

Code	Name	Service	Recommended action
NRLY	Available Audit Relays	BADC, BARC, BCLB, BCMN, BLDR, BNMS, BDDS	<p>If audit relays are not connected to ADC services, audit events cannot be reported. They are queued and unavailable to users until the connection is restored.</p> <p>Restore connectivity to an ADC service as soon as possible.</p> <p>If the problem persists, contact technical support.</p>
NSCA	NMS Status	NMS	<p>If the value of NMS Status is DB Connectivity Error, restart the service. If the problem persists, contact technical support.</p>
NSCE	NMS State	NMS	<p>If the value of NMS State is Standby, continue monitoring and if the problem persists, contact technical support.</p> <p>If the value of NMS State is Offline, restart the service. If the problem persists, contact technical support.</p>
NSPD	Speed	SSM	<p>This can be caused by network connectivity or driver compatibility issues. If the problem persists, contact technical support.</p>
NTBR	Free Tablespace	NMS	<p>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.</p> <p>Adjusting the alarm threshold allows you to proactively manage when additional storage needs to be allocated.</p> <p>If the available space reaches a low threshold (see alarm threshold), contact technical support to change the database allocation.</p>
NTER	Transmit Errors	SSM	<p>These errors can clear without being manually reset. If they do not clear, check network hardware. Check that the adapter hardware and driver are correctly installed and configured to work with your network routers and switches.</p> <p>When the underlying problem is resolved, reset the counter. Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site</i> > <i>grid node</i> > SSM > Resources > Configuration > Main, select Reset Transmit Error Count, and click Apply Changes.</p>
NTFQ	NTP Frequency Offset	SSM	<p>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.</p>
NTLK	NTP Lock	SSM	<p>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.</p>

Code	Name	Service	Recommended action
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	These values give an indication of the reliability and stability of the time source that NTP on the local server is using as its reference. If an alarm is triggered, it can be an indication that the time source's oscillator is defective, or that there is a problem with the WAN link to the time source.
NTSO	Chosen Time Source Offset		
NTSU	NTP Status	SSM	If the value of NTP Status is Not Running, contact technical support.
OPST	Overall Power Status	SSM	An alarm is triggered if the power of a StorageGRID appliance deviates from the recommended operating voltage. Check the status of Power Supply A or B to determine which power supply is operating abnormally. If necessary, replace the power supply.
OQRT	Objects Quarantined	LDR	After the objects are automatically restored by the StorageGRID system, the quarantined objects can be removed from the quarantine directory. <ol style="list-style-type: none"> 1. Select Support. Then, in the Tools section of the menu, select Grid Topology. 2. Select <i>site</i> > <i>Storage Node</i> > LDR > Verification > Configuration > Main. 3. Select Delete Quarantined Objects. 4. Click Apply Changes. <p>The quarantined objects are removed, and the count is reset to zero.</p>
ORSU	Outbound Replication Status	BLDR, BARC	An alarm indicates that outbound replication is not possible: storage is in a state where objects cannot be retrieved. An alarm is triggered if outbound replication is disabled manually. Select Support . Then, in the Tools section of the menu, select Grid Topology . Then select <i>site</i> > <i>grid node</i> > LDR > Replication > Configuration . An alarm is triggered if the LDR service is unavailable for replication. Select Support . Then, in the Tools section of the menu, select Grid Topology . Then select <i>site</i> > <i>grid node</i> > 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	<p>Can have a value of Over Y% RAM, where Y represents the percentage of memory being used by the server.</p> <p>Figures under 80% are normal. Over 90% is considered a problem.</p> <p>If memory usage is high for a single service, monitor the situation and investigate.</p> <p>If the problem persists, contact technical support.</p>
PSAS	Power Supply A Status	SSM	<p>An alarm is triggered if power supply A in a StorageGRID appliance deviates from the recommended operating voltage.</p> <p>If necessary, replace power supply A.</p>
PSBS	Power Supply B Status	SSM	<p>An alarm is triggered if power supply B in a StorageGRID appliance deviates from the recommended operating voltage.</p> <p>If necessary, replace the power supply B.</p>
RDTE	Tivoli Storage Manager State	BARC	<p>Only available for Archive Nodes with a Target Type of Tivoli Storage Manager (TSM).</p> <p>If the value of Tivoli Storage Manager State is Offline, check Tivoli Storage Manager Status and resolve any problems.</p> <p>Bring the component back online. Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site</i> > <i>grid node</i> > ARC > Target > Configuration > Main, select Tivoli Storage Manager State > Online, and click Apply Changes.</p>

