



Monitor and troubleshoot

StorageGRID

NetApp
October 03, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-115/monitor/web-browser-requirements.html> on October 03, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Monitor and troubleshoot	1
Monitor a StorageGRID system	1
Using the Grid Manager for monitoring	1
Information you should monitor regularly	41
Managing alerts and alarms	80
Using SNMP monitoring	127
Collecting additional StorageGRID data	141
Alerts reference	177
Alarms reference (legacy system)	216
Log files reference	266
Troubleshoot a StorageGRID system	283
Overview of problem determination	283
Troubleshooting object and storage issues	291
Troubleshooting metadata issues	320
Troubleshooting certificate errors	326
Troubleshooting Admin Node and user interface issues	328
Troubleshooting network, hardware, and platform issues	333
Review audit logs	341
Audit message overview	341
Audit log file and message formats	348
Audit messages and the object lifecycle	366
Audit messages	372

Monitor and troubleshoot

Monitor a StorageGRID system

Learn how to monitor a StorageGRID system and how to assess issues that might occur. Lists all system alerts.

- [Using the Grid Manager for monitoring](#)
- [Information you should monitor regularly](#)
- [Managing alerts and alarms](#)
- [Using SNMP monitoring](#)
- [Collecting additional StorageGRID data](#)
- [Troubleshooting a StorageGRID system](#)
- [Alerts reference](#)
- [Alarms reference \(legacy system\)](#)
- [Log files reference](#)

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](#)
- [Viewing the Dashboard](#)
- [Viewing the Nodes page](#)

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

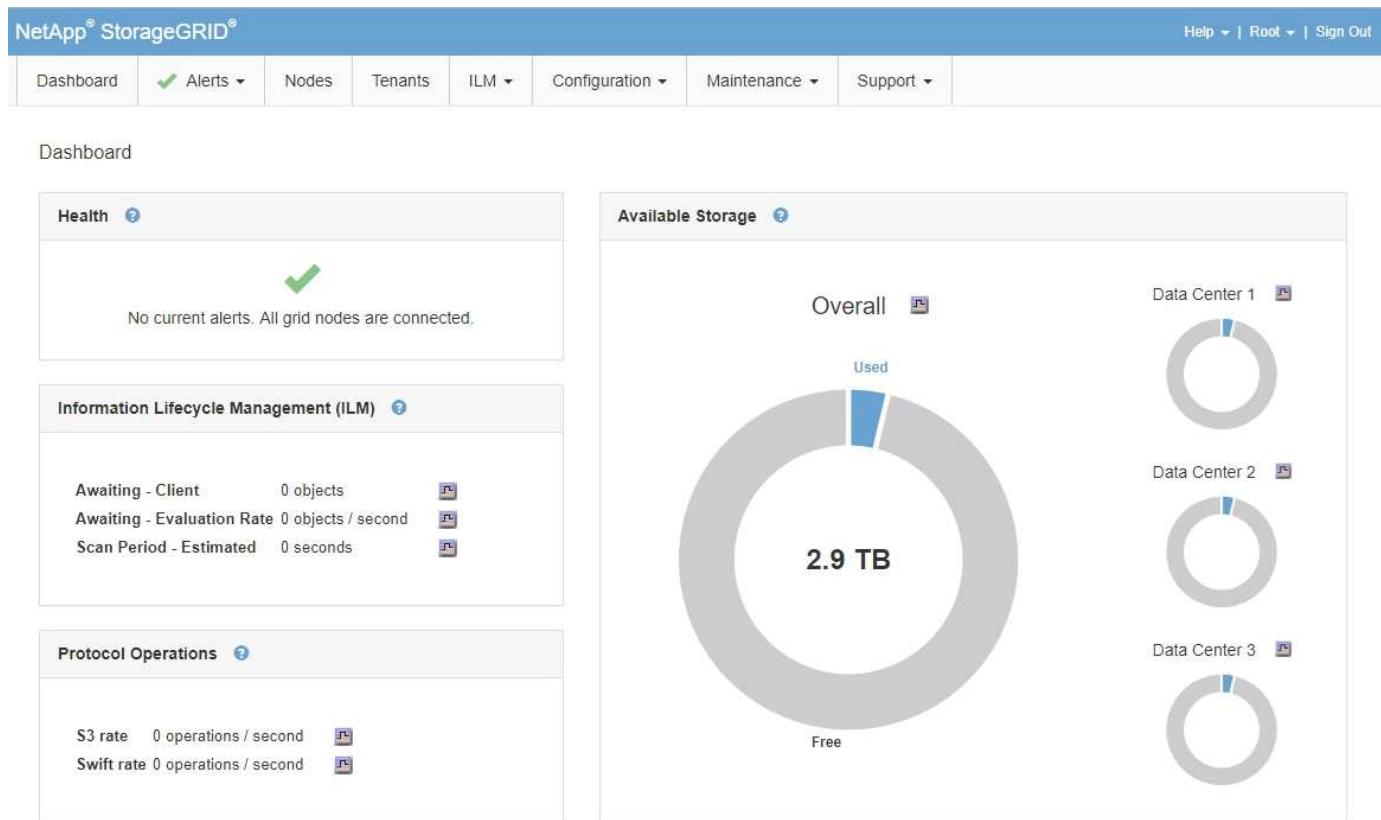
You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024

Browser width	Pixels
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 > Alarms (Legacy) > 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 > System > License page. 	<ul style="list-style-type: none"> • Monitoring node connection states • Viewing current alerts • Viewing resolved alerts • Viewing legacy alarms • Administer StorageGRID

Available Storage panel

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

Information Lifecycle Management (ILM) panel

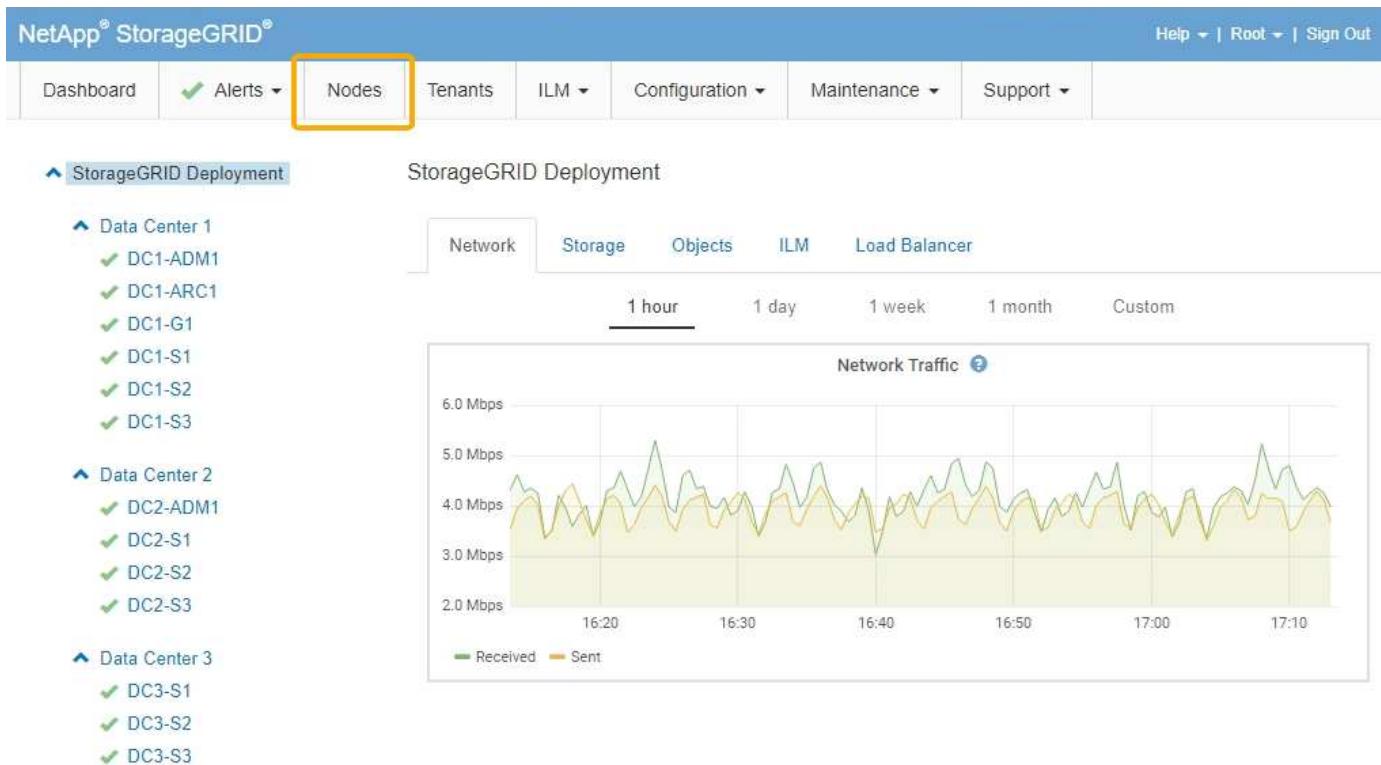
Description	View additional details	Learn more
<p>Displays current ILM operations and ILM queues for your system. You can use this information to monitor your system's workload.</p> <ul style="list-style-type: none"> Awaiting - Client: The total number of objects awaiting ILM evaluation from client operations (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 Administer 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.</p> <p>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 Use S3 Use Swift

Viewing the Nodes page

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



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

Connection state icons

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

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



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

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

Alert icons

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

- **Critical** : An abnormal condition exists that has stopped the normal operations of a StorageGRID node or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved.
- **Major** : An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service.
- **Minor** : The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that do not clear on their own to ensure they do not result in a more serious problem.
- **Normal** : No alerts are active, and the node is connected to the grid.

Viewing details for a system, site, or node

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

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

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)

Overview **Hardware** Network Storage Objects ILM Events Tasks

Node Information 

Name	DC1-S1
Type	Storage Node
ID	5bf57bd4-a68d-467e-b866-bfe09a5c6b96
Connection State	 Connected
Software Version	11.4.0 (build 20200328.0051.269ac98)
IP Addresses	10.96.101.111 Show more 

Alerts 

 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.



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.

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 information

[Monitoring node connection states](#)

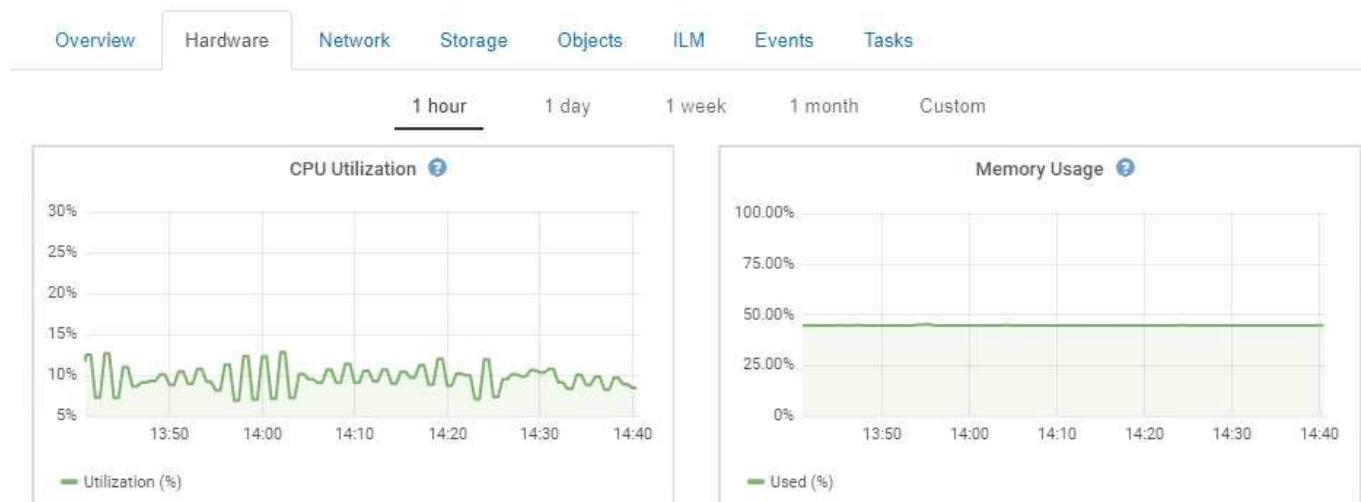
[Viewing current alerts](#)

[Viewing a specific alert](#)

Viewing the Hardware tab

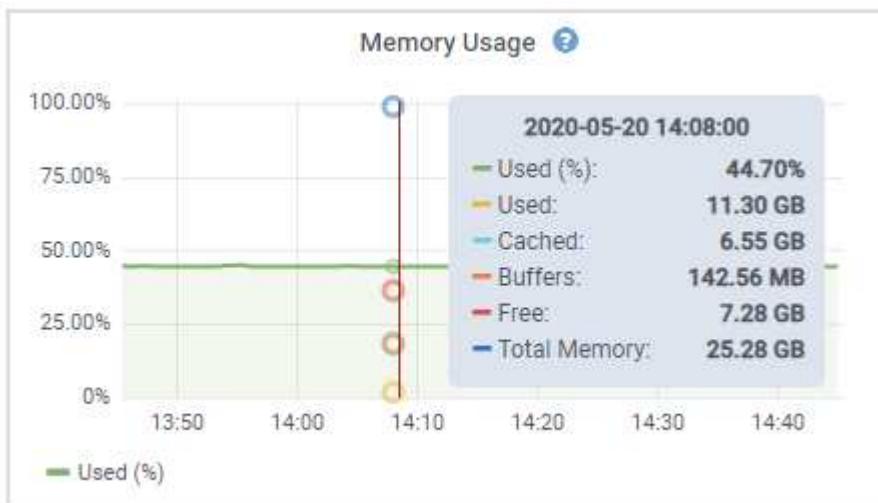
The Hardware tab displays CPU utilization and memory usage for each node, and additional hardware information about appliances.

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

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



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

Related information

[Viewing information about appliance Storage Nodes](#)

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

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

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

DC1-S1-226 (Storage Node)

[Overview](#)[Hardware](#)[Network](#)[Storage](#)[Objects](#)[ILM](#)[Events](#)[1 hour](#)[1 day](#)[1 week](#)[1 month](#)[1 year](#)[Custom](#)

Network Traffic



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 information

[Monitoring network connections and performance](#)

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.

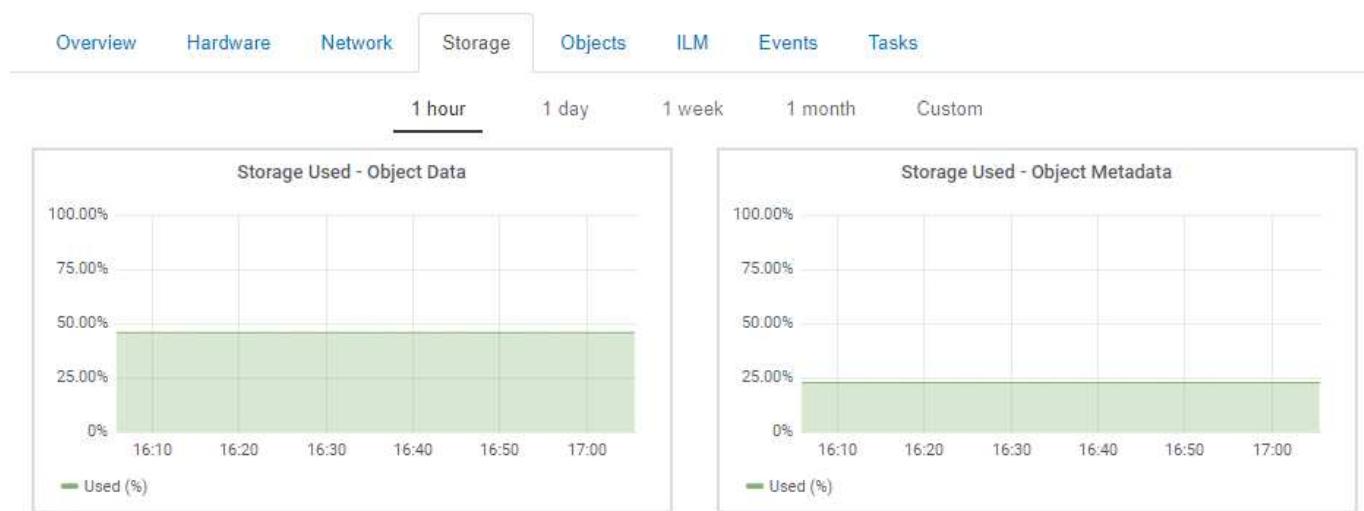
Storage Used graphs

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.



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

DC1-SN1-99-88 (Storage Node)



Disk Devices, Volumes, and Object Store tables

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

Disk Devices				
Name	World Wide Name	I/O Load	Read Rate	Write Rate
croot(8:1,sda1)	N/A	0.03%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.85%	0 bytes/s	58 KB/s
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	81 bytes/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s

Volumes						
Mount Point	Device	Status	Size	Available	Write Cache Status	
/	croot	Online	21.00 GB	14.90 GB	 Unknown	
/var/local	cvloc	Online	85.86 GB	84.10 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	sde	Online	107.32 GB	107.18 GB	 Enabled	

Object Stores							
ID	Size	Available	Replicated Data	EC Data	Object Data (%)	Health	
0000	107.32 GB	96.45 GB	 250.90 KB	 0 bytes	 0.00%	No Errors	
0001	107.32 GB	107.18 GB	 0 bytes	 0 bytes	 0.00%	No Errors	
0002	107.32 GB	107.18 GB	 0 bytes	 0 bytes	 0.00%	No Errors	

Related information

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

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

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

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 information

[Monitoring events](#)

[Displaying charts and graphs](#)

[Resetting event counts](#)

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.

What you'll need

- 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.

Reboot

Shuts down and restarts the node.

Reboot

Maintenance Mode

Places the appliance's compute controller 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.



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.

DC3-S3 (Storage Node)



Reboot

Reboot shuts down and restarts the node.

Reboot

4. Click **Reboot**.

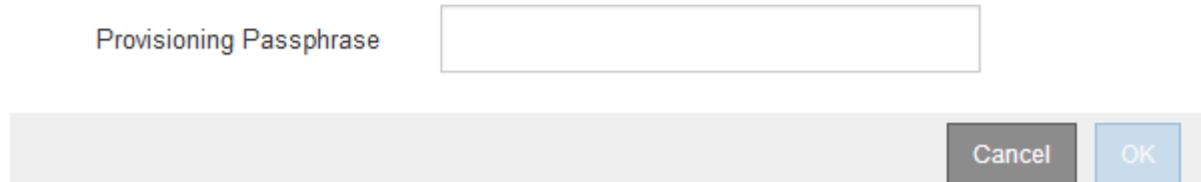
A confirmation dialog box appears.



Reboot shuts down and restarts a node, based on where the node is installed:

- Rebooting a VMware node reboots the virtual machine.
- Rebooting a Linux node reboots the container.
- Rebooting a StorageGRID Appliance node reboots the compute controller.

If you are ready to reboot this node, enter the provisioning passphrase and click OK.



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 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

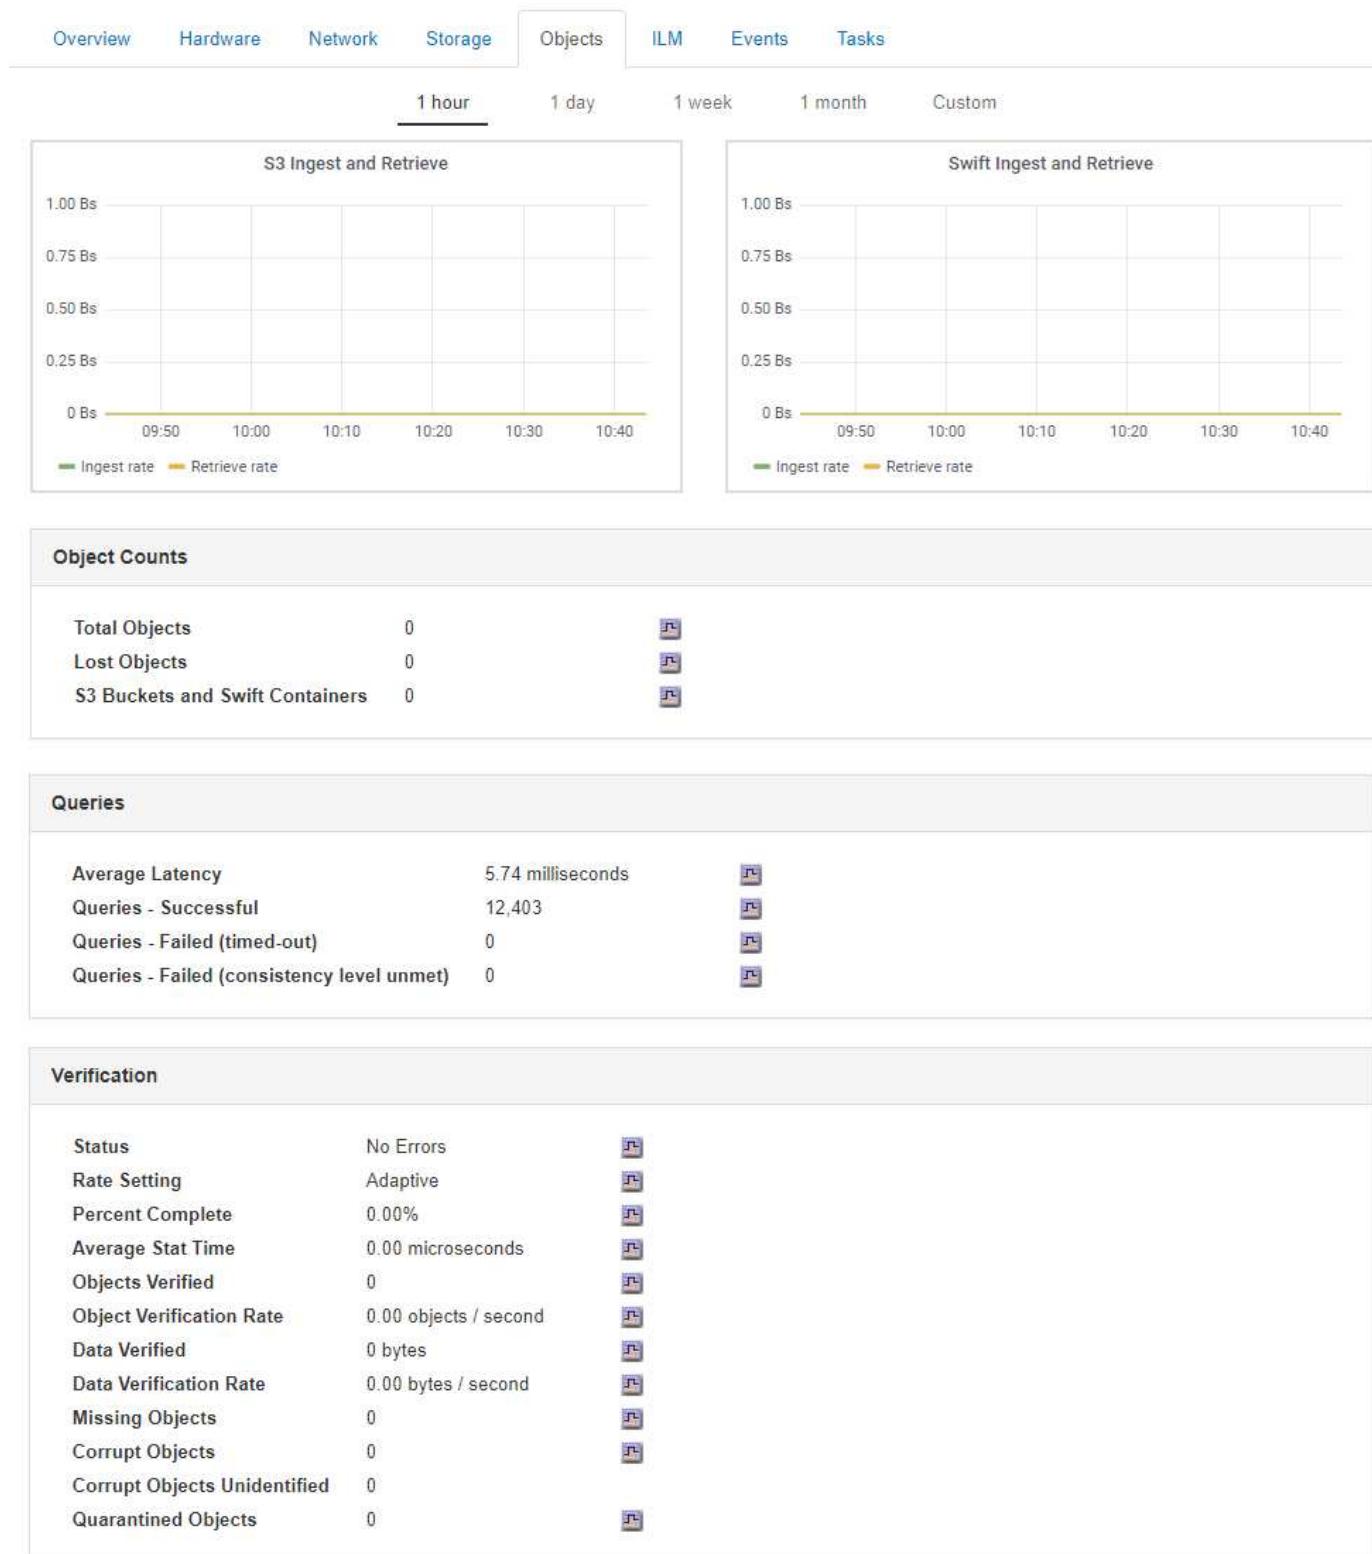
[SG100 & SG1000 services appliances](#)