Code	Name	Service	Recommended action
RDTU	Tivoli Storage Manager Status	BARC	<p>Only available for Archive Nodes with a Target Type of Tivoli Storage Manager (TSM).</p> <p>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.</p> <p>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.</p> <p>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 cannot authenticate the connection. Check that the TSM middleware server is configured with the correct user, password, and permissions, and restart the service.</p> <p>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.</p> <p>If the value of Tivoli Storage Manager Status is Unknown Error, contact technical support.</p>
RIRF	Inbound Replications – Failed	BLDR, BARC	<p>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.</p> <p>To reset the count, select Support. Then, in the Tools section of the menu, select Grid Topology, then select <i>site > grid node > LDR > Replication > Configuration > Main</i>. Select Reset Inbound Replication Failure Count, and click Apply Changes.</p>
RIRQ	Inbound Replications – Queued	BLDR, BARC	<p>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.</p>

Code	Name	Service	Recommended action
RORQ	Outbound Replications – Queued	BLDR, BARC	<p>The outbound replication queue contains object data being copied to satisfy ILM rules and objects requested by clients.</p> <p>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.</p>
SAVP	Total Usable Space (Percent)	LDR	<p>If usable space reaches a low threshold, options include expanding the StorageGRID system or move object data to archive through an Archive Node.</p> <p>Troubleshooting the Total Usable Space (Percent) (SAVP) alarm on page 176</p>
SCAS	Status	CMN	<p>If the value of Status for the active grid task is Error, look up the grid task message. Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site > grid node > CMN > Grid Tasks > Overview > Main</i>. The grid task message displays information about the error (for example, “check failed on node 12130011”).</p> <p>After you have investigated and corrected the problem, restart the grid task. Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site > grid node > CMN > Grid Tasks > Configuration > Main</i>, and select Actions > Run.</p> <p>If the value of Status for a grid task being aborted is Error, retry aborting the grid task.</p> <p>If the problem persists, contact technical support.</p>
SCEP	Storage API Service Endpoints Certificate Expiry	CMN	<p>Triggered when the certificate used for accessing storage API endpoints is about to expire.</p> <ol style="list-style-type: none"> Go to Configuration > Server Certificates. In the Object Storage API Service Endpoints Server Certificate section, upload a new certificate. <p>Administering StorageGRID</p>
SCHR	Status	CMN	<p>If the value of Status for the historical grid task is Aborted, investigate the reason and run the task again if required.</p> <p>If the problem persists, contact technical support.</p>
SCSA	Storage Controller A	SSM	<p>An alarm is triggered if there is an issue with storage controller A in a StorageGRID appliance.</p> <p>If necessary, replace the component.</p>
SCSB	Storage Controller B	SSM	<p>An alarm is triggered if there is an issue with storage controller B in a StorageGRID appliance.</p> <p>If necessary, replace the component.</p> <p>Some appliance models do not have a storage controller B.</p>

Code	Name	Service	Recommended action
SHLH	Health	LDR	<p>If the value of Health for an object store is Error, check and correct:</p> <ul style="list-style-type: none"> • problems with the volume being mounted • file system errors
SLSA	CPU Load Average	SSM	<p>The higher the value the busier the system.</p> <p>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. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site > grid node > SSM > Resources > Reports > Charts</i>.</p> <p>If the load on the system is not heavy and the problem persists, contact technical support.</p>
SMST	Log Monitor State	SSM	<p>If the value of Log Monitor State is not Connected for a persistent period of time, contact technical support.</p>
SMTT	Total Events	SSM	<p>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.</p> <p>When an issue is resolved, reset the counter to clear the alarm. Select Nodes > site > grid node > Events > Reset event counts.</p> <p>Note: To reset event counts, you must have the Grid Topology Page Configuration permission.</p> <p>If the value of Total Events is zero, or the number increases and the problem persists, contact technical support.</p>
SNST	Status	CMN	<p>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.</p> <p>If the problem persists, contact technical support.</p>
SOSS	Storage Operating System Status	SSM	<p>An alarm is triggered if SANtricity software indicates that there is a "Needs attention" issue with a component in a StorageGRID appliance.</p> <p>Select Nodes. Then select <i>appliance Storage Node > Hardware</i>. Scroll down to view the status of each component. In SANtricity software, check other appliance components to isolate the issue.</p>

Code	Name	Service	Recommended action
SSMA	SSM Status	SSM	<p>If the value of SSM Status is Error, select Support. Then, in the Tools section of the menu, select Grid Topology, then select <i>site</i> > <i>grid node</i> > SSM > Overview > Main and SSM > Overview > Alarms to determine the cause of the alarm.</p> <p>If the problem persists, contact technical support.</p>
SSME	SSM State	SSM	<p>If the value of SSM State is Standby, continue monitoring, and if the problem persists, contact technical support.</p> <p>If the value of SSM State is Offline, restart the service.</p> <p>If the problem persists, contact technical support.</p>
SSTS	Storage Status	BLDR	<p>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.</p> <p>Additional storage should be added. It is not impacting end user functionality, but the alarm persists until additional storage is added.</p> <p>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. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site</i> > <i>grid node</i> > LDR > Storage > Overview > Main. The volume's Health is listed under Object Stores.</p> <p>If the value of Storage Status is Error, contact technical support.</p> <p>Troubleshooting the Storage Status (SSTS) alarm on page 177</p>

Code	Name	Service	Recommended action
SVST	Status	SSM	<p>This alarm clears when other alarms related to a non-running service are resolved. Track the source service alarms to restore operation.</p> <p>Select Support. Then, in the Tools section of the menu, select Grid Topology. Then select <i>site</i> > <i>grid node</i> > 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:</p> <ul style="list-style-type: none"> • The service has been manually stopped (<code>/etc/init.d/<service> stop</code>). • There is an issue with the MySQL database and Server Manager shuts down the MI service. • A grid node has been added, but not started. • During installation, a grid node has not yet connected to the Admin Node. <p>If a service is listed as Not Running, restart the service (<code>/etc/init.d/<service> restart</code>).</p> <p>This alarm might also indicate that the metadata store (Cassandra database) for a Storage Node requires rebuilding.</p> <p>If the problem persists, contact technical support. Troubleshooting the Services: Status - Cassandra (SVST) alarm on page 185</p>
TMEM	Installed Memory	SSM	<p>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.</p>
TPOP	Pending Operations	ADC	<p>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.</p>
UMEM	Available Memory	SSM	<p>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.</p>
VMFI	Entries Available	SSM	<p>This is an indication that additional storage is required. Contact technical support.</p>
VMFR	Space Available	SSM	<p>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.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
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 cannot 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 . Then, in the Tools section of the menu, select Grid Topology . Then select <i>site > grid node > LDR > Storage > Overview > Main</i> . Check the operating system for any signs of block-device or file system errors. If the value of Object Verification Status is Unknown Error, it usually indicates a low-level file system or hardware problem (I/O error) that prevents the Storage Verification task from accessing stored content. Contact technical support.
XAMS	Unreachable Audit Repositories	BADC, BARC, BCLB, BCMN, BLDR, BNMS	Check network connectivity to the server hosting the Admin Node. If the problem persists, contact technical support.

Alarms that generate SNMP notifications (legacy system)

The following table lists the legacy alarms that generate SNMP notifications. Unlike alerts, not all alarms generate SNMP notifications. Only the alarms listed generate SNMP notifications and only at the indicated severity or higher.

Note: While the alarm system continues to be supported in StorageGRID 11.4, the new alert system offers significant benefits and is easier to use.