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

[Use S3](#)

[Use Swift](#)

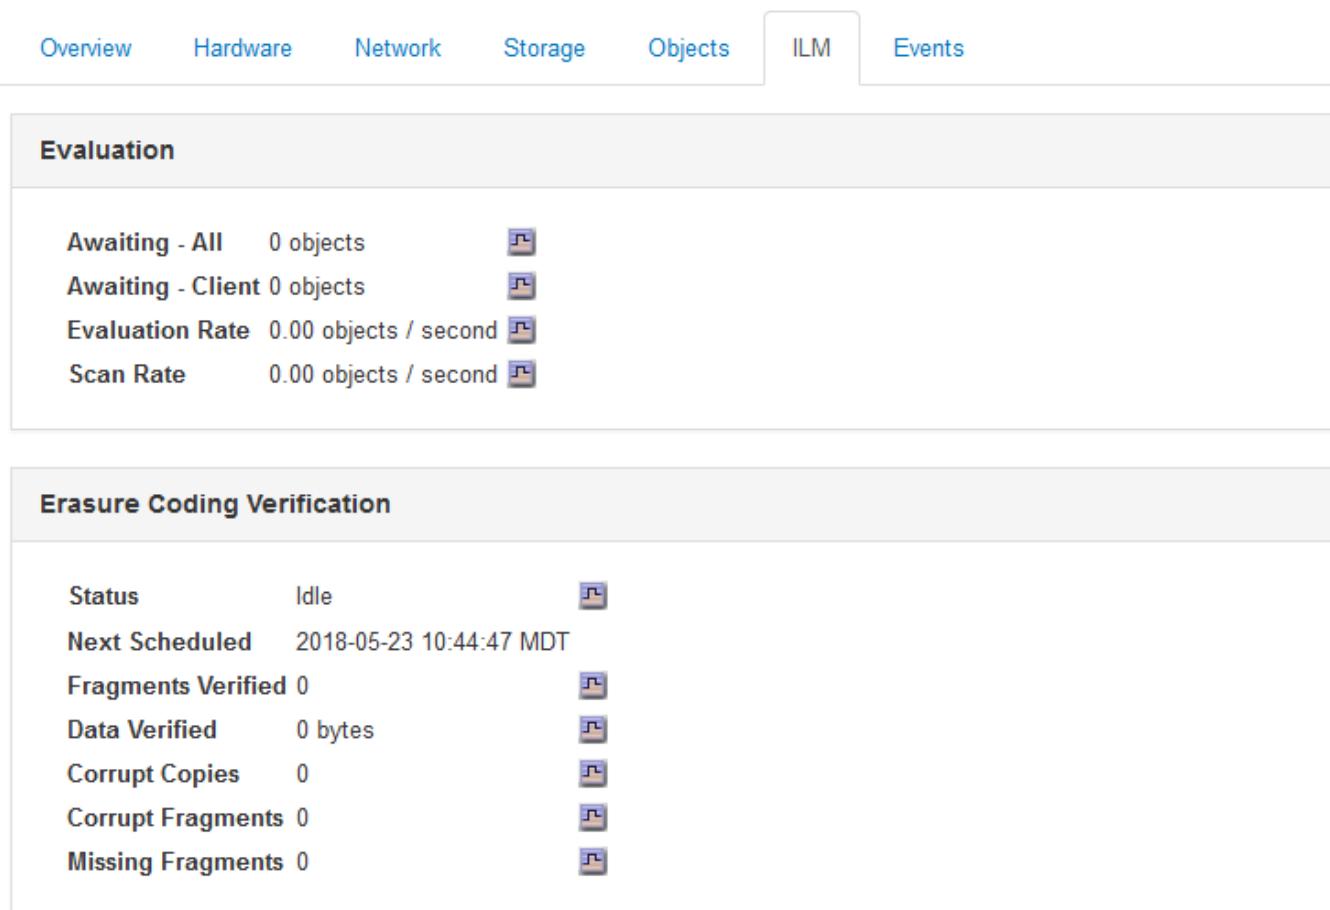
Viewing the ILM tab

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)



The screenshot shows the ILM tab for the DC1-S1 Storage Node. The top navigation bar includes tabs for Overview, Hardware, Network, Storage, Objects, ILM (which is selected and highlighted in blue), and Events. The ILM section is divided into two main sections: Evaluation and Erasure Coding Verification.

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 information

[Monitoring information lifecycle management](#)

[Administer StorageGRID](#)

Viewing the Load Balancer tab

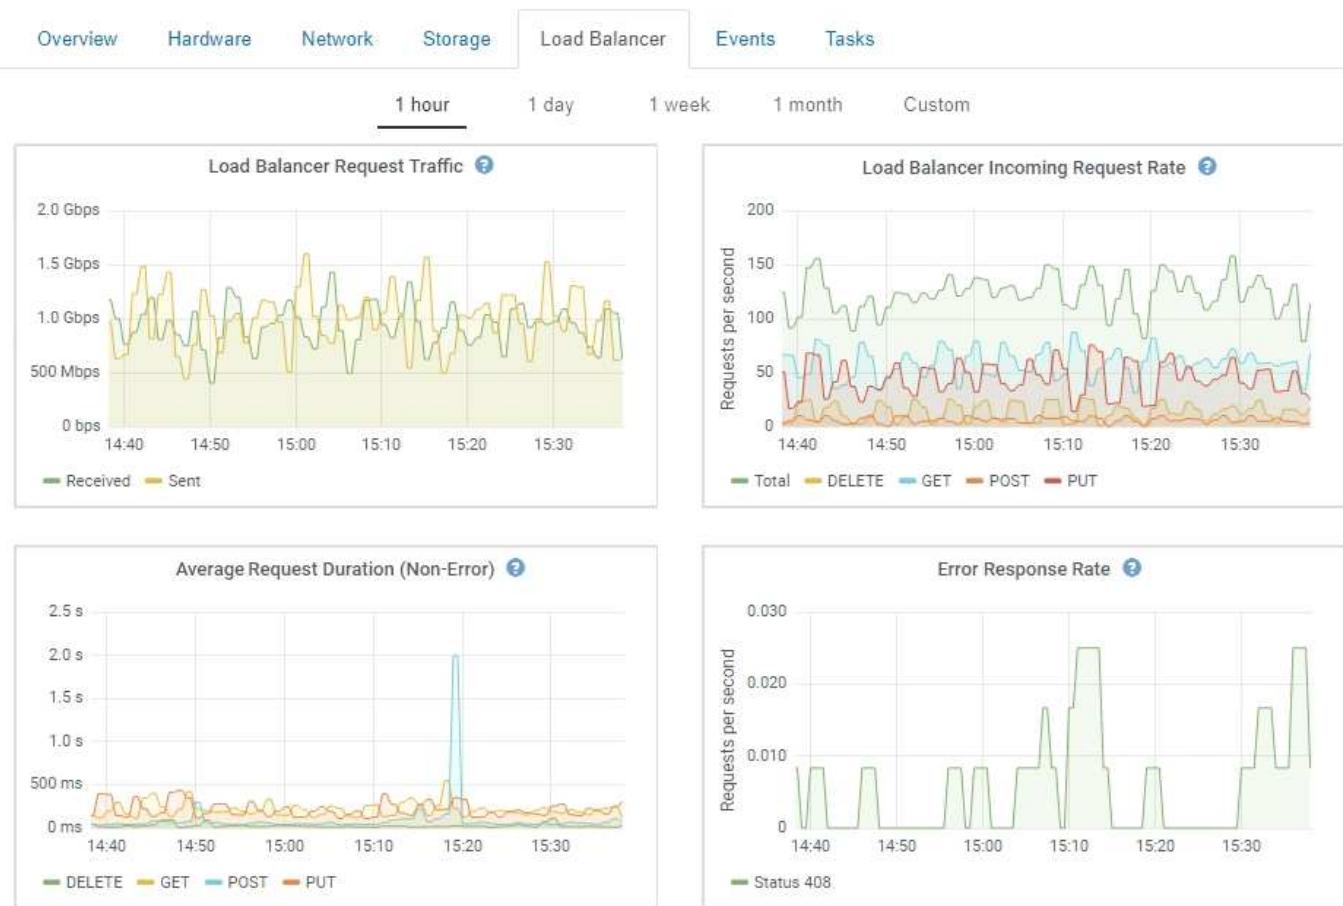
The Load Balancer tab includes performance and diagnostic graphs related to the operation of the Load Balancer service.

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.”

DC1-SG1000-ADM (Admin Node)



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.



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 information

[Monitoring load balancing operations](#)

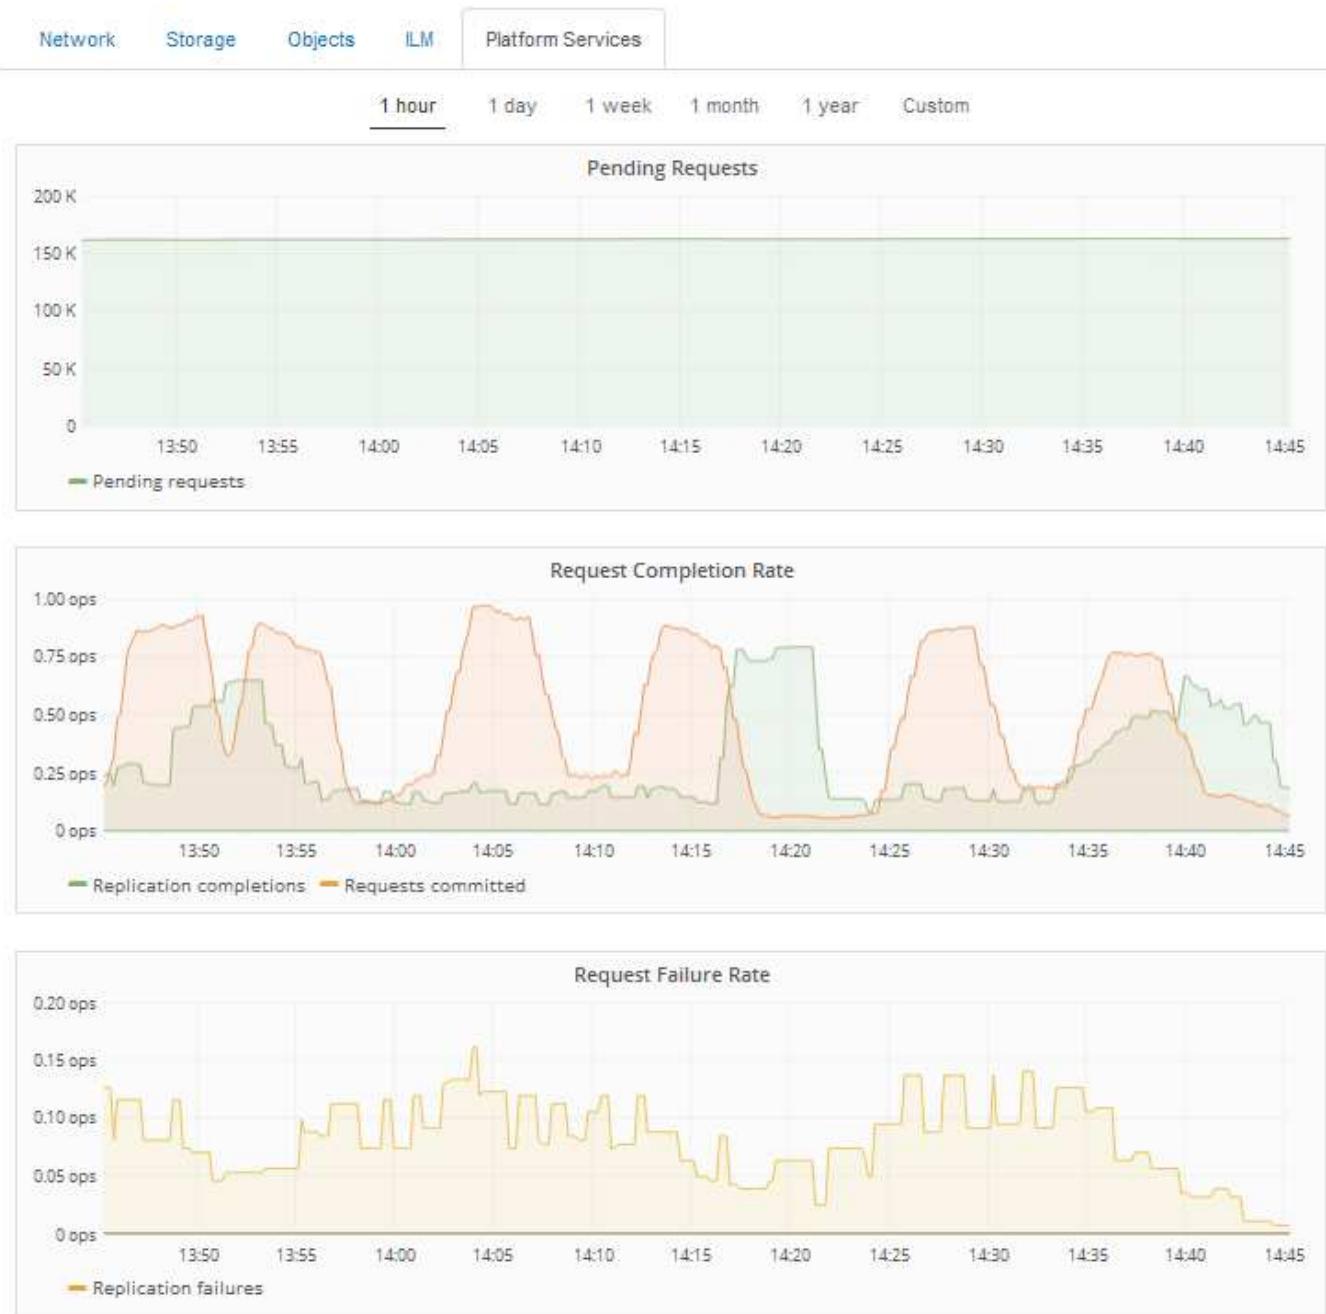
[Administer StorageGRID](#)

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.

Data Center 1



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

Related information

[Administer 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, controller firmware version, network resources, network interfaces,

network addresses, and receive and transmit data.

Steps

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

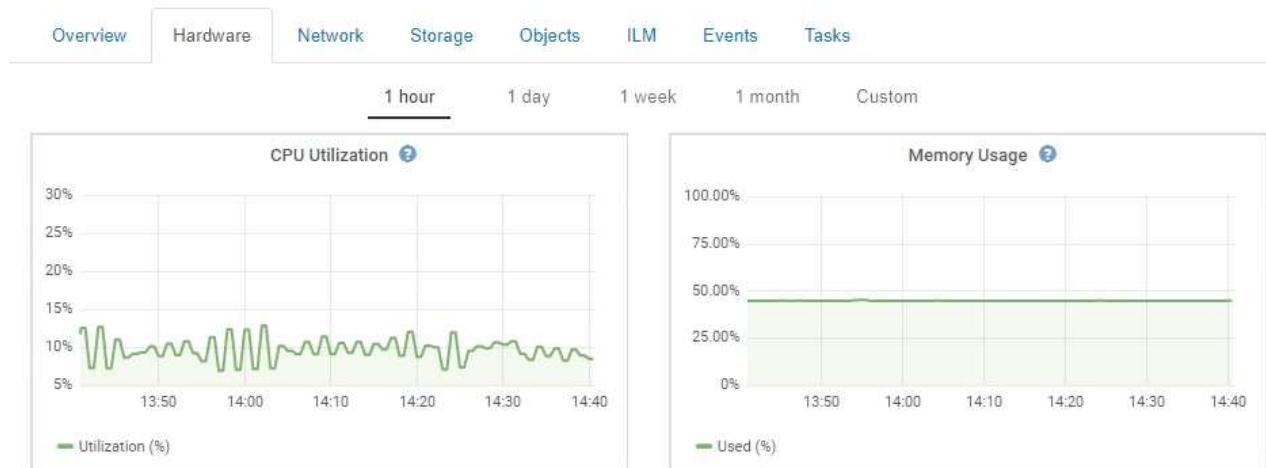
The Node Information table on the Overview tab displays the node's ID and name, the node type, the software version installed, and the IP addresses associated with the node. The Interface column contains the name of the interface, as follows:

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



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



Some fields, such as Compute Controller BMC IP and Compute Hardware, appear only for appliances with that feature.

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

StorageGRID Appliance

Appliance Model	SG6060
Storage Controller Name	StorageGRID-NetApp-SGA-000-012
Storage Controller A Management IP	10.224.1.79
Storage Controller B Management IP	10.224.1.80
Storage Controller WWID	6d039ea000016fc700000005fac58f4
Storage Appliance Chassis Serial Number	721924500062
Storage Controller Firmware Version	08.70.00.02
Storage Hardware	Needs Attention 
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	DDP
Storage Connectivity	Nominal 
Overall Power Supply	Nominal 
Compute Controller BMC IP	10.224.0.13
Compute Controller Serial Number	721917500067
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
721924500062	99	Nominal 	N/A	Nominal	Nominal	Nominal	60	58	4.00 TB	2	800.17 GB	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.

Field in the Appliance table	Description
Storage Appliance Chassis Serial Number	The chassis serial number of the appliance.
Storage Controller Firmware Version	The version of the firmware on the storage controller for this appliance.
Storage Hardware	<p>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.</p> <p>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.</p>
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.

Field in the Appliance table	Description
Compute Controller BMC IP	<p>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.</p> <p>This field is not displayed for appliance models that do not contain a BMC.</p>
Compute Controller Serial Number	The serial number of the compute controller.
Compute Hardware	<p>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.</p>
Compute Controller CPU Temperature	The temperature status of the compute controller's CPU.
Compute Controller Chassis Temperature	The temperature status of the compute controller.

Column in the Storage Shelves table	Description
Shelf Chassis Serial Number	The serial number for the storage shelf chassis.
Shelf ID	<p>The numeric identifier for the storage shelf.</p> <ul style="list-style-type: none"> • 99: Storage controller shelf • 0: First expansion shelf • 1: Second expansion shelf <p>Note: Expansion shelves apply to the SG6060 only.</p>
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.

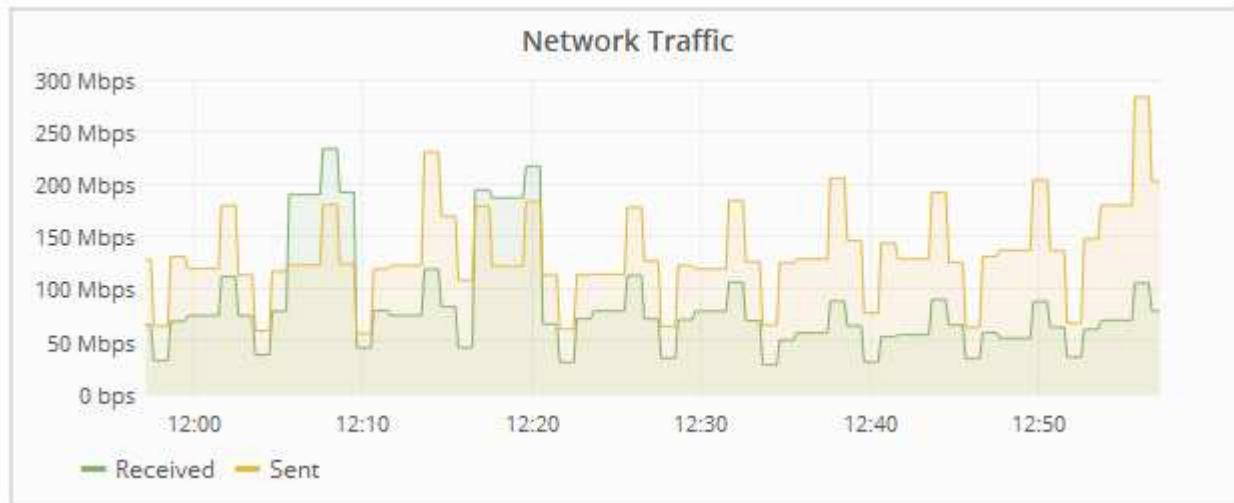
Column in the Storage Shelves table	Description
Fan Status	The overall status of the cooling fans in the storage shelf.
Drive Slots	The total number of drive slots in the storage shelf.
Data Drives	The number of drives in the storage shelf that are used for data storage.
Data Drive Size	The effective 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.



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.

- 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

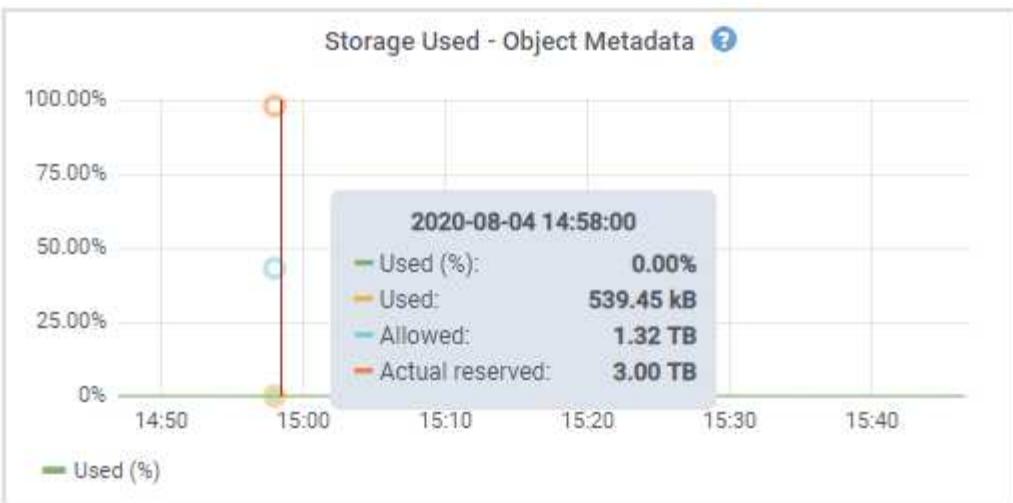
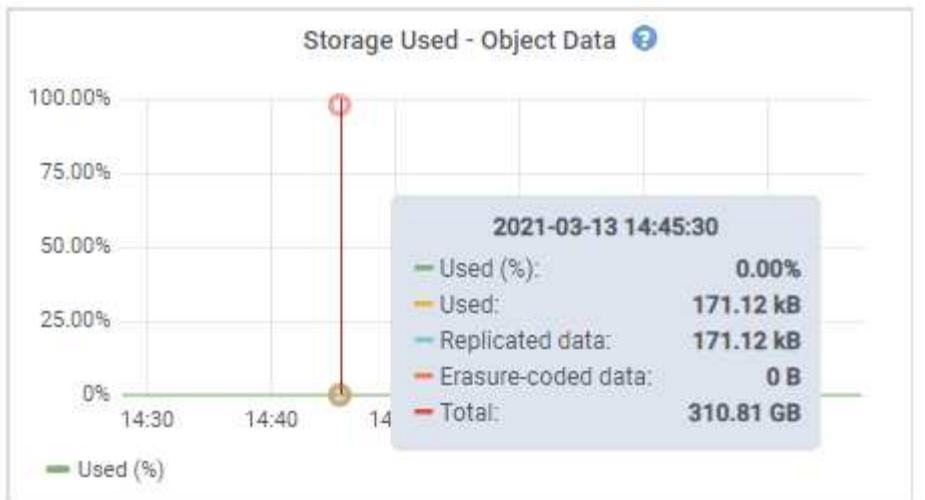
Receive

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	3 KB/s	
cvloc(8:2,sda2)	N/A	0.85%	0 bytes/s	58 KB/s	
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	81 bytes/s	
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s	
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s	

Volumes					
Mount Point	Device	Status	Size	Available	Write Cache Status
/	croot	Online	21.00 GB	14.90 GB	 Unknown
/var/local	cvloc	Online	85.86 GB	84.10 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	sde	Online	107.32 GB	107.18 GB	 Enabled

Object Stores							
ID	Size	Available	Replicated Data	EC Data	Object Data (%)	Health	
0000	107.32 GB	96.45 GB	 250.90 KB	 0 bytes	 0.00%	No Errors	
0001	107.32 GB	107.18 GB	 0 bytes	 0 bytes	 0.00%	No Errors	
0002	107.32 GB	107.18 GB	 0 bytes	 0 bytes	 0.00%	No Errors	

Related information

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

Viewing the SANtricity System Manager tab

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

The SANtricity System Manager tab is shown for storage appliance nodes.

Using SANtricity System Manager, you can do the following:

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



To use SANtricity System Manager to configure a proxy for E-Series AutoSupport, see the instructions in [administeringStorageGRID](#).

Administer StorageGRID

To access SANtricity System Manager through Grid Manager, you must have the Storage Appliance Administrator permission or Root Access permission.



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



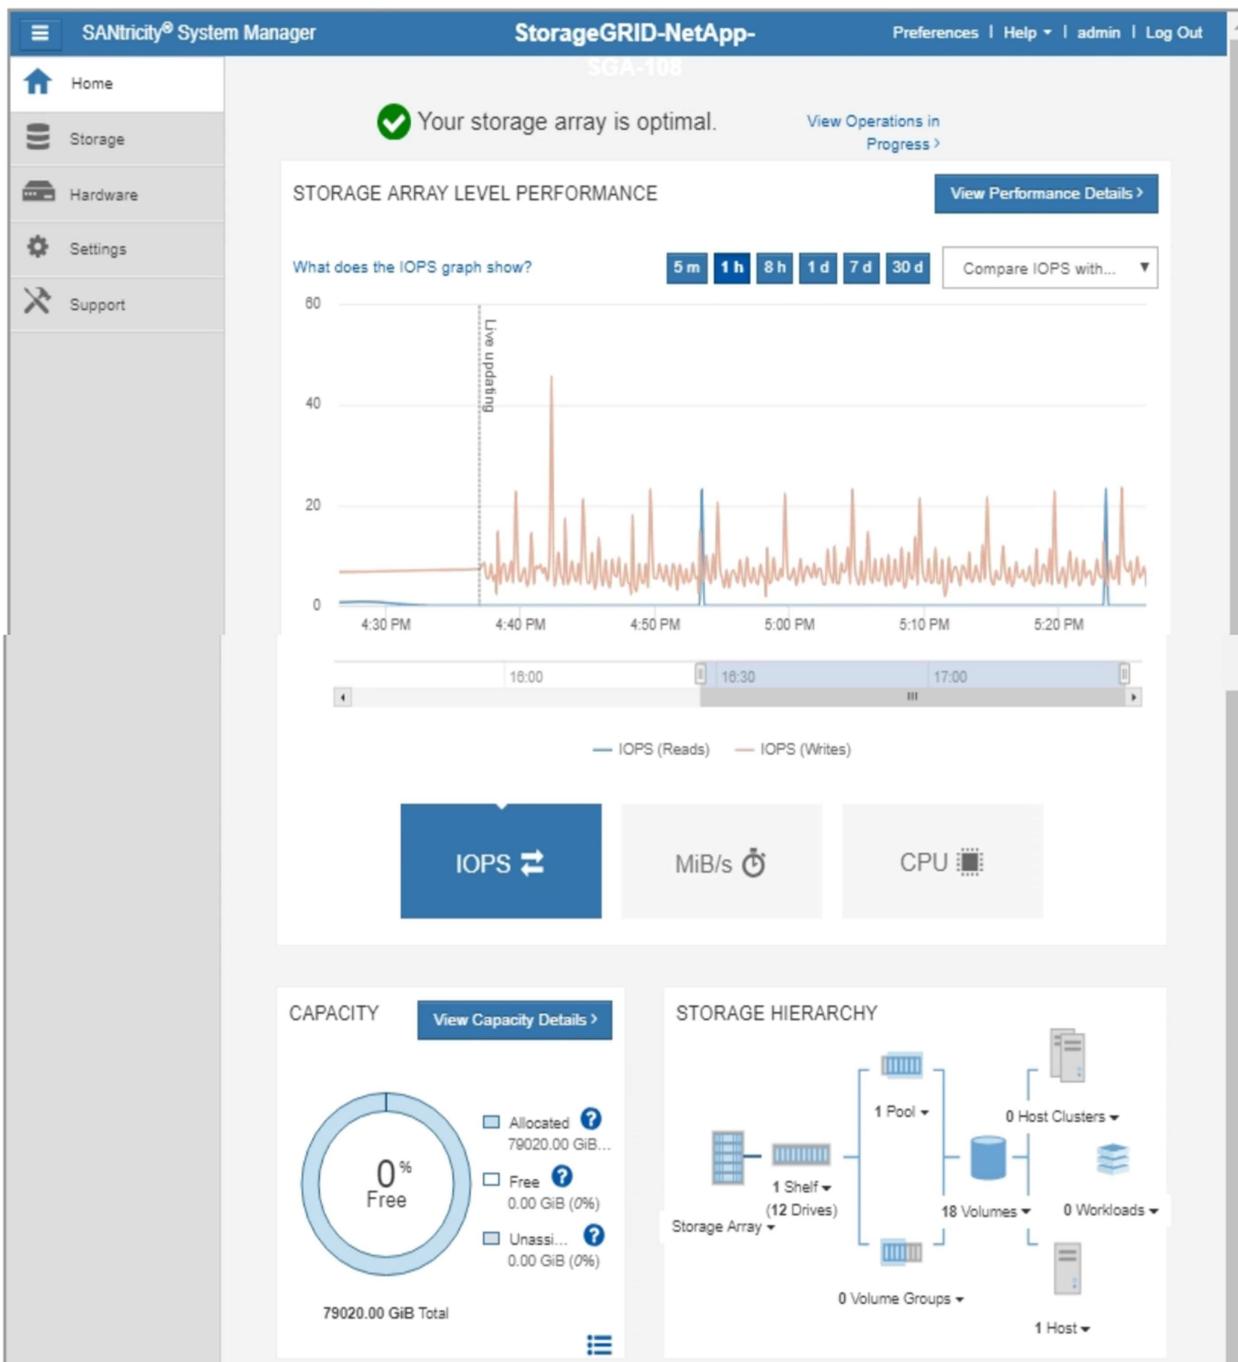
Accessing SANtricity System Manager from the Grid Manager is generally meant only to monitor appliance hardware and configure E-Series AutoSupport. Many features and operations within SANtricity System Manager such as upgrading firmware do not apply to monitoring your StorageGRID appliance. To avoid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

The tab displays the home page of SANtricity System Manager

Use SANtricity System Manager to monitor and manage the hardware components in this storage appliance. From SANtricity System Manager, you can review hardware diagnostic and environmental information as well as issues related to the drives.

Note: Many features and operations within SANtricity Storage Manager do not apply to your StorageGRID appliance. To avoid issues, always follow the hardware installation and maintenance instructions for your appliance model.

Open [SANtricity System Manager](#) in a new browser tab.



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

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

For more details on viewing the information accessible from the SANtricity System Manager tab, see the information in the [NetApp E-Series Systems Documentation Center](#)

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).

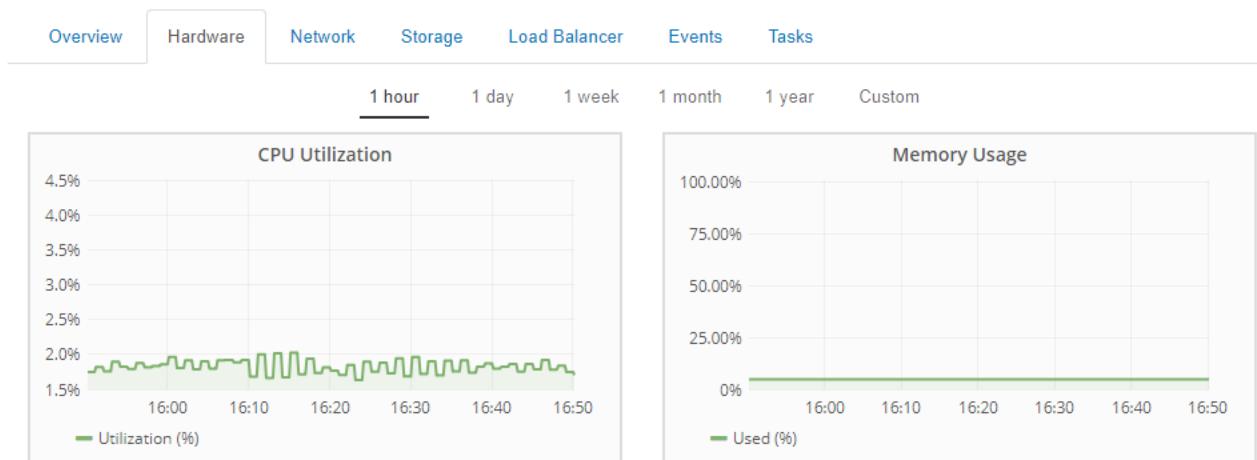
Node Information ?

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.

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



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

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.

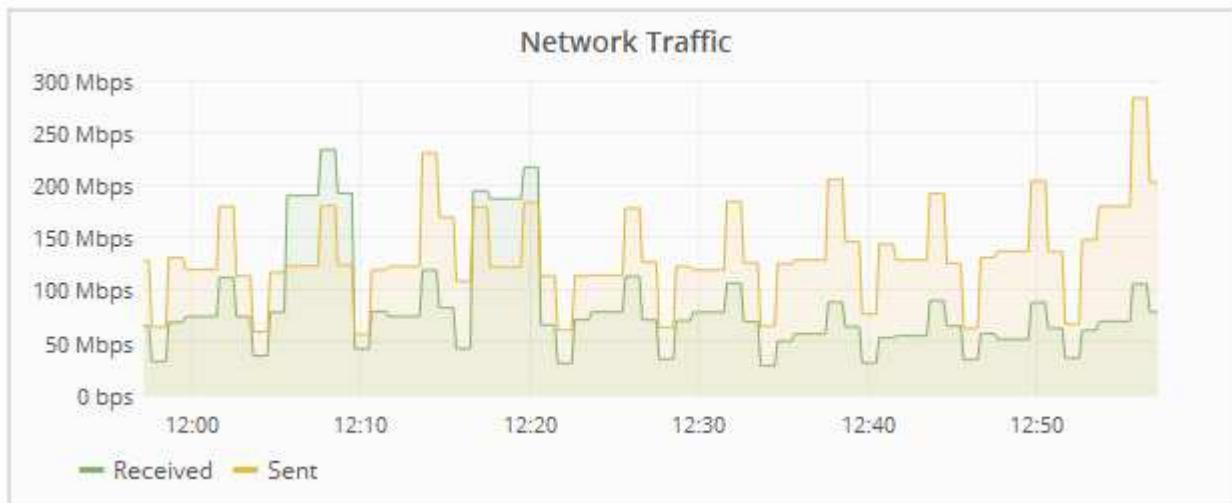
Field in the Appliance table	Description
Storage RAID Mode	The RAID mode for the appliance.
Overall Power Supply	The status of all power supplies in the appliance.
Compute Controller BMC IP	<p>The IP address of the baseboard management controller (BMC) port in the compute controller. You can use this IP to connect to the BMC interface to monitor and diagnose the appliance hardware.</p> <p>This field is not displayed for appliance models that do not contain a BMC.</p>
Compute Controller Serial Number	The serial number of the compute controller.
Compute Hardware	The status of the compute controller hardware.
Compute Controller CPU Temperature	The temperature status of the compute controller's CPU.
Compute Controller Chassis Temperature	The temperature status of the compute controller.

c. Confirm that all statuses are "Nominal."

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

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

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



a. Review the Network Interfaces section.

Network Interfaces					
Name	Hardware Address	Speed	Duplex	Auto 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.



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.

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 & SG1000 services appliances](#)

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. This section contains the following sub-sections:

- [Monitoring system health](#)
- [Monitoring storage capacity](#)
- [Monitoring information lifecycle management](#)
- [Monitoring performance, networking, and system resources](#)
- [Monitoring tenant activity](#)
- [Monitoring archival capacity](#)
- [Monitoring load balancing operations](#)
- [Applying hotfixes or upgrading software if necessary](#)

What to monitor	Frequency
The system health data shown on the Grid Manager DashboardNote 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"> <li data-bbox="159 572 747 604">• Query latency <li data-bbox="159 620 747 652">• Connectivity and networking <li data-bbox="159 669 747 701">• 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

Monitoring system health

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

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 issues that might be affecting your system. You should investigate any issues that are shown on the Dashboard.



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

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	<p>Appears if any alarms (legacy system) are currently active. Click the link to see the details on the Support > Alarms (legacy) > Current Alarms page.</p> <p>Note: While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.</p>
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 > System > License page.

Related information

[Administer StorageGRID](#)

[Setting up email notifications for alerts](#)

[Using SNMP monitoring](#)

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.

What you'll need

- 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.



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

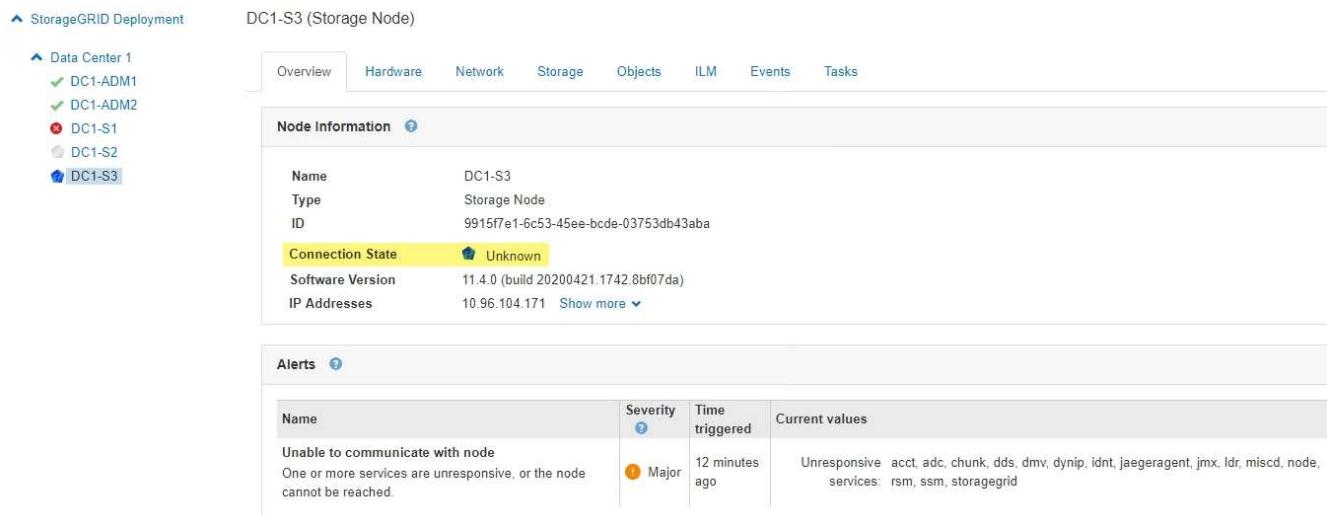
- **Not connected - Administratively down** : The node is not connected to the grid for an expected reason. For example, the node, or services on the node, has been gracefully shut down, the node is rebooting, or the software is being upgraded. One or more alerts might also be active.
- **Connected** : 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 Node Information panel for DC1-S3. The Connection State is listed as Unknown. An alert for 'Unable to communicate with node' is active, with a Major severity, triggered 12 minutes ago, and current values indicating unresponsive services: acct, adc, chunk, dds, dmv, dynip, idnt, jaegeragent, jmx, ldr, miscd, node, rsm, ssm, storagegrid.

2. If a node has a blue icon, follow these steps:

- 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.

- 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.

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

Related information

[Alerts reference](#)

[Maintain & recover](#)

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.

What you'll need

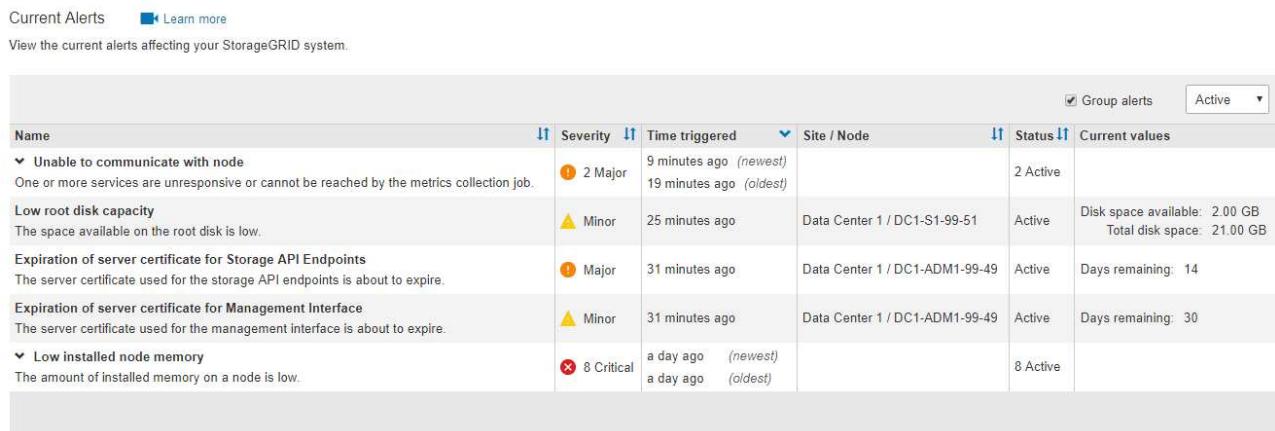
- You must be signed in to the Grid Manager using a supported browser.

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.



Name	Severity	Time triggered	Site / Node	Status	Current values
Unable to communicate with node	2 Major	9 minutes ago (newest) 19 minutes ago (oldest)		2 Active	
Low root disk capacity	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	Major	31 minutes ago	Data Center 1 / DC1-ADM1-99-49	Active	Days remaining: 14
Expiration of server certificate for Management Interface	Minor	31 minutes ago	Data Center 1 / DC1-ADM1-99-49	Active	Days remaining: 30
Low installed node memory	8 Critical	a day ago (newest) a day ago (oldest)		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.

Column header	Description
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.

							<input checked="" type="checkbox"/> Group alerts	Active
Name	Severity	Time triggered	Site / Node	Status	Current values			
▲ <u>Low object data storage</u> 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 information

[Viewing a specific alert](#)

[Silencing alert notifications](#)

Viewing resolved alerts

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

What you'll need

- 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	Search	
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

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.</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.
Time resolved	How long ago the alert was resolved.
Site/Node	The name of the site and node where the alert occurred.

Column header	Description
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.

- 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)

- Select one or more severities from the **Severity** drop-down menu to filter on resolved alerts of a specific severity.

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

- Select one or more nodes from the **Node** drop-down menu to filter on resolved alerts related to a specific node.

- 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 information

[Viewing a specific alert](#)

[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.

What you'll need

- 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:

Column header	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.
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.



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) (oldest)		8 Active
Low installed node memory The amount of installed memory on a node is low.	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

X 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	<p>The severity of the alert.</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.
data values	<p>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.</p>

4. 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.



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.

5. To view the current conditions for the alert rule:

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

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

Low installed node memory

Major node_memory_MemTotal_bytes < 24000000000

Critical node_memory_MemTotal_bytes < 12000000000

Total RAM size
8.38 GB
Condition
View conditions | **Edit rule**

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

6. 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.



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.

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

Related information

[Silencing alert notifications](#)

[Editing an alert rule](#)

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.

What you'll need

- 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.

Health

Administratively Down 1

Critical 5

License Status 1

Grid details Current alerts (5) Recently resolved alerts (1) **Legacy alarms (5)** License

The **Legacy alarms** count on the Dashboard is incremented whenever a legacy alarm is triggered. 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.



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

Steps

1. 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 > Alarms (legacy) > 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

<input type="checkbox"/> Show Acknowledged Alarms (1 - 1 of 1)						
Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value
	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	

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.

Icon	Color	Alarm severity	Meaning
	Red	Critical	The node is connected to the grid, but an abnormal condition exists that has stopped normal operations. You should address the issue immediately.

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

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

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

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

Apply Changes

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

Related information

[Alarms reference \(legacy system\)](#)

[Acknowledging current alarms \(legacy system\)](#)

[Disabling alarms \(legacy system\)](#)

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 information

[Viewing the Storage tab](#)

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.

What you'll need

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

About this task

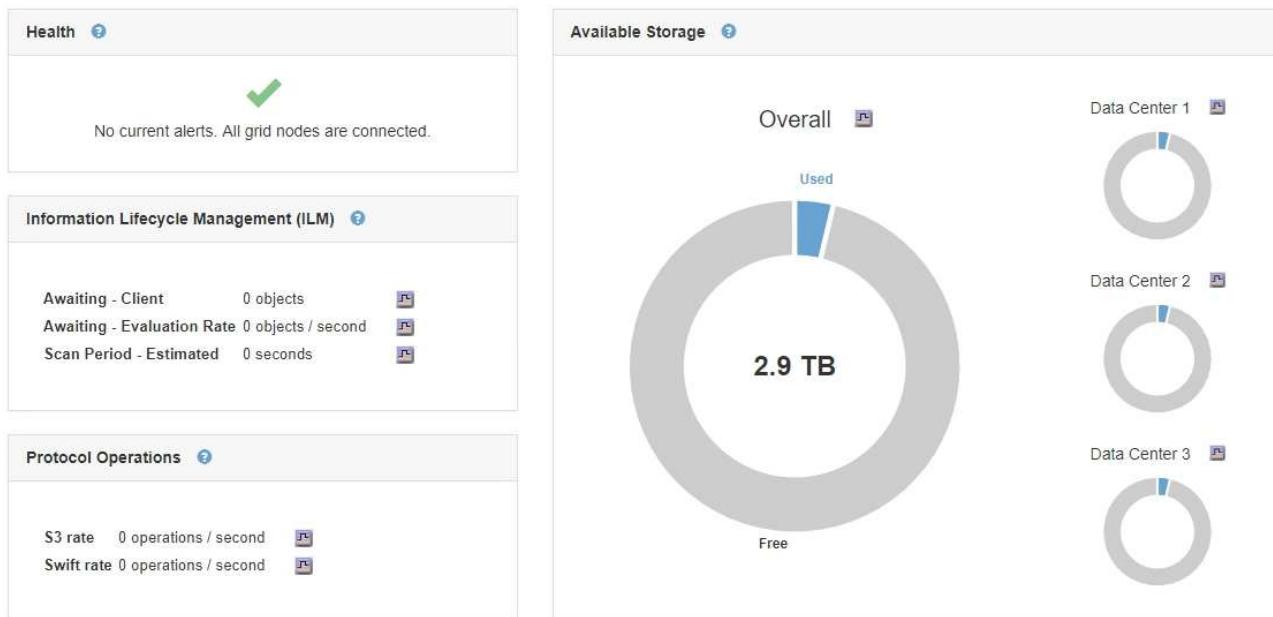
The Dashboard in the Grid Manager lets you quickly assess how much storage is available for the entire grid and for each data center. The Nodes page provides more detailed values for object data and object metadata.

Steps

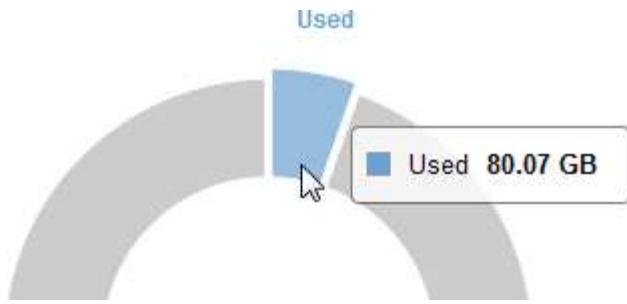
1. Assess how much storage is available for the entire grid and for each data center.
 - a. Select **Dashboard**.
 - b. In the Available Storage panel, note the overall summary of free and used storage capacity.