Code	Name	Severity
ACMS	Available Metadata Services	Critical
AITE	Retrieve State	Minor
AITU	Retrieve Status	Major
AMQS	Audit Messages Queued	Notice
AOTE	Store State	Minor
AOTU	Store Status	Major
AROQ	Objects Queued	Minor
ARRF	Request Failures	Major
ARRV	Verification Failures	Major
ARVF	Store Failures	Major

Code	Name	Severity
ASXP	Audit Shares	Minor
AUMA	AMS Status	Minor
AUXS	Audit Export Status	Minor
BTOF	Offset	Notice
CAHP	Java Heap Usage Percent	Major
CAQH	Number Available Destinations	Notice
CASA	Data Store Status	Major
CDLP	Metadata Used Space (Percent)	Major
CLBE	CLB State	Critical
DNST	DNS Status	Critical
ECST	Verification Status	Major
HSTE	HTTP State	Major
HTAS	Auto-Start HTTP	Notice
LOST	Lost Objects	Major
MINQ	E-mail Notifications Queued	Notice
MINS	E-mail Notifications Status	Minor
NANG	Network Auto Negotiate Setting	Notice
NDUP	Network Duplex Setting	Minor
NLNK	Network Link Detect	Minor
NRER	Receive Errors	Notice
NSPD	Speed	Notice
NTER	Transmit Errors	Notice
NTFQ	NTP Frequency Offset	Minor
NTLK	NTP Lock	Minor
NTOF	NTP Time Offset	Minor
NTSJ	Chosen Time Source Jitter	Minor
NTSO	Chosen Time Source Offset	Minor
NTSU	NTP Status	Major
OPST	Overall Power Status	Major
ORSU	Outbound Replication Status	Notice
PSAS	Power Supply A Status	Major
PSBS	Power Supply B Status	Major
RDTE	Tivoli Storage Manager State	Notice
RDTU	Tivoli Storage Manager Status	Major

Code	Name	Severity
SAVP	Total Usable Space (Percent)	Notice
SHLH	Health	Notice
SLSA	CPU Load Average	Notice
SMTT	Total Events	Notice
SNST	Status	
SOSS	Storage Operating System Status	Notice
SSTS	Storage Status	Notice
SVST	Status	Notice
TMEM	Installed Memory	Minor
UMEM	Available Memory	Minor
VMST	Status	Minor
VPRI	Verification Priority	Notice
VSTU	Object Verification Status	Notice

Log files reference

The following sections list the logs 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.



Attention: These tables 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 this guide.

To access these logs, you can collect log files and system data (**Support > Logs**). Or, if the primary Admin Node is unavailable or unable to reach a specific node, you can access the logs 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.

Related tasks

[Collecting log files and system data](#) on page 139

You can use the Grid Manager to retrieve log files and system data (including configuration data) for your StorageGRID system.

StorageGRID software logs

You can use StorageGRID logs to troubleshoot issues.

General StorageGRID logs

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

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.	
/var/local/log/gridstat.errlog	Log file for the Server Manager GUI.	

Logs for StorageGRID services

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/clb.errlog	Error information for the CLB service. Note: The CLB service is deprecated.	Gateway Nodes
/var/local/log/cmn.errlog		Admin Nodes
/var/local/log/cms.errlog	This log file might be present on systems that have been upgraded from an older version of StorageGRID. It contains legacy information.	Storage Nodes
/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

File name	Notes	Found on
/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/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 logs for the restricted admin ports on Admin Nodes and for the Load Balancer service, which provides load balancing of S3 and Swift traffic from clients to Storage Nodes.	Admin Nodes and Gateway Nodes
/var/local/log/persistence*	Contains logs for the Persistence service, which manages files on the root disk that need to persist across a reboot.	All nodes
/var/local/log/prometheus.log	For all nodes, contains the node exporter service log and the ade-exporter metrics service log. For Admin Nodes, also contains logs for the Prometheus and Alert Manager services.	All nodes
/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