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.



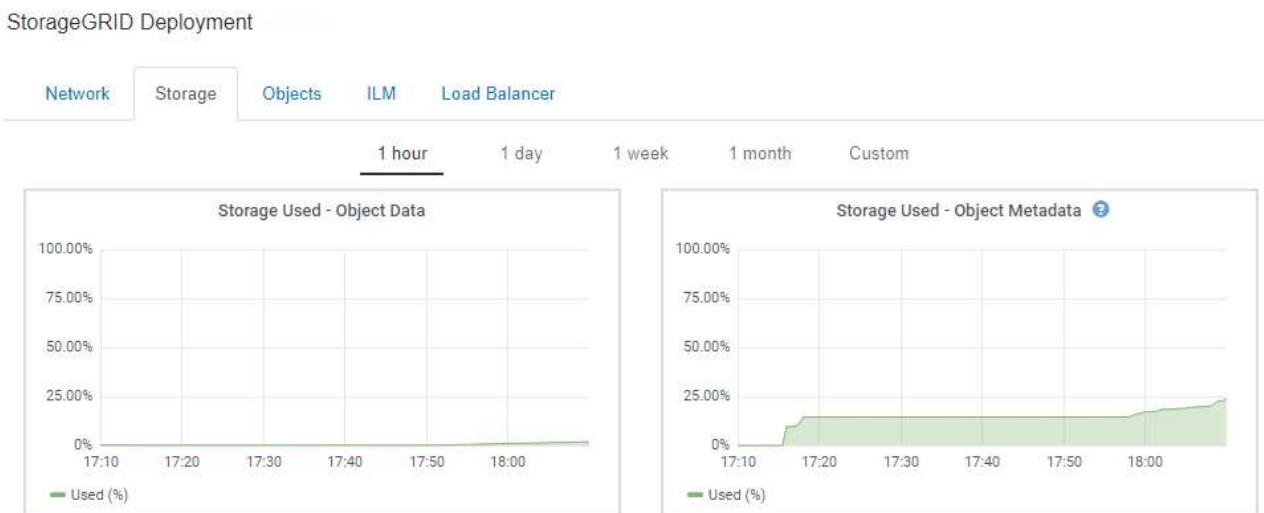
- d. For multi-site grids, review the chart for each data center.
- e. Click the chart icon  for the overall chart or for an individual data center to view a graph showing capacity usage over time.

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.



The total values for a site or the grid do not include nodes that have not 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 > Tools > Grid Topology**.

b. Select **grid > Overview > Main**.

Grid Topology

- StorageGRID Deployment
 - Data Center 1
 - Data Center 2
 - Data Center 3

Overview: Summary - StorageGRID Deployment
Updated: 2019-03-01 11:50:40 MST

Storage Capacity

Storage Nodes Installed:	9	
Storage Nodes Readable:	9	
Storage Nodes Writable:	9	
Installed Storage Capacity:	2,898 GB	
Used Storage Capacity:	100 GB	
Used Storage Capacity for Data:	2.31 MB	
Used Storage Capacity for Metadata:	5.82 MB	
Usable Storage Capacity:	2,797 GB	
Percentage Storage Capacity Used:	3.465 %	
Percentage Usable Storage Capacity:	96.535 %	

ILM Activity

Awaiting - All:	0	
Awaiting - Client:	0	
Scan Rate:	0 Objects/s	
Scan Period - Estimated:	0 us	
Awaiting - Evaluation Rate:	0 Objects/s	
Repairs Attempted:	0	

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.



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

[Expand your grid](#)

Monitoring storage capacity for each Storage Node

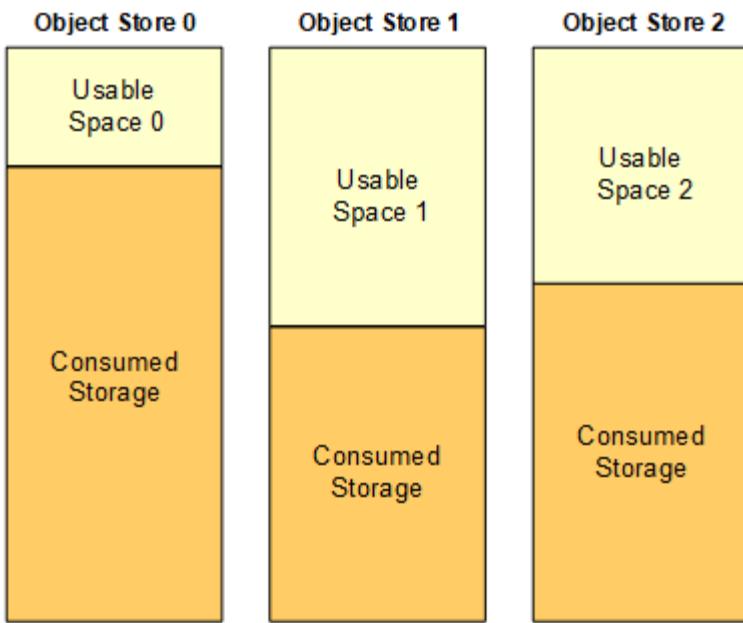
You must monitor the total usable space for each Storage Node to ensure that the node has enough space for new object data.

What you'll need

- 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

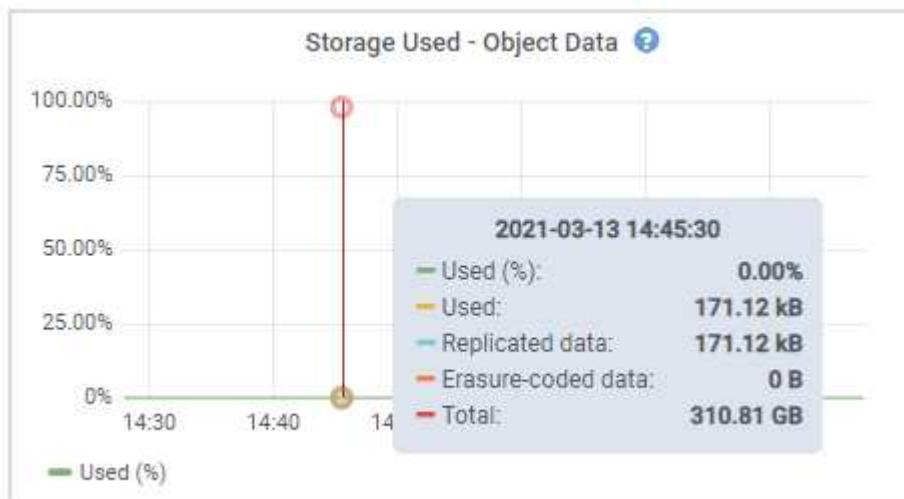
1. Select **Nodes > Storage Node > Storage**.

The graphs and tables for the node appear.

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

The following values are shown:

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



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



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

Disk Devices					
Name	World Wide Name	I/O Load	Read Rate	Write Rate	
croot(8:1,sda1)	N/A	0.03%	0 bytes/s	3 KB/s	
cvloc(8:2,sda2)	N/A	0.85%	0 bytes/s	58 KB/s	
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	81 bytes/s	
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s	
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s	

Volumes					
Mount Point	Device	Status	Size	Available	Write Cache Status
/	croot	Online	21.00 GB	14.90 GB	Unknown
/var/local	cvloc	Online	85.86 GB	84.10 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	sde	Online	107.32 GB	107.18 GB	Enabled

Object Stores						
ID	Size	Available	Replicated Data	EC Data	Object Data (%)	Health
0000	107.32 GB	96.45 GB	250.90 KB	0 bytes	0.00%	No Errors
0001	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors
0002	107.32 GB	107.18 GB	0 bytes	0 bytes	0.00%	No Errors

4. Monitor the values over time to estimate the rate at which usable storage space is being consumed.

5. To maintain normal system operations, add Storage Nodes, add storage volumes, or archive object data before usable space is consumed.

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



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

For more information 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 information

[Administer StorageGRID](#)

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

[Expand your grid](#)

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.

What you'll need

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

About this task

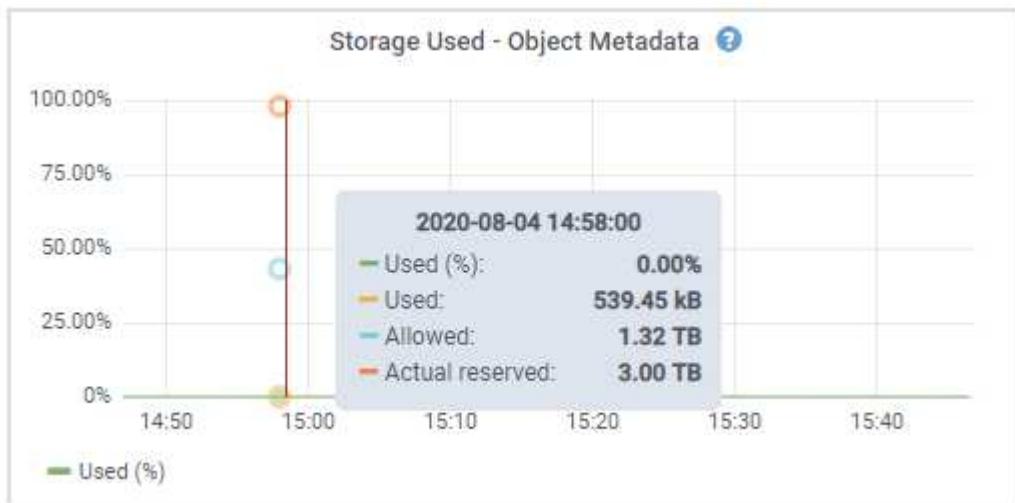
StorageGRID maintains three copies of object metadata at each site to provide redundancy and to protect object metadata from loss. The three copies are evenly distributed across all Storage Nodes at each site using the space reserved for metadata on storage volume 0 of each Storage Node.

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

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

Steps

1. Select **Nodes > Storage Node > Storage**.
2. Hover your cursor over the Storage Used - Object Metadata graph to see the values for a specific time.



Value	Description	Prometheus metric
Used (%)	The percentage of the allowed metadata space that has been used on this Storage Node.	storagegrid_storage_utilization_metadata_bytes/storagegrid_storage_utilization_metadata_allowed_bytes

Value	Description	Prometheus metric
Used	The bytes of the allowed metadata space that have been used on this Storage Node.	storagegrid_storage_utilization_metadata_bytes
Allowed	The space allowed for object metadata on this Storage Node. To learn how this value is determined for each Storage Node, see the instructions for administering StorageGRID.	storagegrid_storage_utilization_metadata_allowed_bytes
Actual reserved	The actual space reserved for metadata on this Storage Node. Includes the allowed space and the required space for essential metadata operations. To learn how this value is calculated for each Storage Node, see the instructions for administering StorageGRID.	storagegrid_storage_utilization_metadata_reserved_bytes



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

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



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

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

Related information

[Troubleshooting the Low metadata storage alert](#)

[Administer StorageGRID](#)

[Expand your grid](#)

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.

What you'll need

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

About this task

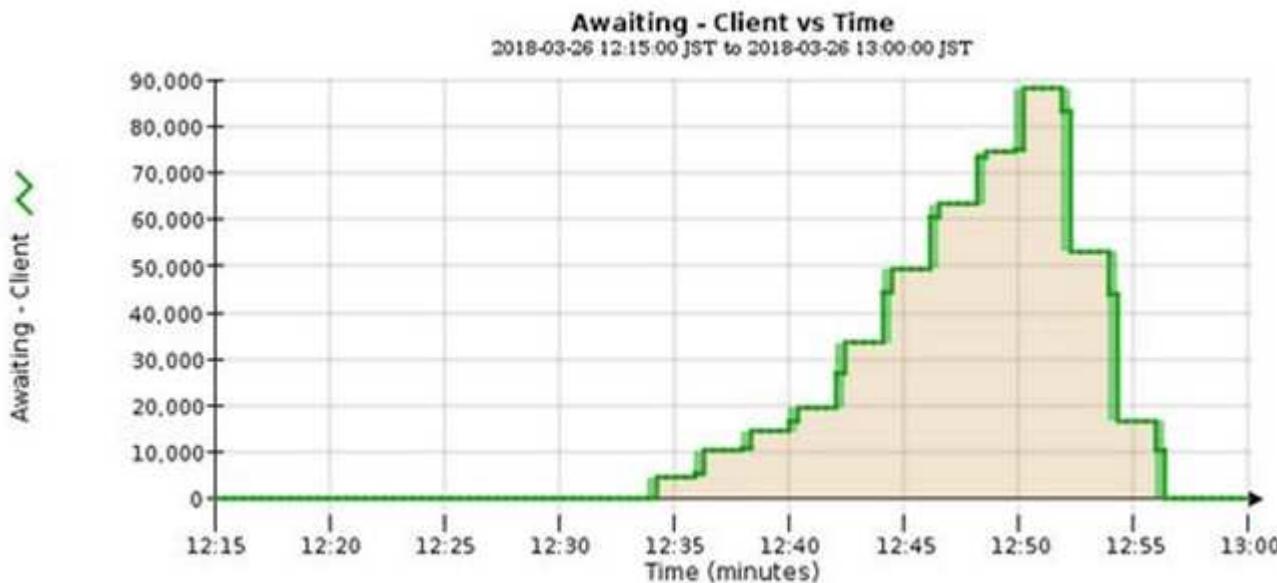
The StorageGRID system manages objects by applying the active ILM policy. The ILM policy and associated ILM rules determine how many copies are made, the type of copies that are created, where copies are placed, and the length of time each copy is retained.

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.



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

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



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

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



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

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

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.

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.

What you'll need

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.

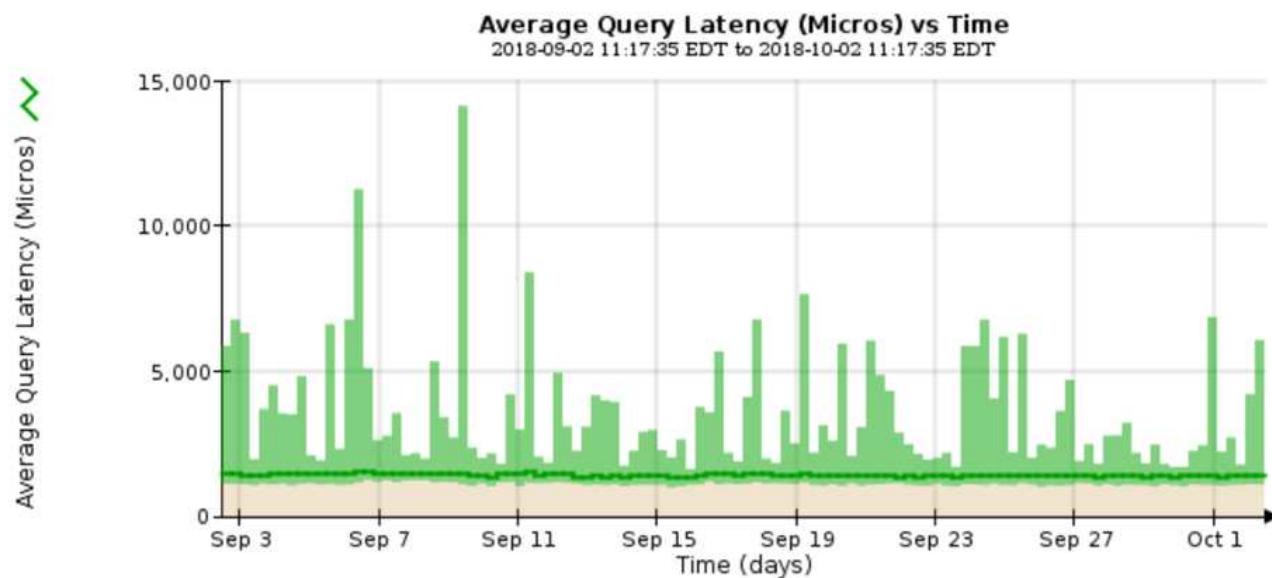
The **High latency for metadata queries** alert is triggered if the average time for queries is too long.

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

[Expand your grid](#)

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.

What you'll need

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

- You must have specific access permissions.

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.

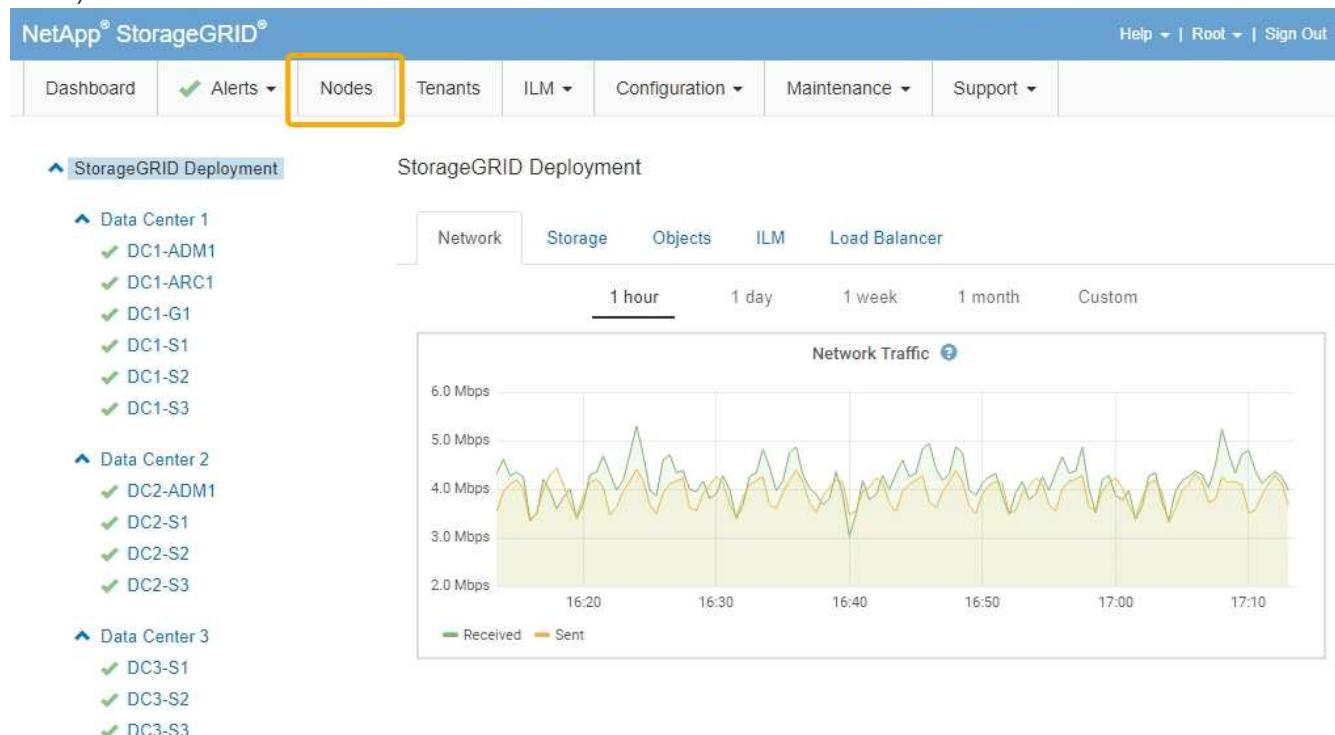
You can use the Grid Manager to monitor connectivity and network performance, so you can address any issues promptly.

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

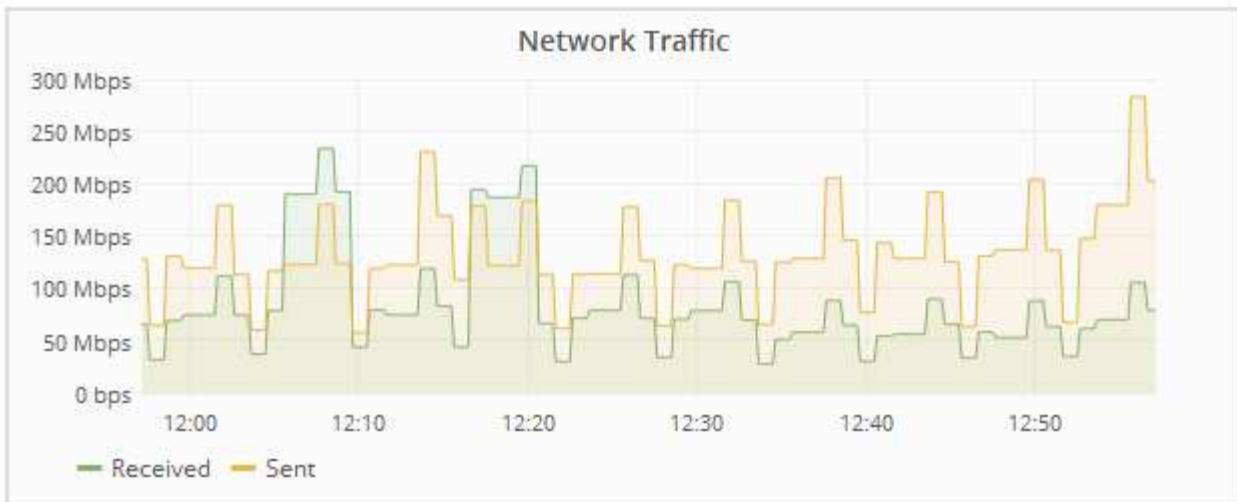
1. Select **Nodes**.

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

- Use the metrics associated with your traffic classification policies to monitor network traffic.
 - Select **Configuration > Network Settings > 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.



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

Displaying 2 traffic classification policies.

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

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

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

Related information

[Viewing the Network tab](#)

[Monitoring node connection states](#)

[Administer StorageGRID](#)

Monitoring node-level resources

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

What you'll need

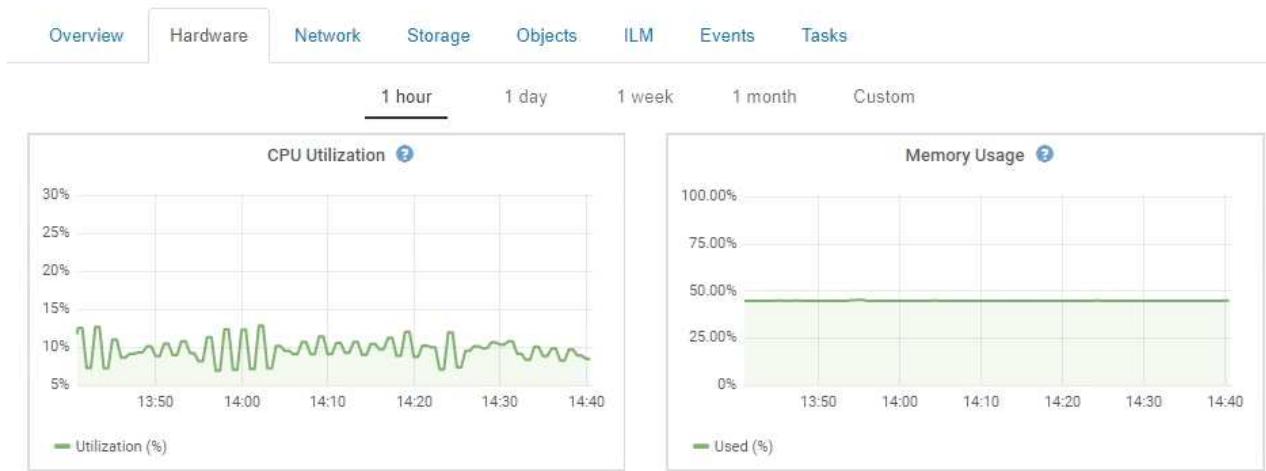
- You must be signed in to the Grid Manager using a supported browser.

About this task

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

Steps

1. 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, or 1 month. 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 information

[Viewing information about appliance Storage Nodes](#)

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

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.

What you'll need

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

About this task



The Space used values are estimates. These estimates are affected by the timing of ingestions, network connectivity, and node status.

Steps

1. Select **Tenants** to review the amount of storage used by all tenants.

The Space Used, Quota Utilization, Quota, and Object Count are listed for each tenant. If a quota is not set for a tenant, the Quota Utilization field contains a dash (--) and the Quota field indicates “Unlimited.”