File name	Notes	Found on
/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

NMS logs

File name	Notes	Found on
/var/local/log/nms.log	<ul style="list-style-type: none"> • Captures notifications from the Grid Manager and the Tenant Manager. • Captures events related to the operation of the NMS service, for example, alarm processing, email notifications, and configuration changes. • Contains XML bundle updates resulting from configuration changes made in the system. • Contains error messages related to the attribute downsampling done once a day. • Contains Java web server error messages, for example, page generation errors and HTTP Status 500 errors. 	Admin Nodes
/var/local/log/nms.errlog	<p>Contains error messages related to MySQL database upgrades.</p> <p>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.</p>	

Related concepts

[About the bycast.log](#) on page 251

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.

Related information

[Implementing S3 client applications](#)

Deployment and maintenance logs

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

File name	Notes	Found on
<code>/var/local/log/install.log</code>	Created during software installation. Contains a record of the installation events.	All nodes
<code>/var/local/log/expansion-progress.log</code>	Created during expansion operations. Contains a record of the expansion events.	Storage Nodes
<code>/var/local/log/gdu-server.log</code>	Created by the GDU service. Contains events related to provisioning and maintenance procedures managed by the primary Admin Node.	Primary Admin Node
<code>/var/local/log/send_admin_hw.log</code>	Created during installation. Contains debugging information related to a node's communications with the primary Admin Node.	All nodes
<code>/var/local/log/upgrade.log</code>	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
apache2 logs	<code>/var/local/log/apache2/access.log</code> <code>/var/local/log/apache2/error.log</code> <code>/var/local/log/apache2/other_vhosts_access.log</code>	Log files for apache2.	Admin Nodes
Archiving	<code>/var/local/log/dsierror.log</code>	Error information for TSM Client APIs.	Archive Nodes
MySQL	<code>/var/local/log/mysql.err</code> <code>/var/local/log/mysql-slow.log</code>	Log files generated by MySQL. The file <code>mysql.err</code> captures database errors and events such as startups and shutdowns. The file <code>mysql-slow.log</code> (the slow query log) captures the SQL statements that took more than 10 seconds to execute.	Admin Nodes

Category	File name	Notes	Found on
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 . Then, in the Tools section of the menu, select Grid Topology . Then select Topology > Site > Node > SSM > Events .	All nodes
NTP	/var/local/log/ntp.log	Log file for NTP error messages.	All nodes
	/var/lib/ntp/var/log/ntpstats/	The directory that contains NTP timing statistics. loopstats records loop filter statistics information. peerstats records peer statistics information.	
Samba	/var/local/log/samba/	The Samba log directory includes a log file for each Samba process (smb, nmb, and winbind) and every client hostname/IP.	Admin Node configured to export the audit share over CIFS

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.

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.

Note: 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.

Related tasks

[Collecting log files and system data](#) on page 139

You can use the Grid Manager to retrieve log files and system data (including configuration data) for your StorageGRID system.

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.

This is an example of an 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**. Then, in the Tools section of the menu, select **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

Error code	Meaning
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	
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	

Error number	Error code	Meaning	
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	
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	ENOCCSI	No CSI structure available	
051	EL2HLT	Level 2 halted	
052	EBADE	Invalid exchange	
053	EBADR	Invalid request descriptor	
054	EXFULL	Exchange full	
055	ENOANO	No anode	

Error number	Error code	Meaning	
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	
073	EDOTDOT	RFS specific error	
074	EBADMSG	Not a data message	
075	E_OVERFLOW	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	Cannot access a needed shared library	
080	ELIBBAD	Accessing a corrupted shared library	
081	ELIBSCN	.lib section in a.out corrupted	
082	ELIBMAX	Attempting to link in too many shared libraries	
083	ELIBEXEC	Cannot 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	

Error number	Error code	Meaning	
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	
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	Cannot 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 abort	
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	Cannot send after transport endpoint shutdown	
109	ETOOMANYREFS	Too many references: cannot 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			
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	

Error number	Error code	Meaning	
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	

Copyright and trademark

Copyright

Copyright © 2020 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>