Tenant Accounts

View information for each tenant account.

Note: Depending on the timing of ingests, network connectivity, and node status, the usage data shown might be out of date. To view more recent values, select the tenant and select **View Details**.



	Display Name	Space Used	Quota Utilization	Quota	Object Count	Sign in
●	Account01	500.00 KB	0.00%	20.00 GB	100	Sign in
○	Account02	2.50 MB	0.01%	30.00 GB	500	Sign in
○	Account03	605.00 MB	4.03%	15.00 GB	31,000	Sign in
○	Account04	1.00 GB	10.00%	10.00 GB	200,000	Sign in
○	Account05	0 bytes	—	Unlimited	0	Sign in

Show 20 rows per page

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. Use the search box to search for a tenant account by display name or tenant ID.

You can sign in to a tenant account by selecting the link in the **Sign in** column of the table.

2. Optionally, select **Export to CSV** to view and export a .csv file containing the usage values for all tenants.

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

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

Tenant ID	Display Name	Space Used (Bytes)	Quota utilization (%)	Quota (Bytes)	Object Count	Protocol
56243391454153665591	Account01	500000	0	20000000000	100	S3
82457136581801590515	Account02	2500000	0.01	30000000000	500	S3
04489086912300179118	Account03	605000000	4.03	15000000000	31000	S3
26417581662098345719	Account04	1000000000	10	10000000000	200000	S3
78472447501213318575	Account05	0				0 S3

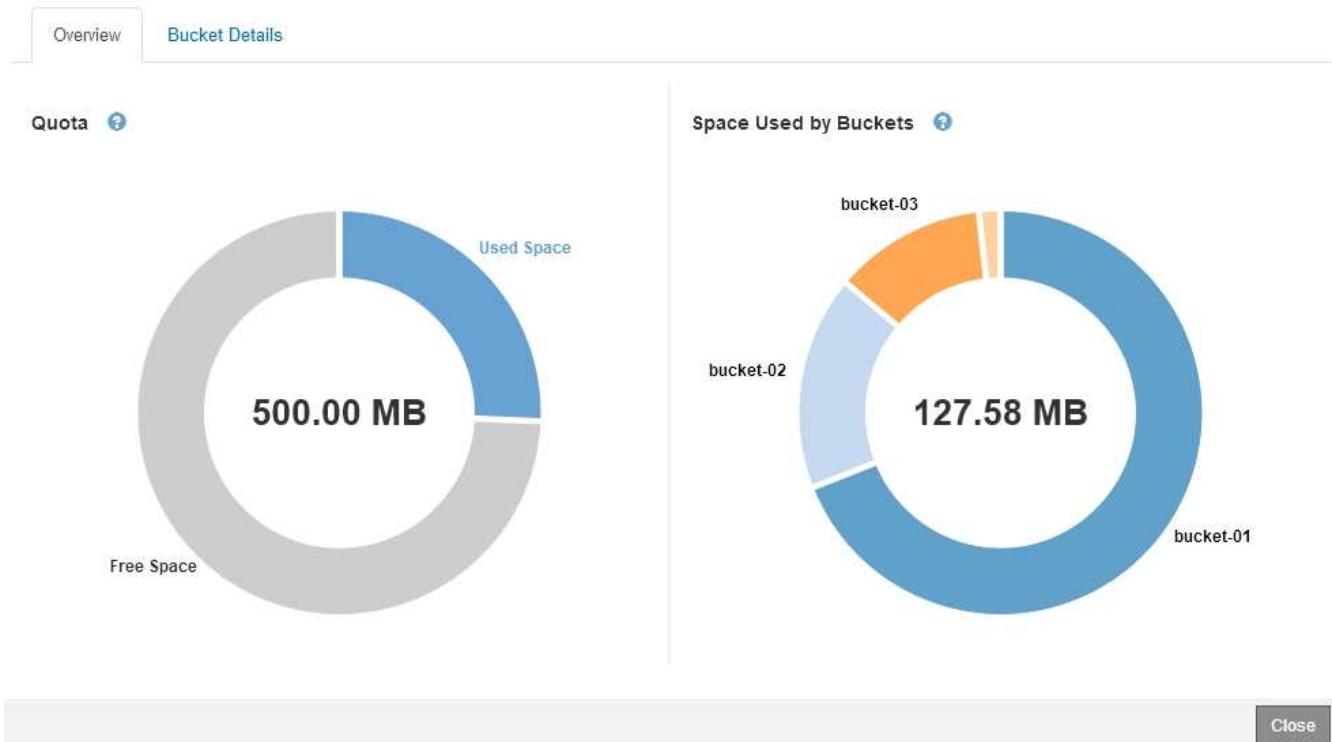
You can open the .csv file in a spreadsheet application or use it in automation.

3. To view details for a specific tenant, including usage charts, select the tenant account from the Tenant Accounts page, and then select **View details**.

The Account Details page appears and shows summary information, a chart that represents the amount of quota used and remaining, and a chart that represents the amount of object data in buckets (S3) or containers (Swift).

Account Details - Account01

Display Name:	Account01	Sign in	Quota Utilization <small>?</small> :	25.52%
Tenant ID:	6479 6966 4290 3892 3647		Logical Space Used <small>?</small> :	127.58 MB
Protocol <small>?</small> :	S3		Quota <small>?</small> :	500.00 MB
Allow Platform Services <small>?</small> :	Yes		Bucket Count <small>?</small> :	5
Uses Own Identity Source <small>?</small> :	No		Object Count <small>?</small> :	30



◦ Quota

If a quota was set for this tenant, the **Quota** chart shows how much of that quota this tenant has used and how much is still available. If no quota was set, the tenant has an unlimited quota, and an informational message is displayed. If the tenant has exceeded the storage quota by more than 1% and by at least 1 GB, the chart shows the total quota and the excess amount.

You can place your cursor over the Used Space segment to see the number of stored objects and the total bytes used. You can place your cursor over the Free Space segment to see how many bytes of storage quota are available.



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



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



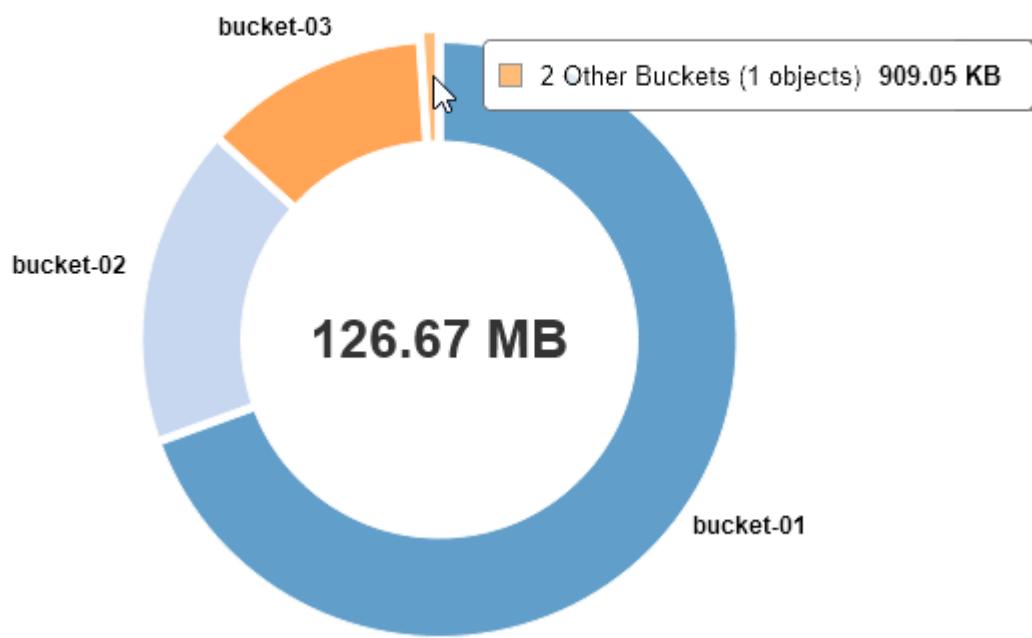
You can enable the **Tenant quota usage high** alert to determine if tenants are consuming their quotas. If enabled, this alert is triggered when a tenant has used 90% of its quota. For more information, see the alerts reference.

- **Space Used**

The **Space Used by Buckets** (S3) or **Space Used by Containers** (Swift) chart shows the largest buckets for the tenant. Space used is the total amount of object data in the bucket. This value does not represent the storage space required for ILM copies and object metadata.

If the tenant has more than nine buckets or containers, they are combined into a segment called Other. Some chart segments might be too small to include a label. You can place your cursor over any of the segments to see the label and obtain more information, including the number of stored objects and total bytes for each bucket or container.

Space Used by Buckets



4. Select **Bucket Details** (S3) or **Container Details** (Swift) to view a list of the space used and number of objects for each of the tenant's buckets or containers.

Account Details - Account01

Display Name:	Account01	Sign in	Quota Utilization <small>?</small> :	84.22%
Tenant ID:	6479 6966 4290 3892 3647		Logical Space Used <small>?</small> :	84.22 MB
Protocol <small>?</small> :	S3		Quota <small>?</small> :	100.00 MB
Allow Platform Services <small>?</small> :	Yes		Bucket Count <small>?</small> :	3
Uses Own Identity Source <small>?</small> :	No		Object Count <small>?</small> :	13

Overview Bucket Details

[Export to CSV](#)

Bucket Name	Space Used	Number of Objects
bucket-01	88.72 MB	14
bucket-02	21.75 MB	11
bucket-03	15.29 MB	3

[Close](#)

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

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

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

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

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

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

- Select **Configuration > Network Settings > Traffic Classification**.

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

Traffic Classification Policies

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

[+ Create](#) [Edit](#) [Remove](#) [Metrics](#)

Name	Description	ID
ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
Fabric Pools	Monitor Fabric Pools	223b0ccb-6968-4646-b32d-7665bddc894b

Displaying 2 traffic classification policies.

- Review the list of policies to identify the ones that apply to a specific tenant.
- To view metrics associated with a policy, select the radio button to the left of the policy, and then click **Metrics**.
- 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.

7. 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.

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

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



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

- 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 information

[Alerts reference](#)

[Review audit logs](#)

[Administer StorageGRID](#)

[Reviewing support metrics](#)

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.

What you'll need

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

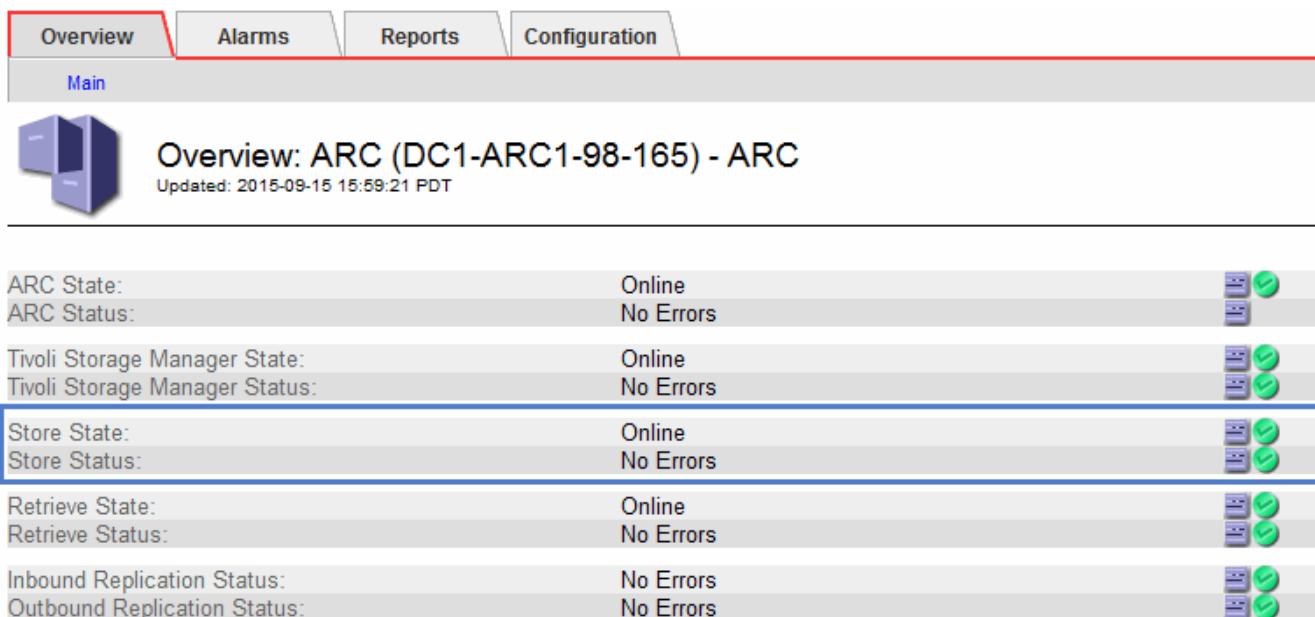
About this task

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



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

[Administer 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.

What you'll need

- 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.



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.
- h. Optionally, use traffic classification policies to view a more detailed breakdown of traffic being served by the Load Balancer service.

2. If S3 or Swift clients connect using the CLB service (deprecated), 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 > Tools > 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 information

[Administer StorageGRID](#)

[Viewing the Load Balancer tab](#)

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 the End User License Agreement, select the check box, and then select **Accept & Continue**.

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

[Upgrade software](#)

[Maintain & recover](#)

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. This section contains the following sub-sections:

- [Comparing alerts and alarms](#)
- [Managing alerts](#)
- [Managing alarms \(legacy system\)](#)

StorageGRID includes two systems for informing you about issues.

Alert system

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

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

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

Legacy alarm system

The alarm system is supported, 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 > Alarms (legacy) > Current Alarms** page.
- An email notification is sent, assuming you have configured an SMTP server and configured one or more mailing lists.
- An SNMP notification might be sent, assuming you have configured the StorageGRID SNMP agent. (SNMP notifications 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.

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

	Alerts	Alarms (legacy system)
How do I see which alerts or alarms are active?	<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>Viewing current alerts</p>	<ul style="list-style-type: none"> Click the Legacy alarms link on the Dashboard. Select Support > Alarms (legacy) > Current Alarms. <p>Viewing legacy alarms</p>
What causes an alert or an alert to be triggered?	<p>Alerts are triggered when a Prometheus expression in an alert rule evaluates as true for the specific trigger condition and duration.</p> <p>Viewing alert rules</p>	<p>Alarms are triggered when a StorageGRID attribute reaches a threshold value.</p> <p>Alarm triggering logic (legacy system)</p>
If an alert or alarm is triggered, how do I resolve the underlying problem?	<p>The recommended actions for an alert are included in email notifications and are available from the Alerts pages in the Grid Manager.</p> <p>As required, additional information is provided in the StorageGRID documentation.</p> <p>Alerts reference</p>	<p>You can learn about an alarm by clicking the attribute name, or you can search for an alarm code in the StorageGRID documentation.</p> <p>Alarms reference (legacy system)</p>
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. <p>Viewing resolved alerts</p>	<p>Select Support > Alarms (legacy) > Historical Alarms.</p> <p>Reviewing historical alarms and alarm frequency (legacy system)</p>
Where do I manage the settings?	<p>Select Alerts. Then, use the options in the Alerts menu.</p> <p>Managing alerts</p>	<p>Select Support. Then, use the options in the Alarms (legacy) section of the menu.</p> <p>Managing alarms (legacy system)</p>

	Alerts	Alarms (legacy system)
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. <p>Administer StorageGRID</p>	<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. <p>Administer StorageGRID</p>
How do I manage email notifications?	<p>Select Alerts > Email Setup.</p> <p>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.</p> <p>Managing alert notifications</p>	<p>Select Support > Alarms (legacy) > Legacy Email Setup.</p> <p>Configuring notifications for alarms (legacy system)</p>
How do I manage SNMP notifications?	<p>Select Configuration > Monitoring > SNMP Agent. Using SNMP monitoring</p>	<p>Select Configuration > Monitoring > SNMP Agent. Using SNMP monitoring</p> <p>Note: SNMP notifications are not sent for every alarm or alarm severity.</p> <p>Alarms that generate SNMP notifications (legacy system)</p>

	Alerts	Alarms (legacy system)
How do I control who receives notifications?	<p>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> <p>Setting up email notifications for alerts</p>	<p>1. Select Support > Alarms (legacy) > Legacy Email Setup. 2. Creating a mailing list. 3. Select Notifications. 4. Select the mailing list.</p> <p>Creating mailing lists for alarm notifications (legacy system)</p> <p>Configuring email notifications for alarms (legacy system)</p>
Which Admin Nodes send notifications?	<p>A single Admin Node (the “preferred sender”).</p> <p>Administer StorageGRID</p>	<p>A single Admin Node (the “preferred sender”).</p> <p>Administer StorageGRID</p>
How do I suppress some notifications?	<p>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> <p>Note: If you have enabled the SNMP agent, silences also suppress SNMP traps and informs.</p> <p>Silencing alert notifications</p>	<p>1. Select Support > Alarms (legacy) > Legacy Email Setup. 2. Select Notifications. 3. Select a mailing list, and select Silence.</p> <p>Suppressing alarm notifications for a mailing list (legacy system)</p>

	Alerts	Alarms (legacy system)
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>Silencing alert notifications</p>	<ol style="list-style-type: none"> 1. Select Configuration > System Settings > 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>Suppressing email notifications system wide</p>
How do I customize the conditions and triggers?	<ol style="list-style-type: none"> 1. Select Alerts > Alert Rules. 2. Select a default rule to edit, or select Create custom rule. <p>Editing an alert rule</p> <p>Creating custom alert rules</p>	<ol style="list-style-type: none"> 1. Select Support > Alarms (legacy) > Global Alarms. 2. Create a Global Custom alarm to override a Default alarm or to monitor an attribute that does not have a Default alarm. <p>Creating Global Custom alarms (legacy system)</p>
How do I disable an individual alert or alarm?	<ol style="list-style-type: none"> 1. Select Alerts > Alert Rules. 2. Select the rule, and click Edit rule. 3. Unselect the Enabled check box. <p>Disabling an alert rule</p>	<ol style="list-style-type: none"> 1. Select Support > Alarms (legacy) > Global Alarms. 2. Select the rule, and click the Edit icon. 3. Unselect the Enabled check box. <p>Disabling a Default alarm (legacy system)</p> <p>Disabling Global Custom alarms (legacy system)</p>

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 information

[Viewing current alerts](#)

[Viewing resolved alerts](#)

Viewing a specific alert

Alerts reference

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.

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.

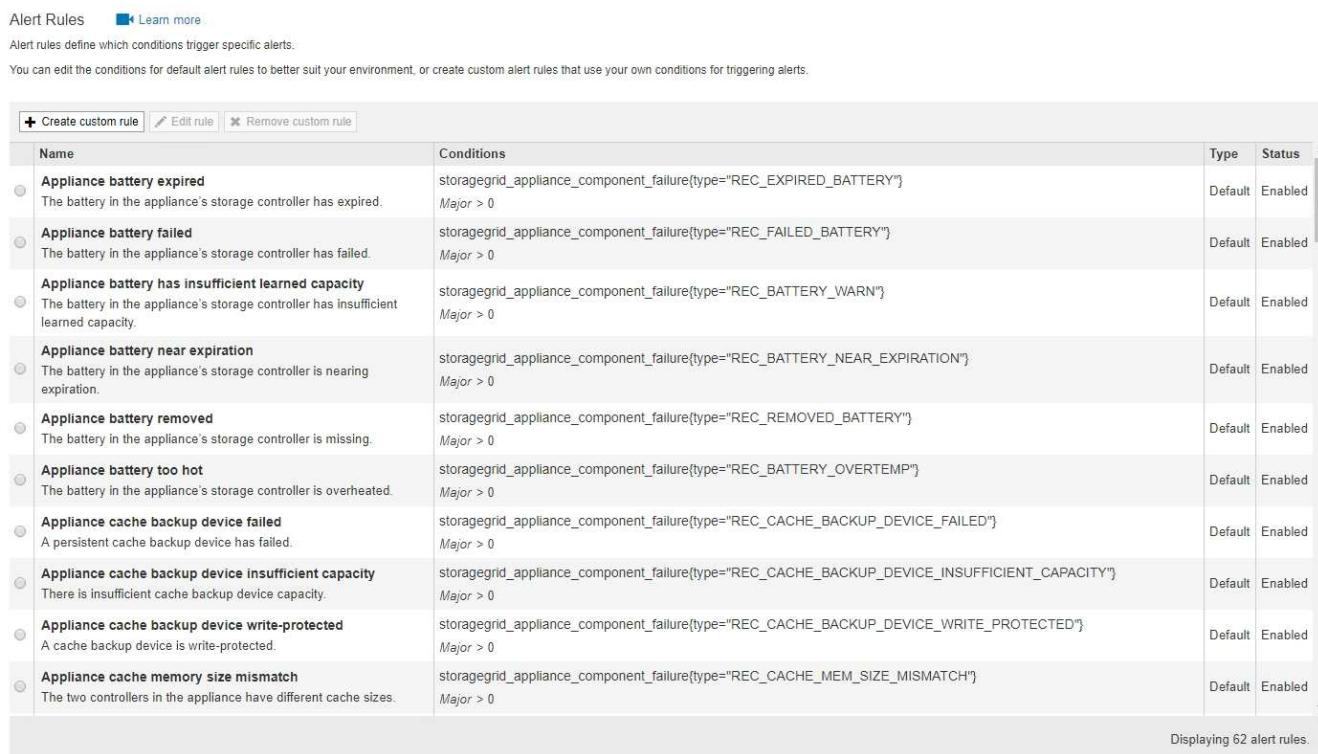
What you'll need

- 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 > Alert Rules**.

The Alert Rules page appears.



Alert Rules		Learn more	
Alert rules define which conditions trigger specific alerts.			
You can edit the conditions for default alert rules to better suit your environment, or create custom alert rules that use your own conditions for triggering alerts.			
+ Create custom rule	Edit rule	Remove custom rule	
Name	Conditions	Type	Status
Appliance battery expired The battery in the appliance's storage controller has expired.	storagegrid_appliance_component_failure(type="REC_EXPIRED_BATTERY") Major > 0	Default	Enabled
Appliance battery failed The battery in the appliance's storage controller has failed.	storagegrid_appliance_component_failure(type="REC_FAILED_BATTERY") Major > 0	Default	Enabled
Appliance battery has insufficient learned capacity The battery in the appliance's storage controller has insufficient learned capacity.	storagegrid_appliance_component_failure(type="REC_BATTERY_WARN") Major > 0	Default	Enabled
Appliance battery near expiration The battery in the appliance's storage controller is nearing expiration.	storagegrid_appliance_component_failure(type="REC_BATTERY_NEAR_EXPIRATION") Major > 0	Default	Enabled
Appliance battery removed The battery in the appliance's storage controller is missing.	storagegrid_appliance_component_failure(type="REC_REMOVED_BATTERY") Major > 0	Default	Enabled
Appliance battery too hot The battery in the appliance's storage controller is overheated.	storagegrid_appliance_component_failure(type="REC_BATTERY_OVERTEMP") Major > 0	Default	Enabled
Appliance cache backup device failed A persistent cache backup device has failed.	storagegrid_appliance_component_failure(type="REC_CACHE_BACKUP_DEVICE_FAILED") Major > 0	Default	Enabled
Appliance cache backup device insufficient capacity There is insufficient cache backup device capacity.	storagegrid_appliance_component_failure(type="REC_CACHE_BACKUP_DEVICE_INSUFFICIENT_CAPACITY") Major > 0	Default	Enabled
Appliance cache backup device write-protected A cache backup device is write-protected.	storagegrid_appliance_component_failure(type="REC_CACHE_BACKUP_DEVICE_WRITE_PROTECTED") Major > 0	Default	Enabled
Appliance cache memory size mismatch The two controllers in the appliance have different cache sizes.	storagegrid_appliance_component_failure(type="REC_CACHE_MEM_SIZE_MISMATCH") Major > 0	Default	Enabled

Displaying 62 alert rules

2. Review the information in the alert rules table:

Column header	Description
Name	The unique name and description of the alert rule. Custom alert rules are listed first, followed by default alert rules. The alert rule name is the subject for email notifications.

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 information

[Alerts reference](#)

Creating custom alert rules

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

What you'll need

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

- Do not assume a custom alert is working unless you have validated that the alert is triggered when expected.

Steps

1. Select **Alerts > Alert 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

 minutes

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.

Field	Description
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.

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

A basic expression is usually of the form:

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

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

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
```

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

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

The default is 5 minutes.

7. Click **Save**.

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

Related information

[Administer StorageGRID](#)

[Commonly used Prometheus metrics](#)

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.

What you'll need

- 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.



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 > Alert 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	<input type="text"/>
Major	<input type="text" value="node_memory_MemTotal_bytes < 24000000000"/>
Critical	<input type="text" value="node_memory_MemTotal_bytes <= 12000000000"/>

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

Duration ▾

[Cancel](#) [Save](#)

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.



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



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

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



You 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.

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

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



If you want to restore a condition for an edited default alert rule back to its original value, 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 < 24000000000"/>
Critical	<input type="text" value="node_memory_MemTotal_bytes <= 14000000000"/> 



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 information

[Administer StorageGRID](#)

[Commonly used Prometheus metrics](#)

[Prometheus: Query basics](#)

Disabling an alert rule

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

What you'll need

- 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.



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 > Alert 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.



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.

What you'll need

- 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 > Alert 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).

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

To learn how to configure the SNMP agent, see 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.



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 information

[Using SNMP monitoring](#)

[Silencing alert notifications](#)

[Administer StorageGRID](#)

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

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.

What you'll need

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

What you'll need

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 ? 10.224.1.250

Port ? 25

Username (optional) ? smtpuser

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 	storagegrid-alerts@example.com
Recipient 1 	recipient1@example.com 
Recipient 2 	recipient2@example.com  

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.



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-----  
1234567890abcdefghijklmnopqrstuvwxyz  
ABCDEFGHIJKLMNPQRSTUVWXYZ1234567890  
-----END CERTIFICATE-----
```

[Browse](#)

Send Client Certificate

Client Certificate

```
-----BEGIN CERTIFICATE-----  
1234567890abcdefghijklmnopqrstuvwxyz  
ABCDEFGHIJKLMNPQRSTUVWXYZ1234567890  
-----END CERTIFICATE-----
```

[Browse](#)

Private Key

```
-----BEGIN PRIVATE KEY-----  
1234567890abcdefghijklmnopqrstuvwxyz  
ABCDEFGHIJKLMNPQRSTUVWXYZ1234567890  
-----BEGIN PRIVATE KEY-----
```

[Browse](#)

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

Severity	Description
Minor, major, critical	An email notification is sent when the minor, major, or critical condition for an alert rule is met.
Major, critical	An email notification is sent when the major or critical condition for an alert rule is met. Notifications 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

Send Test Email

Save

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.



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

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



When you test alert notifications, you must sign in to every Admin Node to verify connectivity. This is in contrast to testing 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 information

[Troubleshooting alert email notifications](#)

[Maintain & recover](#)

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:

Low object data storage (6 alerts) 1The 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

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

Related information[Silencing 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.
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"> Alert A is triggered for node 1. A notification is sent. Alert A is triggered for node 2. A second notification is sent. Alert A is resolved for node 2, but it remains active for node 1. 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"> Alert A is triggered for node 1 on March 8. A notification is sent. 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.

What you'll need

- 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 information

[Collecting log files and system data](#)

Silencing alert notifications

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

What you'll need

- 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.



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



Because alarms and alerts are independent systems, you 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
<i>No results found.</i>								

2. Select **Create**.

The Create Silence dialog box appears.

Create Silence

Alert Rule:

Description (optional):

Duration: Minutes ▾

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.

Field	Description
Duration	<p>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).</p> <p>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.</p>
Severity	<p>Which alert severity or severities should be silenced. If the alert is triggered at one of the selected severities, no notifications are sent.</p>
Nodes	<p>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.</p> <p>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.</p>

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"> Select Alerts > Silences. From the table, select the radio button for the silence you want to edit. Click Edit. Change the description, the amount of time remaining, the selected severities, or the affected node. Click Save.
Remove a silence	<ol style="list-style-type: none"> Select Alerts > Silences. From the table, select the radio button for the silence you want to remove. Click Remove. 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 information

Configuring the SNMP agent

Managing alarms (legacy system)

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



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

Related information

[Alarms reference \(legacy system\)](#)

[Viewing legacy alarms](#)

[Administer StorageGRID](#)

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 information

[Viewing Default alarms \(legacy system\)](#)

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

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

[Disabling Global Custom alarms \(legacy system\)](#)

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.

Icon	Color	Alarm severity	Meaning
	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.

	Global Custom alarm threshold (enabled)	Default alarm threshold (enabled)
Notice	≥ 1500	≥ 1000
Minor	$\geq 15,000$	≥ 1000
Major	$\geq 150,000$	$\geq 250,000$

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	under10	<=	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 shows two critical Entries Available (VMFI) alarms. (Select **Support > Tools > Grid Topology**. Then, select **Storage Node > SSM > Resources**.)

You can disable the VMFI alarm so that the Critical level VMFI alarm is not triggered (both currently Critical alarms would appear in the table as green); however, you 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)

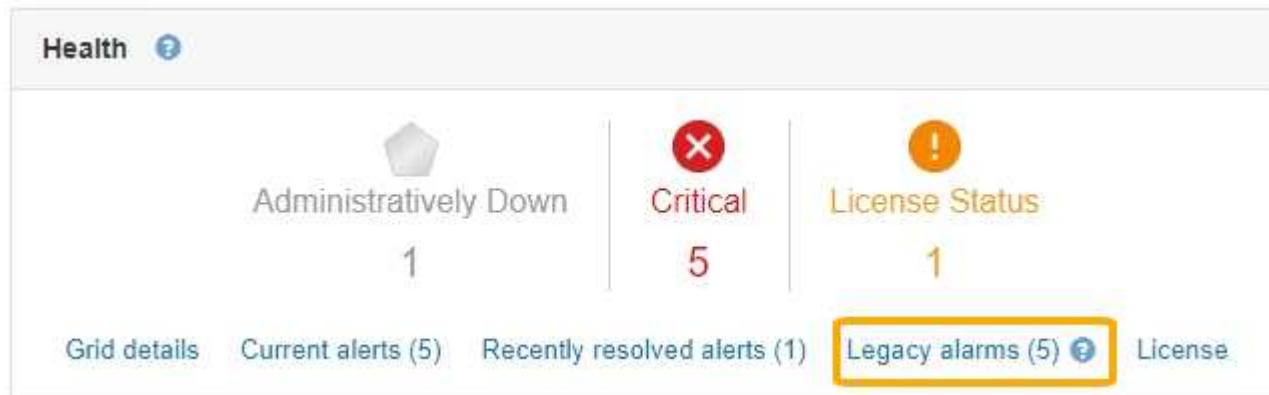
Legacy alarms are triggered when system attributes reach alarm threshold values. If you want to reduce or clear the count of legacy alarms on the Dashboard, you can acknowledge the alarms.

What you'll need

- 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, 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.



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.



While the legacy alarm system continues to be supported, the 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 > Alarms (legacy) > 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

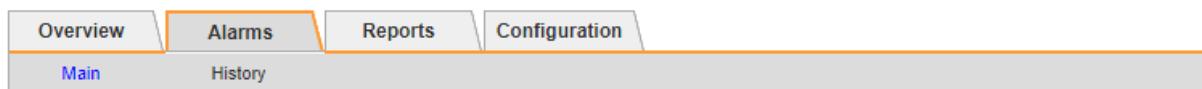
Show Acknowledged Alarms (1 - 1 of 1)

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 > Tools > Grid Topology > Grid Node > Service > Alarms**).



Alarms: ARC (DC1-ARC1) - Replication

Updated: 2019-05-24 10:46:48 MDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledged
 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.



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 > Alarms (legacy) > 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

<input checked="" type="checkbox"/> 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 information

[Alarms reference \(legacy system\)](#)

Viewing Default alarms (legacy system)

You can view the list of all Default legacy alarms.

What you'll need

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



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

Steps

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

All Default alarms are listed.



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"/>		XCVP (% 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.

What you'll need

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



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

Steps

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

2. Follow these steps to find out how often alarms have been triggered for a particular attribute.

- a. Select **Support > Tools > Grid Topology**.
- b. Select **grid node > service or component > Alarms > History**.
- c. Select the attribute from the list.
- d. Do one of the following:
 - Click one of the time periods.
 - Enter a custom range, and click **Custom Query**.

The alarms are listed in reverse chronological order.

- e. To return to the alarms history request form, click **History**.

Related information

[Alarms reference \(legacy system\)](#)

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.

What you'll need

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



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

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



Be 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 > Alarms (legacy) > Global Alarms**.

2. Add a new row to the Global Custom alarms table:

- To add a new alarm, click **Edit**  (if this is the first entry) or **Insert** .



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	
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)		Notice	At least	>=	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	
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)		Notice	At least	>=	3000	
<input checked="" type="checkbox"/>	ARC	ARRF (Request Failures)		Major	At least	>=	1	
<input checked="" type="checkbox"/>	ARC	ARRV (Verification Failures)		Major	At least	>=	1	
<input checked="" type="checkbox"/>	ARC	ARVF (Store Failures)		Major	At least	>=	1	
<input checked="" type="checkbox"/>	NMS	ARRC (Remaining Capacity)		Notice	Below	<=	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.
 - Under Filter by, select either **Attribute Code** or **Attribute Name**.
 - 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.

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

- 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.

Heading	Description
Attribute	<p>Select the name and code of the attribute being monitored from the list of all attributes applicable to the selected service or component.</p> <p>To display information about the attribute, click Info  next to the attribute's name.</p>
Severity	The icon and text indicating the level of the alarm.
Message	The reason for the alarm (connection lost, storage space below 10%, and so on).
Operator	<p>Operators for testing the current attribute value against the Value threshold:</p> <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	<p>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> <p>Note: Mailing lists require SMTP server setup in order to operate. Before adding mailing lists, confirm that SMTP is configured. Notifications for Custom alarms can override notifications from Global Custom or Default alarms.</p>
Actions	<p>Control buttons to:</p> <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 information

Configuring email server settings for alarms (legacy system)

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.



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

Disabling a Default alarm (legacy system)

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

What you'll need

- 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.



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

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

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



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

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



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

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.

What you'll need

- 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 > Alarms (legacy) > 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.



Global Custom Alarms (1 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input type="checkbox"/>	All	RDTE (Tivoli Storage Manager State)		Major	Offline	=	10	

Default Alarms

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
0 Result(s)							

Apply Changes

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.

What you'll need

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

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

1. Disable the alarm.

2. Log in to the primary Admin Node:

- Enter the following command: `ssh admin@primary_Admin_Node_IP`
- Enter the password listed in the `Passwords.txt` file.
- Enter the following command to switch to root: `su -`
- 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 information

[Disabling alarms \(legacy system\)](#)

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 information

[Using SNMP monitoring](#)

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 ones of the following service states:

- Unknown
- Administratively Down

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

Related information

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

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.

What you'll need

- 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.



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

SMTP is the only protocol supported for sending email.

Steps

1. Select **Support > Alarms (legacy) > Legacy Email Setup**.
2. From the Email menu, select **Server**.

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

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: 2016-03-17 11:11:59 PDT

E-mail Server (SMTP) Information

Mail Server	
Port	
Authentication	Off
Authentication Credentials	Username: root Password: *****
From Address	
Test E-mail	To: ***** <input checked="" 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.

Item	Description
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.

- 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.
- Optionally, send a test email to confirm that your SMTP mail server settings are correct.

- 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-delimited 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.

- Select **Send Test E-mail**.

- 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

[Administer 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.

What you'll need

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

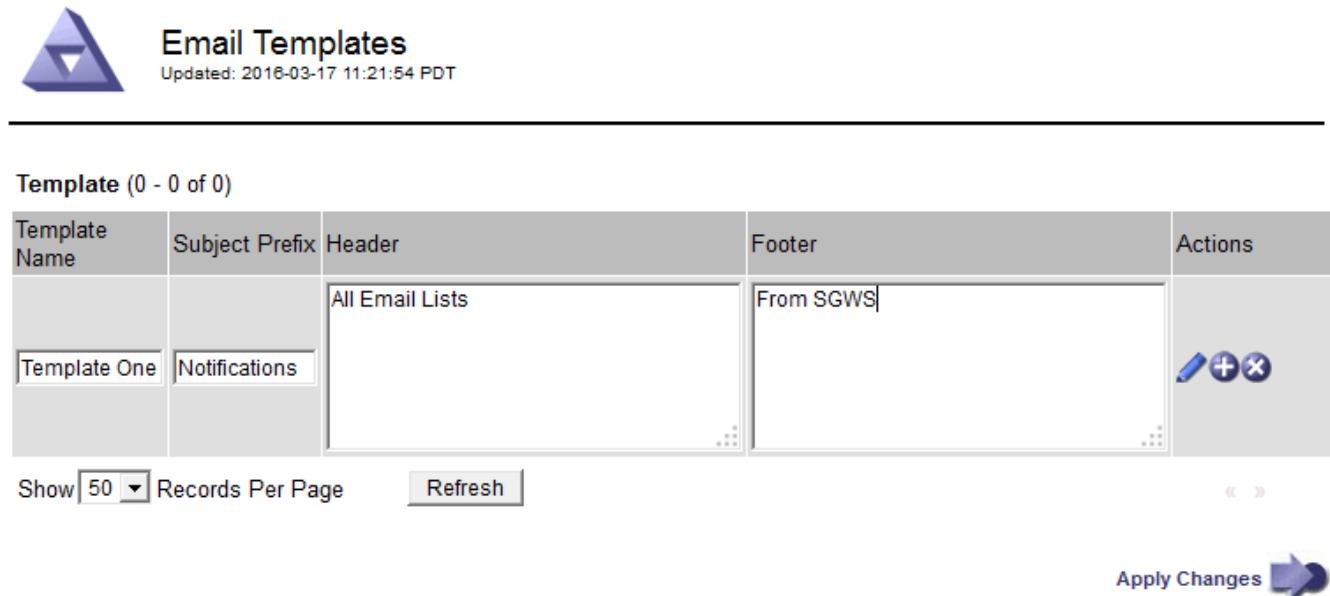
About this task

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



Template (0 - 0 of 0)

Template Name	Subject Prefix	Header	Footer	Actions
Template One	Notifications	All Email Lists	From SGWS	  

Show 50 Records Per Page  

Apply Changes 

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.

What you'll need

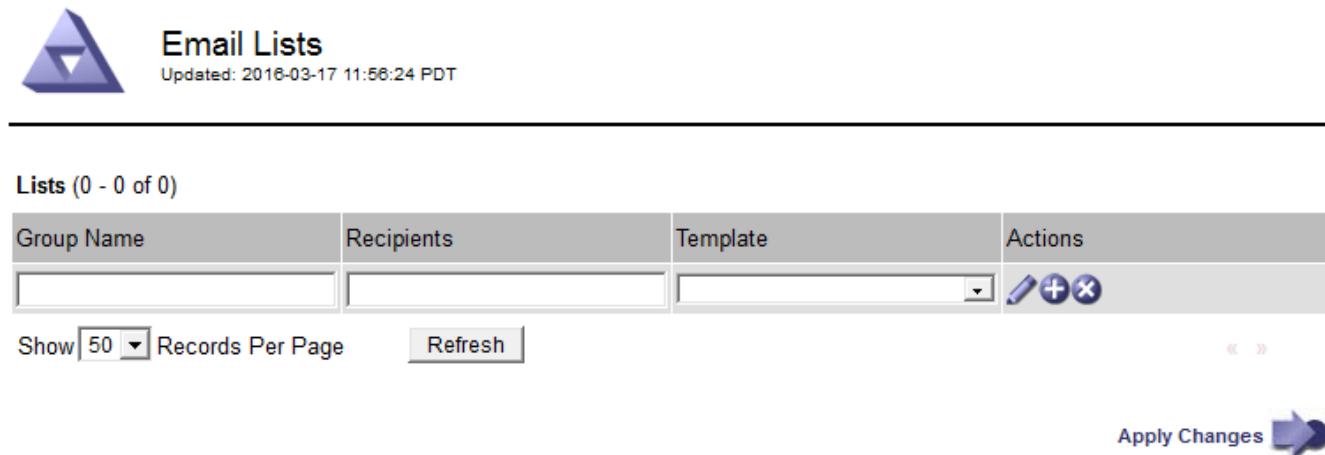
- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- If you want to specify an email template for the mailing list (custom header, footer, and subject line), you must have already created the template.

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



Email Lists
Updated: 2016-03-17 11:56:24 PDT

Lists (0 - 0 of 0)

Group Name	Recipients	Template	Actions
<input type="text"/>	<input type="text"/>	<input type="button" value="▼"/>	  

Show Records Per Page  

Apply Changes 

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.

Item	Description
Recipients	<p>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.</p> <p>Note: If an email address belongs to multiple mailing lists, only one email notification is sent when a notification triggering event occurs.</p>
Template	<p>Optionally, select an email template to add a unique header, footer, and subject line to notifications sent to all recipients of this mailing list.</p>

5. Click **Apply Changes**.

A new mailing list is created.

Related information

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

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.

What you'll need

- 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 > Alarms (legacy) > Legacy Email Setup**.
2. From the Email menu, select **Notifications**.
3. Click **Edit**  (or **Insert** 

125

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 information

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

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

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.

What you'll need

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

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



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

Steps

1. Select **Support > Alarms (legacy) > Legacy Email Setup**.
2. From the Email menu, select **Notifications**.
3. Click **Edit** next to the mailing list for which you want to suppress notifications.
4. Under **Suppress**, select the 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 email notifications for legacy alarms and event-triggered AutoSupport messages.

What you'll need

- 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 email notifications for legacy alarms and event-triggered AutoSupport messages.



This option does not suppress alert email notifications. It also does not suppress weekly or user-triggered AutoSupport messages.

Steps

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

The screenshot shows the 'Display Options' page with the following details:

- Current Sender:** ADMIN-DC1-ADM1
- Preferred Sender:** ADMIN-DC1-ADM1
- GUI Inactivity Timeout:** 900
- Notification Suppress All:**

At the bottom right is a blue 'Apply Changes' button with a right-pointing arrow icon.

4. Click **Apply Changes**.

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

The screenshot shows the 'Notifications' page with the following details:

Notifications (0 - 0 of 0)

All e-mail notifications are now suppressed.

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

Show 50 Records Per Page

At the bottom right is a blue 'Apply Changes' button with a right-pointing arrow icon.

Related information

[Administer 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.

- [Configuring the SNMP agent](#)
- [Updating the SNMP agent](#)

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. The MIB also contains system description information such as platform and model number for each node. Each StorageGRID node also supports a subset of MIB-II objects.

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.

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

Informs are supported in SNMPv2c and SNMPv3.

Trap and inform notifications are sent in the following cases:

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



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

	SNMPv1	SNMPv2c	SNMPv3
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 information

[Alerts reference](#)

[Alarms reference \(legacy system\)](#)

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

[Silencing alert notifications](#)

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.

What you'll need

- 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 > Monitoring > 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

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

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

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

The System Location can be any information that is useful for identifying where your StorageGRID system

is located. For example, you might use the street address of a facility. The value you provide applies to all nodes in the StorageGRID system. **System Location** can be a maximum of 255 characters.

5. Keep the **Enable SNMP Agent Notifications** 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.

6. 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.
7. If you use SNMPv1 or SNMPv2c, complete the Community Strings section.

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

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

As required, you can provide a different (“custom”) community string when you [define a specific trap destination](#).

Default Trap Community can be a maximum of 32 characters and 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.



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.

Create Agent Address

Internet Protocol IPv4 IPv6

Transport Protocol UDP TCP

StorageGRID Network

Port

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

By default, SNMP uses IPv4.

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

By default, SNMP uses UDP.

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

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



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

e. In the **Port** field, optionally enter the port number that the SNMP agent should listen on.

The default UDP port for an SNMP agent is 161, but you can enter any unused port number.



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

f. Click **Create**.

The agent address is created and added to the table.

Other Configurations

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

USM Users (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.



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 SHA

Password

Confirm Password

Privacy

Protocol AES

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.



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.



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.



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

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



The only privacy protocol supported is AES.

h. 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 checked="" type="radio"/>	user2	✓	authNoPriv	
<input checked="" 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 162

Protocol UDP TCP

Community String Use the default trap community: No default found
(Specify the default on the SNMP Agent page.) Use a custom community string

Custom Community String

Cancel **Create**

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> <ol style="list-style-type: none">In the Host field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap.For Port, use the default (162), unless you must use another value. (162 is the standard port for SNMP traps.)For Protocol, use the default (UDP). TCP is also supported. (UDP is the standard SNMP trap protocol.)Use the default trap community, if one was specified on the SNMP Agent page, or enter a custom community string for this trap destination. <p>The custom community string can be a maximum of 32 characters and cannot contain whitespace.</p>

Version	Specify this information
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. iii. For Port, use the default (162), unless you must use another value. (162 is the standard port for SNMP traps.) iv. For Protocol, use the default (UDP). TCP is also supported. (UDP is the standard SNMP trap protocol.) v. Use the default trap community, if one was specified on the SNMP Agent page, or enter a custom community string for this trap destination. <p>The custom community string can be a maximum of 32 characters and cannot contain whitespace.</p>
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. iii. For Port, use the default (162), unless you must use another value. (162 is the standard port for SNMP traps.) iv. For Protocol, use the default (UDP). TCP is also supported. (UDP is the standard SNMP trap protocol.) v. 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.

Other Configurations



Version	Type	Host	Port	Protocol	Community/USM User
SNMPv3	Trap	local		UDP	User: Read only user
SNMPv3	Inform	10.10.10.10	162	UDP	User: Inform user

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

The new SNMP agent configuration becomes active.

Related information

[Silencing 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.

What you'll need

- 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 > Monitoring > 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.



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.

Other Configurations

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

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

USM Users (3)				
+ 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

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.



If you add or remove a user's authoritative engine ID and that user is currently selected for a destination, you must edit or remove the destination, as described in step [SNMP trap destination](#). Otherwise, a validation error occurs when you save the SNMP agent configuration.

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



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



422: Unprocessable Entity

Validation failed. Please check the values you entered for errors.

Undefined trap destination usmUser 'user1'

OK

- 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.

Other Configurations

Agent Addresses (1)	USM Users (2)	Trap Destinations (2)			
		+ Create >Edit Remove			
Version	Type	Host	Port	Protocol	Community/USM User
SNMPv3	Trap	local		UDP	User: Read only user
SNMPv3	Inform	10.10.10.10	162	UDP	User: Inform user

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 information

[Configuring the SNMP agent](#)

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.

- [Using charts and reports](#)
- [Monitoring PUT and GET performance](#)
- [Monitoring object verification operations](#)
- [Monitoring events](#)
- [Reviewing audit messages](#)
- [Collecting log files and system data](#)
- [Manually triggering an AutoSupport message](#)
- [Viewing the Grid Topology tree](#)

- Reviewing support metrics
- Running diagnostics
- Creating custom monitoring applications

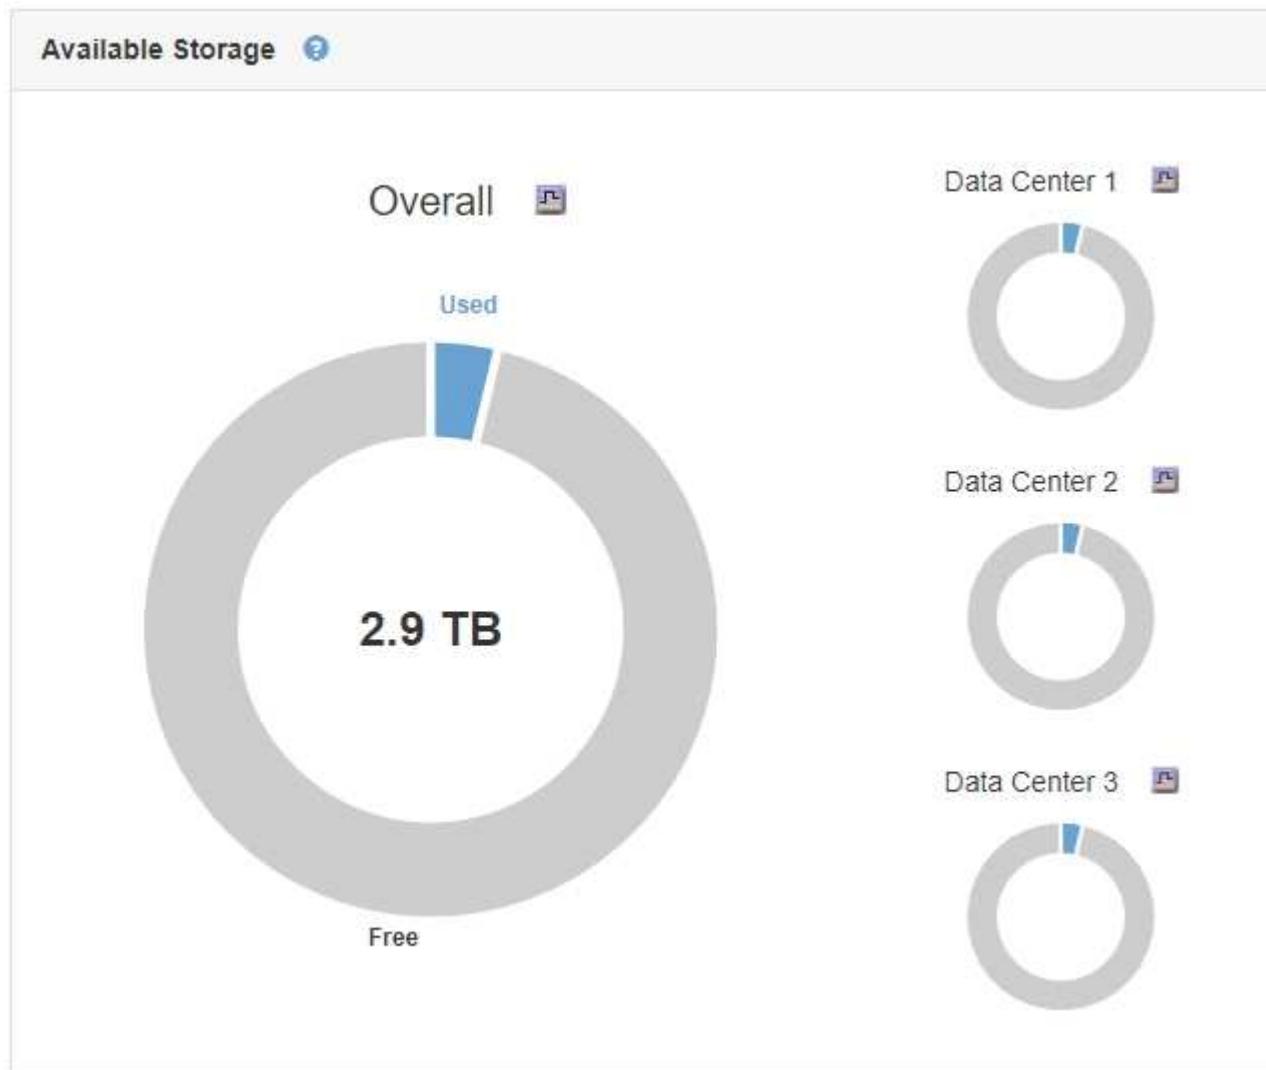
Using charts and reports

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.

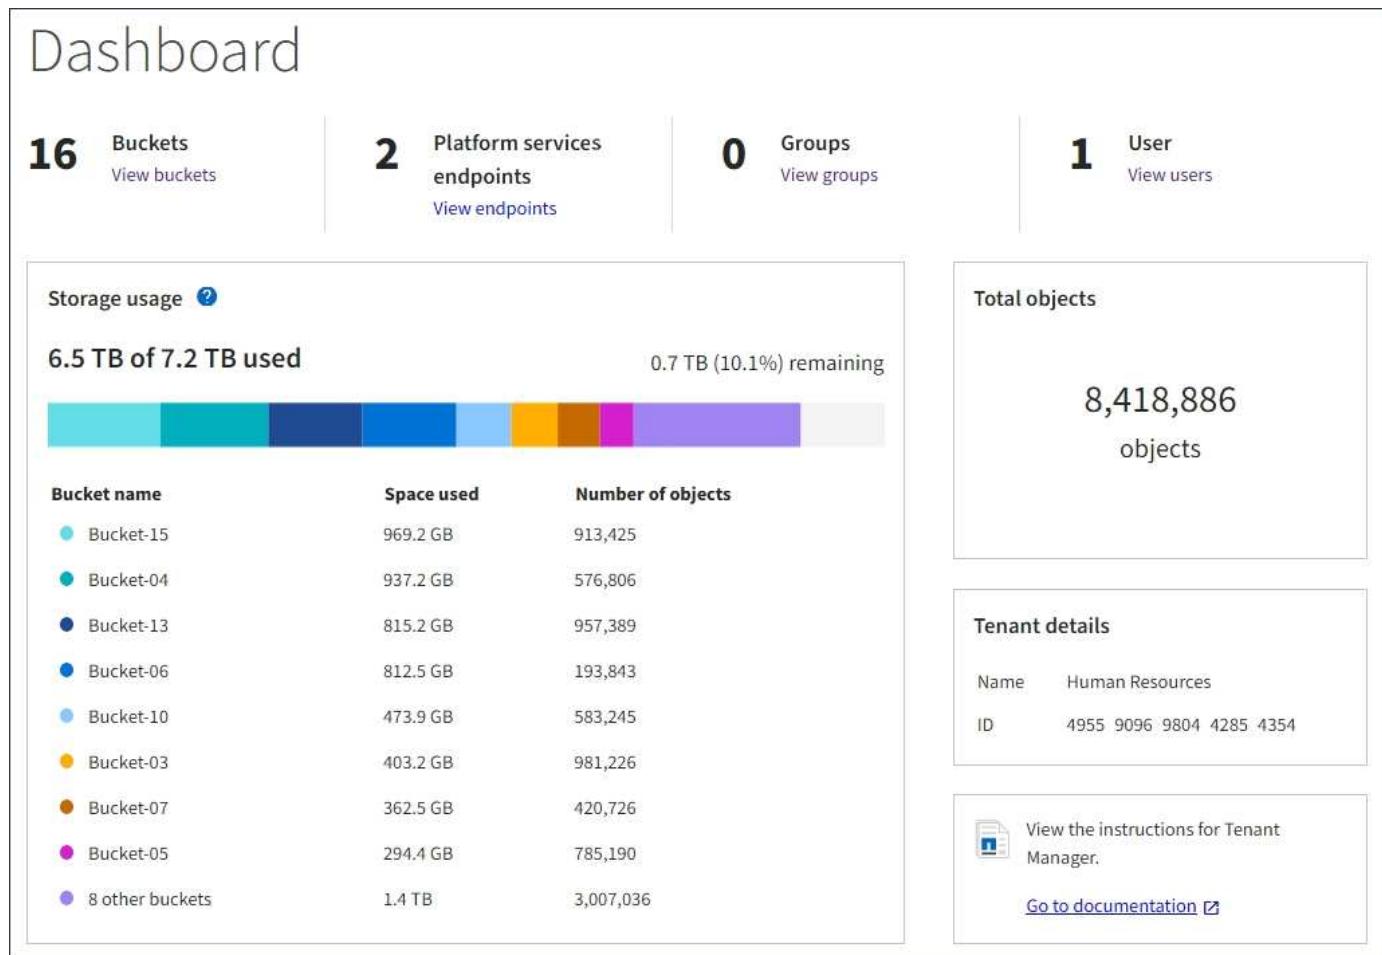
The Grid Manager Dashboard includes pie (doughnut) charts to summarize available storage for the grid and each site.



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

- A list of the largest buckets (S3) or containers (Swift) for the tenant

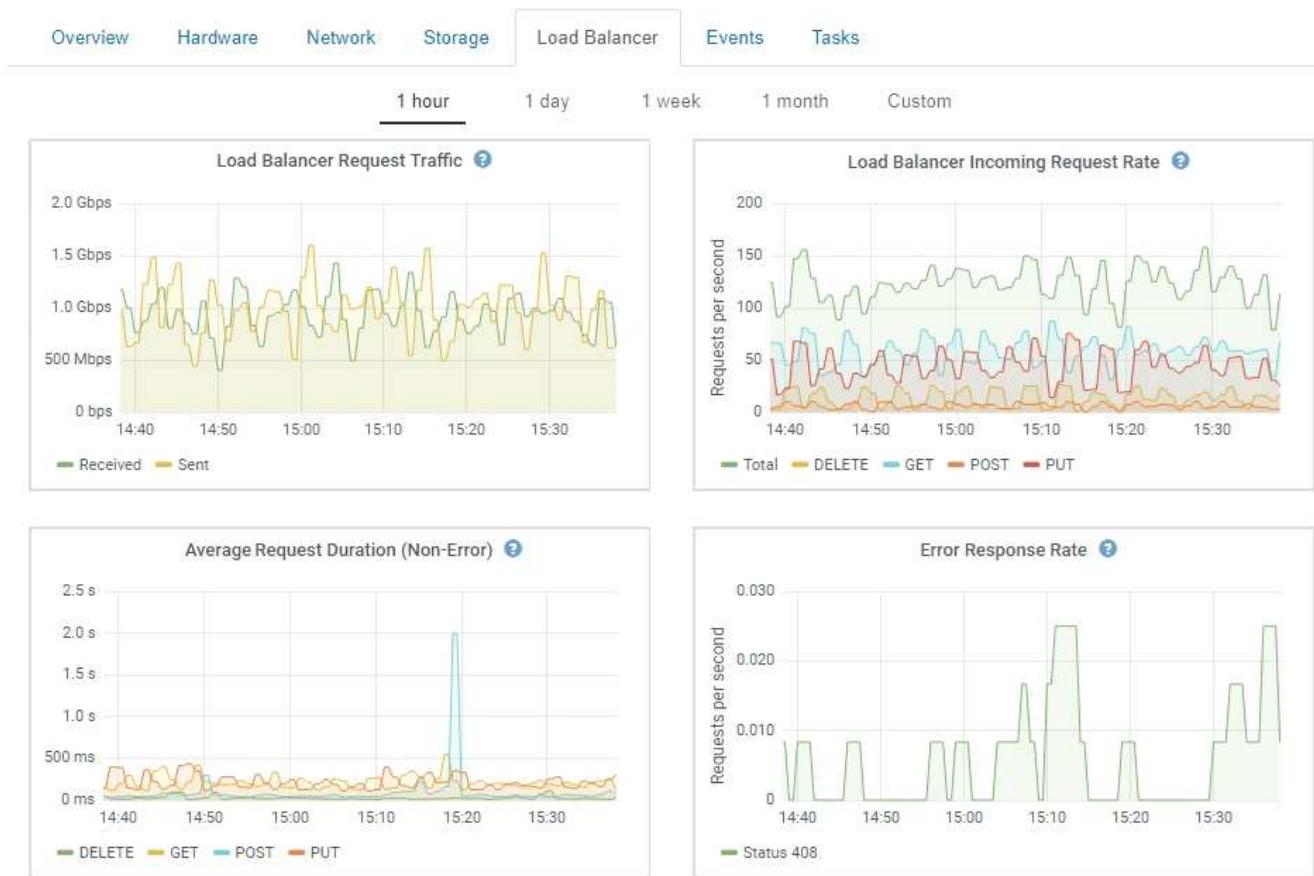
- A bar chart that represents the relative sizes of the largest buckets or containers
- The total amount of space used and, if a quota is set, the amount and percentage of space remaining



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

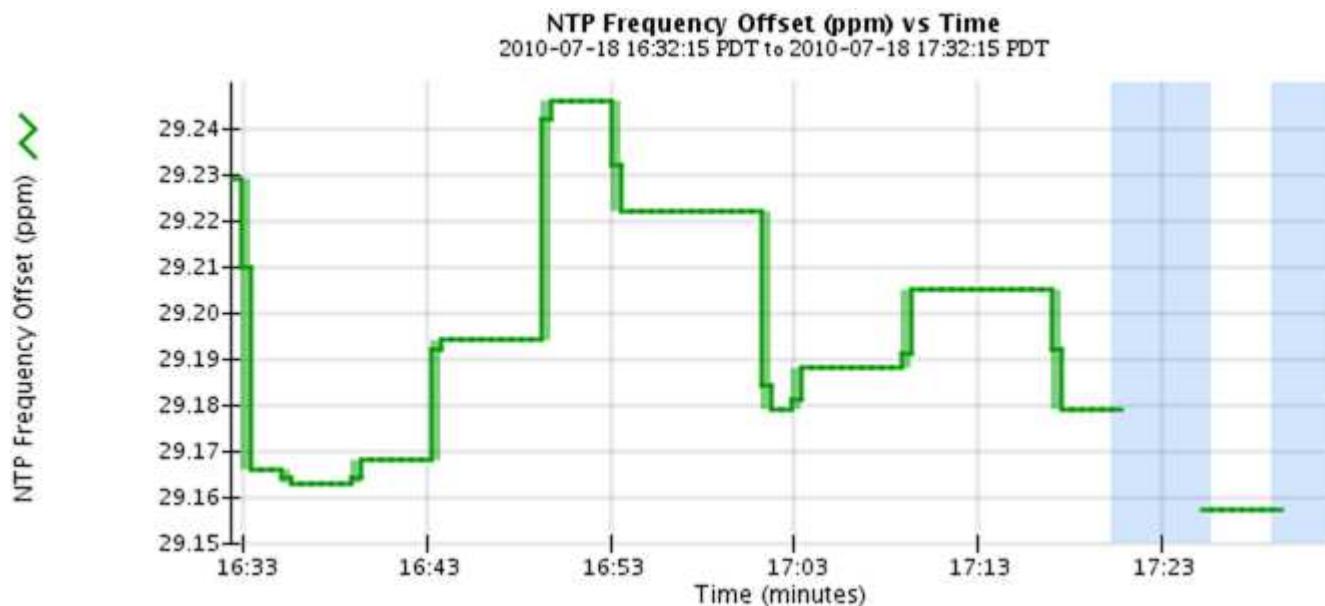
There are four types of graphs:

- **Grafana charts:** Shown on the Nodes page, Grafana charts are used to plot the values of Prometheus metrics over time. For example, the **Nodes > Load Balancer** tab for an Admin Node includes four Grafana charts.

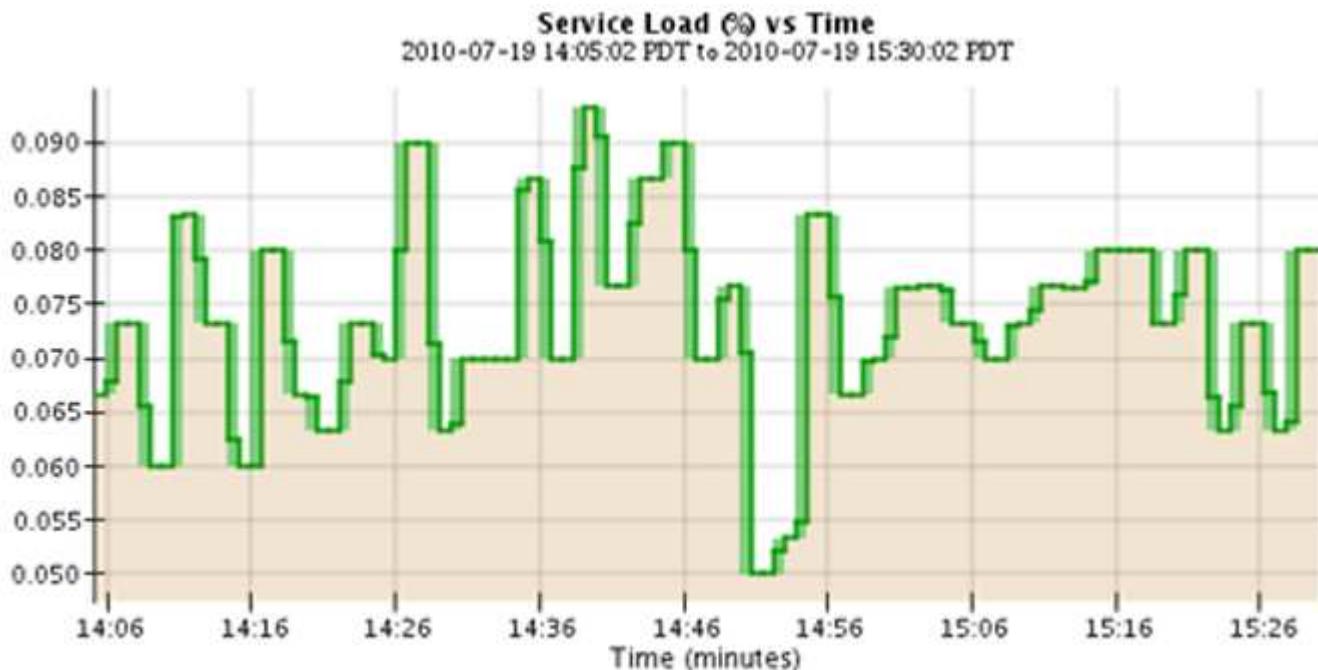


Grafana charts are also included on the pre-constructed dashboards available from the **Support > Tools > Metrics** page.

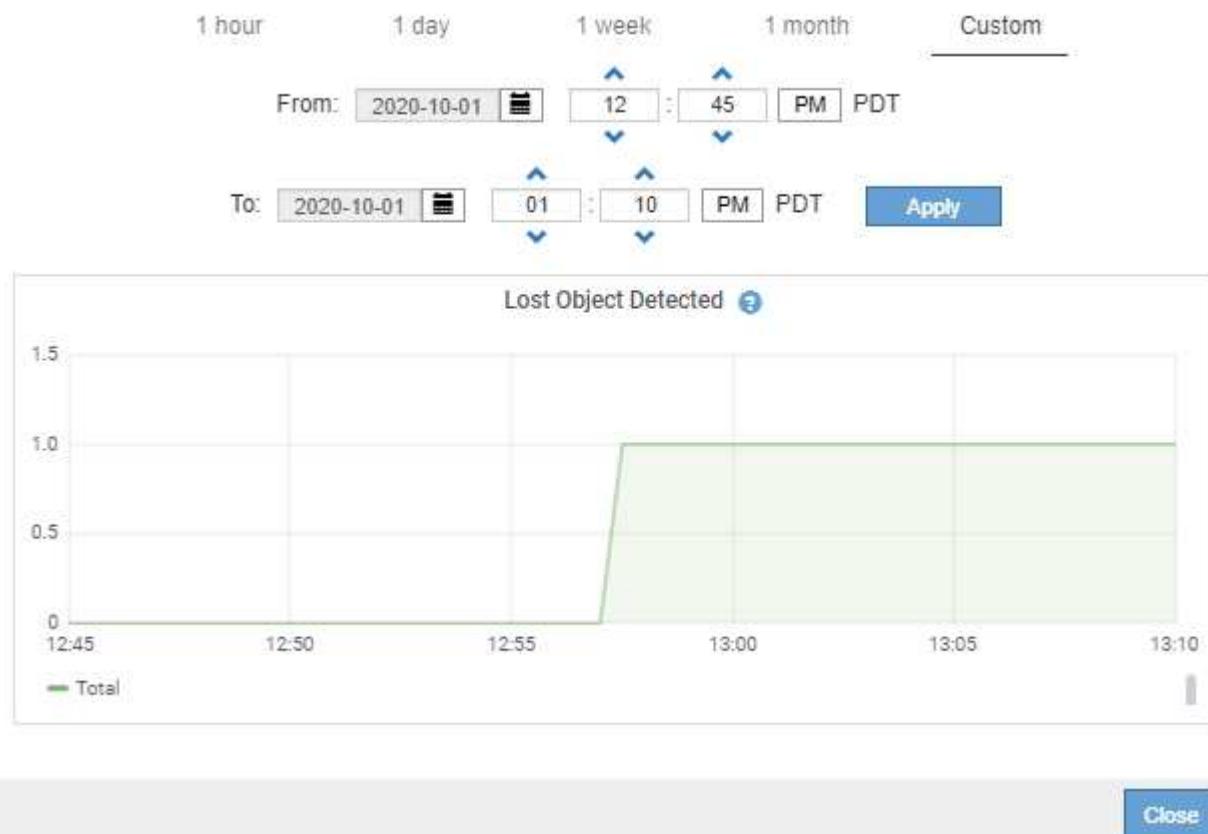
- Line graphs:** Available from the Nodes page and from the **Support > Tools > Grid Topology** page (click the chart icon after a data value), line graphs are used to plot the values of StorageGRID attributes that have a unit value (such as NTP Frequency Offset, in ppm). The changes in the value are plotted in regular data intervals (bins) over time.



- **Area graphs:** Available from the Nodes page and from the **Support > Tools > 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.

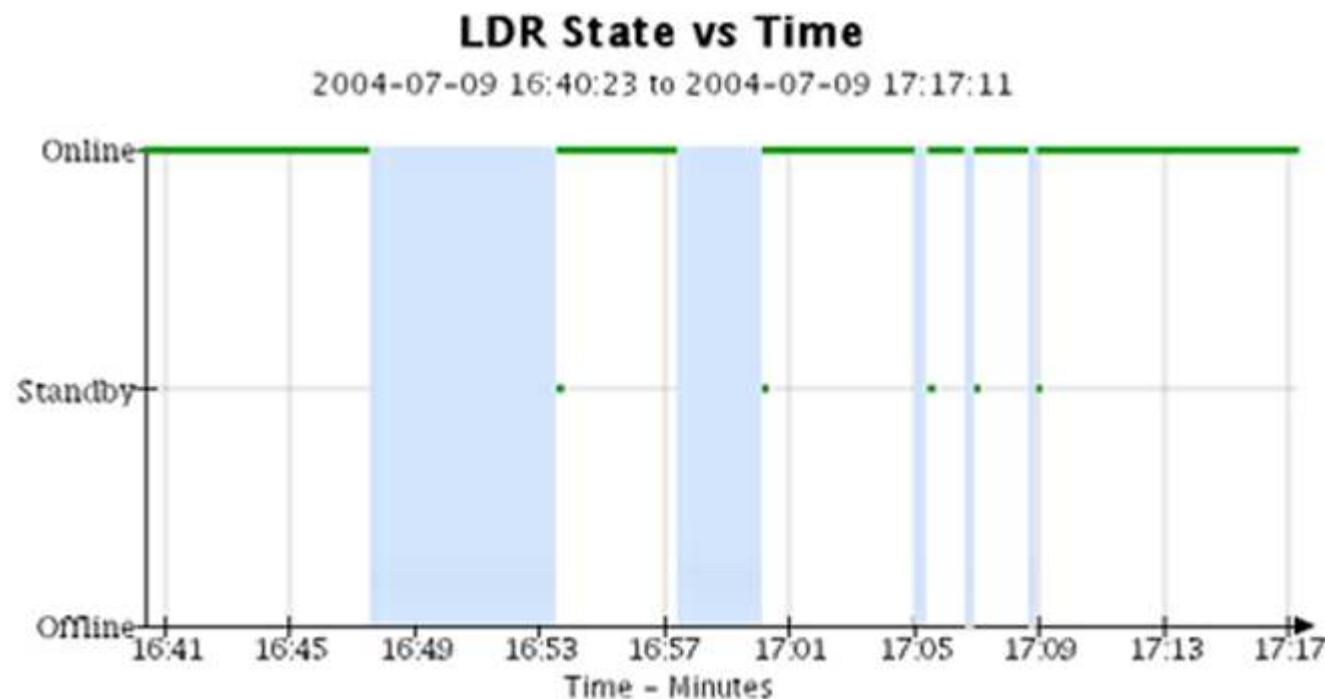


- Some graphs are denoted with a different type of chart icon  and have a different format:



- **State graph:** Available from the **Support > Tools > Grid Topology** page (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 information

[Viewing the Nodes page](#)

[Viewing the Grid Topology tree](#)

[Reviewing support metrics](#)

Chart legend

The lines and colors used to draw charts have specific meaning.

Sample	Meaning
	Reported attribute values are plotted using dark green lines.
	Light green shading around dark green lines indicates that the actual values in that time range vary and have been "binned" for faster plotting. The dark line represents the weighted average. The range in light green indicates the maximum and minimum values within the bin. Light brown shading is used for area graphs to indicate volumetric data.

Sample	Meaning
	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 > Tools > Grid Topology** page to access additional charts.

What you'll need

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

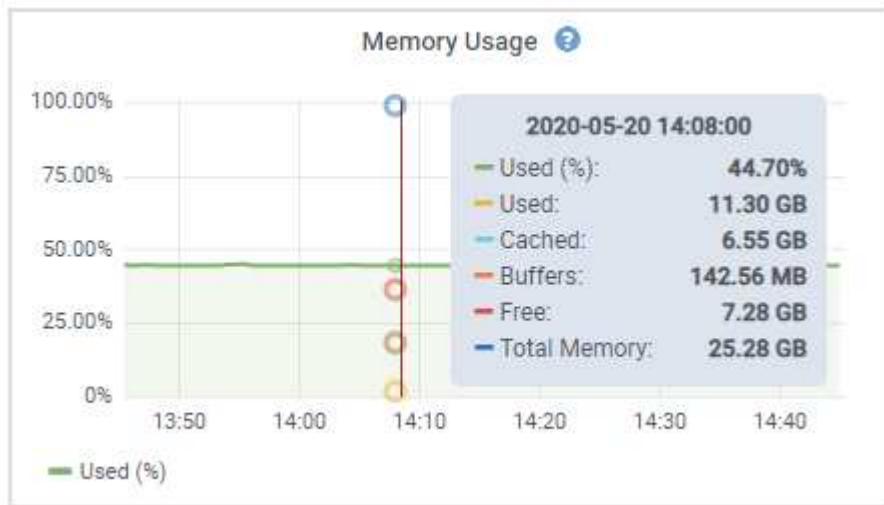
Steps

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** tab for a node includes two Grafana charts.



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



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



Charts are not available for all metrics and attributes.

Example 1: 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.

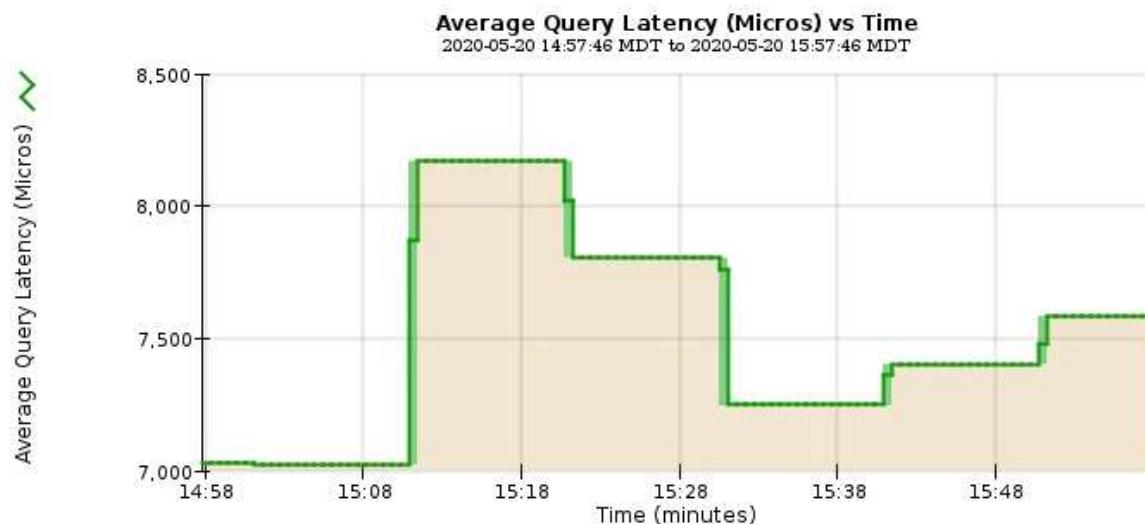
Queries	
Average Latency	14.43 milliseconds
Queries - Successful	19,786
Queries - Failed (timed-out)	0
Queries - Failed (consistency level unmet)	0





Reports (Charts): DDS (DC1-S1) - Data Store

Attribute:	Average Query Latency	Vertical Scaling:	<input checked="" type="checkbox"/>	Start Date:	YYYY/MM/DD HH:MM:SS
Quick Query:	Last Hour	Update	<input type="checkbox"/>	Raw Data:	2020/05/20 14:57:46
				End Date:	2020/05/20 15:57:46

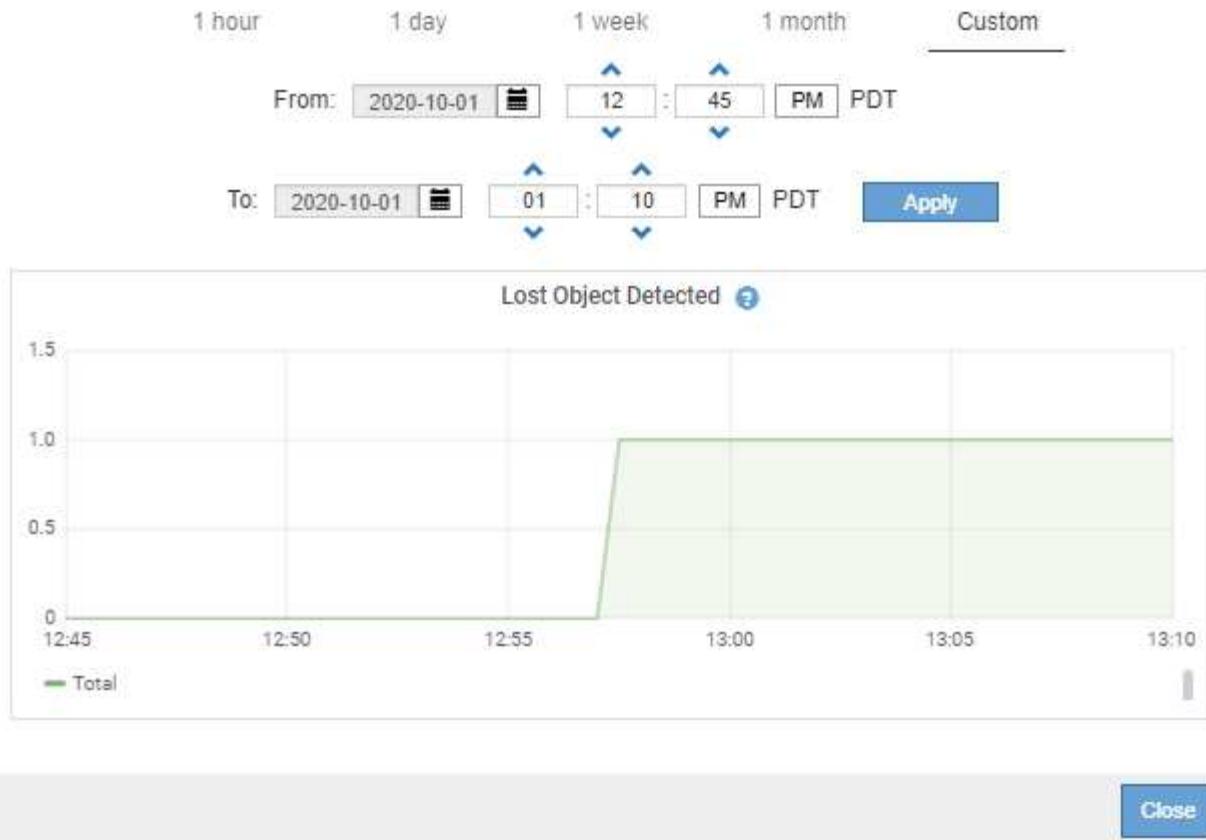


Close

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

Object Counts

Total Objects	1
Lost Objects	1
S3 Buckets and Swift Containers	1



- To display charts for attributes that are not shown on the Node page, select **Support > Tools > 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

- 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.

What you'll need

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

Steps

- Select **Support > Tools > Grid Topology**.
- Select **grid node > component or service > Reports > Charts**.
- Select the attribute to report on from the **Attribute** drop-down list.
- To force the Y-axis to start at zero, deselect the **Vertical Scaling** check box.
- 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/DDHH:MM:SS* in local time. Leading zeros are required to match the format. For example, 2017/4/6 7:30:00 fails validation. The correct format is: 2017/04/06 07:30:00.

8. Click **Update**.

A 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.

What you'll need

- 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 > Tools > Grid Topology**.
2. Select **grid node > component or service > Reports > Text**.
3. Select the attribute to report on from the **Attribute** drop-down list.
4. Select the number of results per page from the **Results per Page** drop-down list.
5. To round values to a maximum of three decimal places (for example, for attributes reported as percentages), 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/DDHH:MM:SS in local time. Leading zeros are required to match the format. For example, 2017/4/6 7:30:00 fails validation. The correct format is: 2017/04/06 07:30:00.

8. Click **Update**.

A text report is generated after a few moments. Allow several minutes for tabulation of long time ranges. Depending on the length of time set for the query, either a raw text report or aggregate text report is displayed.

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

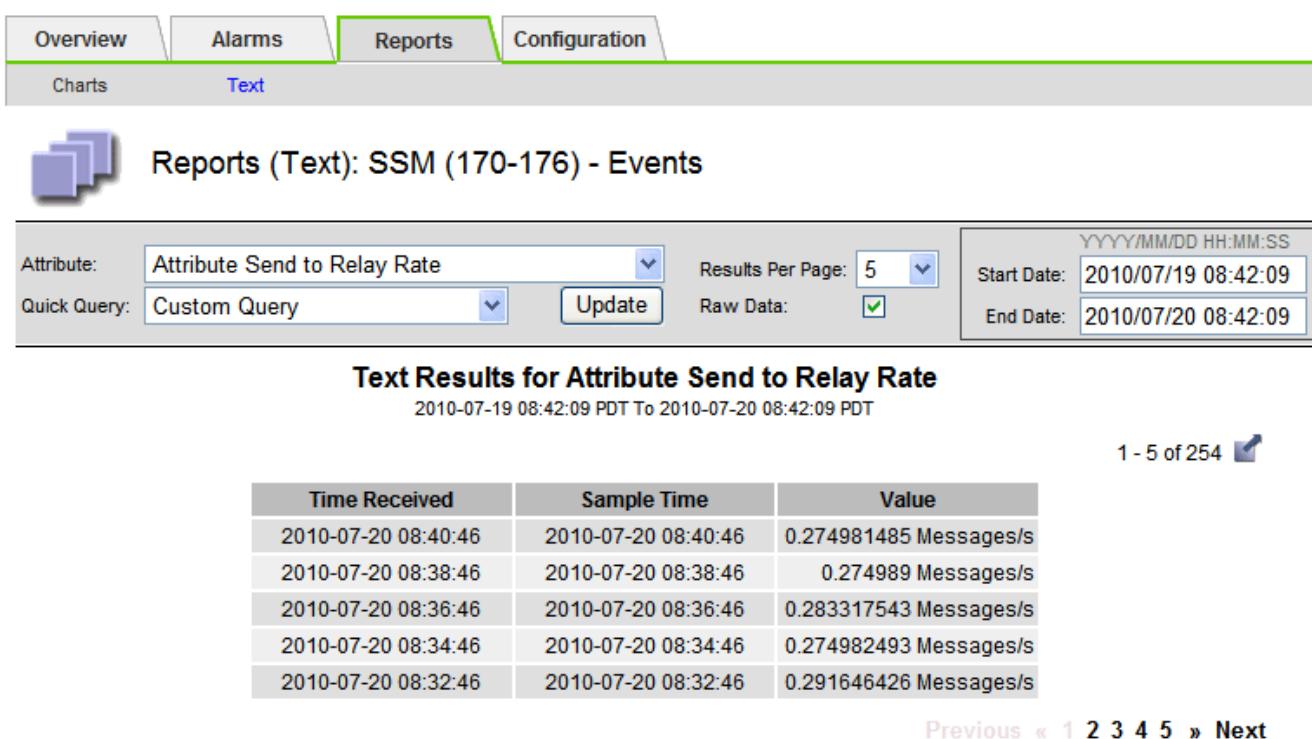
Exported text reports open a new browser tab, which enables you to select and copy the data.

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 > Tools > Grid Topology**.
2. Create a text report.
3. Click ***Export** .



Attribute: Attribute Send to Relay Rate

Quick Query: Custom Query

Results Per Page: 5

Start Date: 2010/07/19 08:42:09

End Date: 2010/07/20 08:42:09

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

Text Results for Attribute Send to Relay Rate

2010-07-19 08:42:09 PDT To 2010-07-20 08:42:09 PDT

1 - 5 of 254

Previous « 1 2 3 4 5 » Next

The Export Text Report window opens displaying the report.

Grid ID: 000 000
OID: 2.16.124.113590.2.1.400019.1.1.1.1.16996732.200
Node Path: Site/170-176/SSM/Events
Attribute: Attribute Send to Relay Rate (ABSR)
Query Start Date: 2010-07-19 08:42:09 PDT
Query End Date: 2010-07-20 08:42:09 PDT
Time Received,Time Received (Epoch),Sample Time,Sample Time (Epoch),Value,Type
2010-07-20 08:40:46,1279640446559000,2010-07-20 08:40:46,1279640446537209,0.274981485 Messages/s,U
2010-07-20 08:38:46,1279640326561000,2010-07-20 08:38:46,1279640326529124,0.274989 Messages/s,U
2010-07-20 08:36:46,1279640206556000,2010-07-20 08:36:46,1279640206524330,0.283317543 Messages/s,U
2010-07-20 08:34:46,1279640086540000,2010-07-20 08:34:46,1279640086517645,0.274982493 Messages/s,U
2010-07-20 08:32:46,1279639966543000,2010-07-20 08:32:46,1279639966510022,0.291646426 Messages/s,U
2010-07-20 08:30:46,1279639846561000,2010-07-20 08:30:46,1279639846501672,0.308315369 Messages/s,U
2010-07-20 08:28:46,1279639726527000,2010-07-20 08:28:46,1279639726494673,0.291657509 Messages/s,U
2010-07-20 08:26:46,1279639606526000,2010-07-20 08:26:46,1279639606490890,0.266627739 Messages/s,U
2010-07-20 08:24:46,1279639486495000,2010-07-20 08:24:46,1279639486473368,0.258318523 Messages/s,U
2010-07-20 08:22:46,1279639366480000,2010-07-20 08:22:46,1279639366466497,0.274985902 Messages/s,U
2010-07-20 08:20:46,1279639246469000,2010-07-20 08:20:46,1279639246460346,0.283253871 Messages/s,U
2010-07-20 08:18:46,1279639126469000,2010-07-20 08:18:46,1279639126426669,0.274982804 Messages/s,U
2010-07-20 08:16:46,1279639006437000,2010-07-20 08:16:46,1279639006419168,0.283315503 Messages/s,U

4. Select and copy the contents of the Export Text Report window.

This data can now be pasted into a third-party document such as a spreadsheet.

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.

- You can download S3tester from github: <https://github.com/s3tester>

Related information

[Review audit logs](#)

Monitoring object verification operations

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

What you'll need

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.



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	



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	



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

Related information

[Verifying object integrity](#)

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 > Tools > Grid Topology**.
2. Select **site > grid node > SSM > Events > Overview > Main**.

Related information

[Resetting event counts](#)

[Log files reference](#)

Reviewing previous events

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

1. Select **Support > Tools > 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**.

Related information

[Using charts and reports](#)

Resetting event counts

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

What you'll need

- 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 !\[\]\(6117e973cfdcd66335cdb4a0bd2694b5_img.jpg\)](#)

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

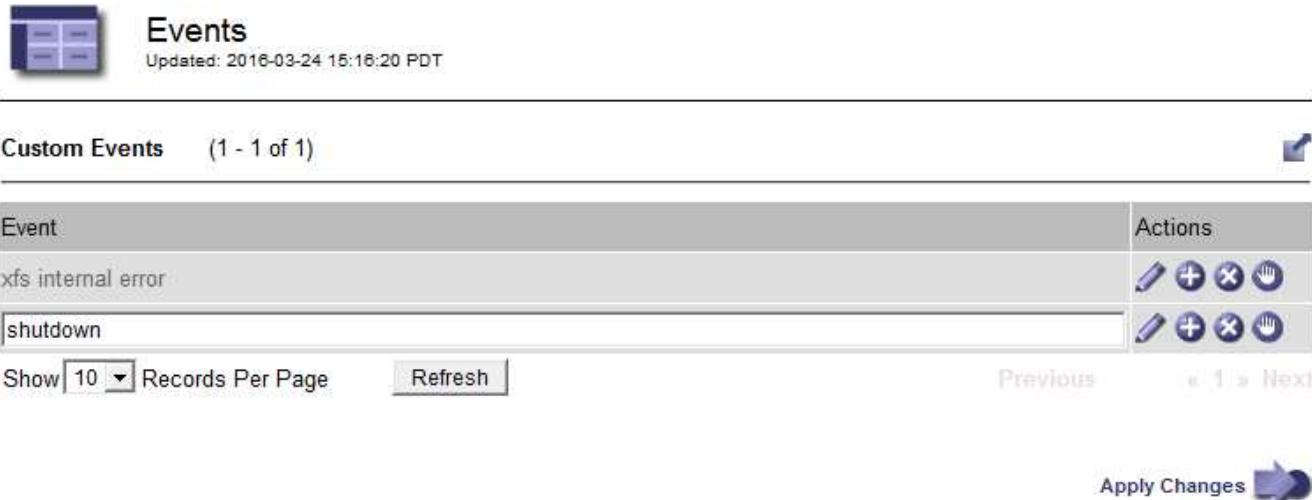
Consider creating custom events to monitor recurring problems. The following considerations apply to custom events.

- 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 > Monitoring > Events**.
2. Click **Edit**  (or **Insert**  if this is not the first event).
3. Enter a custom event string, for example, shutdown



Event	Actions
xfs internal error	   
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.

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	

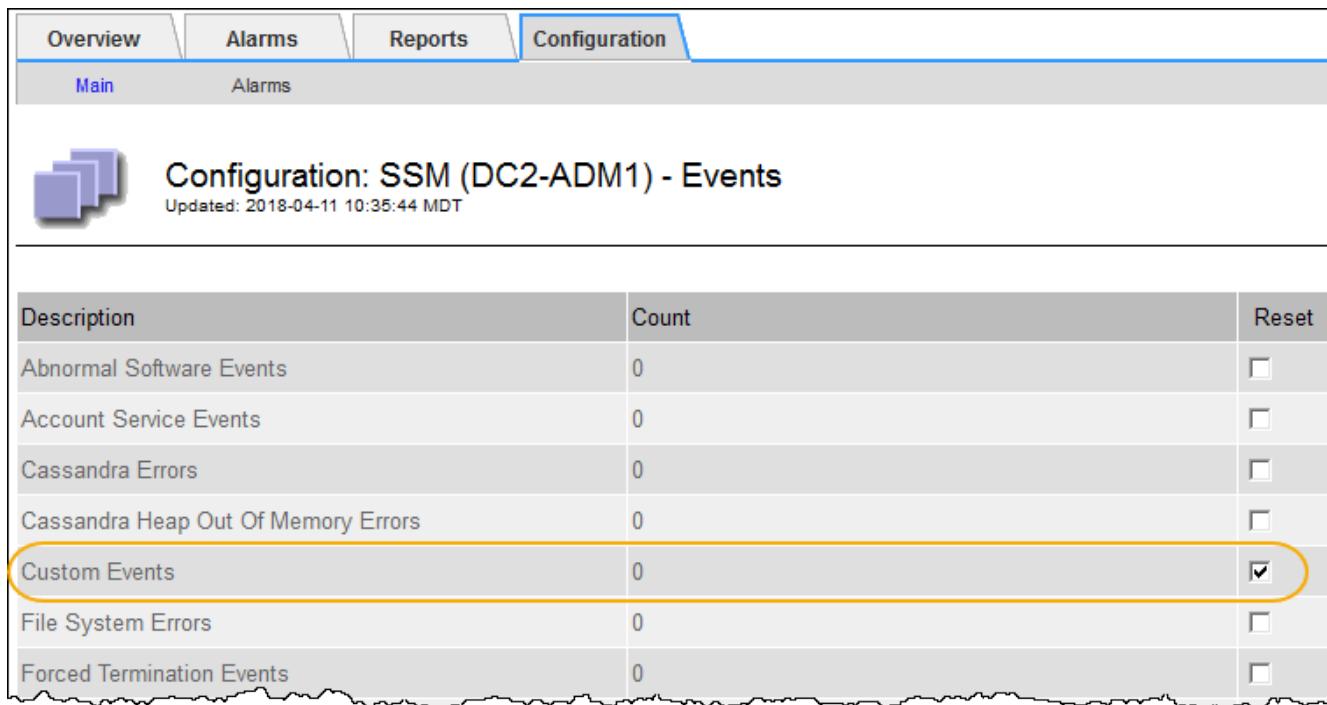
[Reset event counts !\[\]\(e31d5d152cb17b559f8475c09f42ca45_img.jpg\)](#)**Resetting the count of custom events to zero**

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.

1. Select **Support > Tools > Grid Topology**.
2. Select **grid node > SSM > Events > Configuration > Main**.
3. Select the **Reset** check box for Custom Events.



The screenshot shows a user interface for managing events. At the top, there are tabs: Overview, Alarms, Reports, and Configuration. The Configuration tab is selected. Below the tabs, there are sub-tabs: Main and Alarms. The main content area is titled "Configuration: SSM (DC2-ADM1) - Events" and shows the following table:

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

[Review audit logs](#)

[Administer 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.

What you'll need

- 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 > Tools > Logs**.

Logs

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

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:

Collected 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.



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 information

[Log files reference](#)

Manually triggering an AutoSupport message

To assist technical support in troubleshooting issues with your StorageGRID system, you can manually trigger an AutoSupport message to be sent.

What you'll need

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

Steps

1. Select **Support > Tools > AutoSupport**.

The AutoSupport page appears with the **Settings** tab selected.

2. Select **Send User-Triggered AutoSupport**.

StorageGRID attempts to send an AutoSupport message to technical support. If the attempt is successful, the **Most Recent Result** and **Last Successful Time** values on the **Results** tab are updated. If there is a problem, the **Most Recent Result** value updates to "Failed," and StorageGRID does not try to send the AutoSupport message again.



After sending an User-triggered AutoSupport message, refresh the AutoSupport page in your browser after 1 minute to access the most recent results.

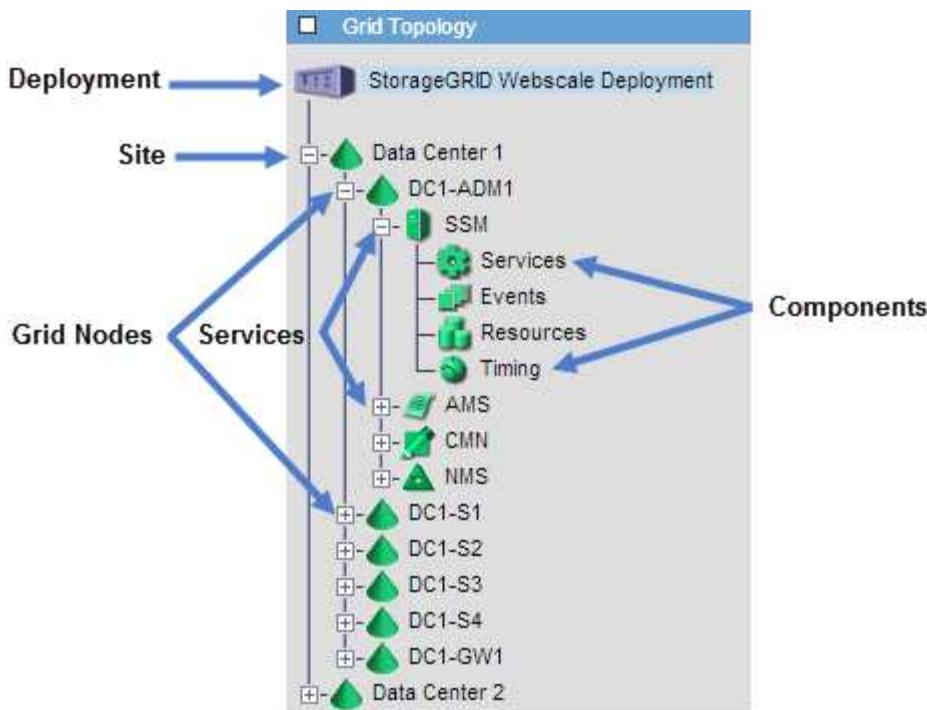
Related information

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

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 > Tools > Grid Topology**.



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

Reviewing support metrics

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

What you'll need

- 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.



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 > Tools > Metrics**.

The Metrics page appears.

Metrics

Access charts and metrics to help troubleshoot issues.

Important: 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://**your_kubelet_ip**:**10255**/metrics/graph](https://your_kubelet_ip:10255/metrics/graph)

Grafana

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

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

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

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.

Enable query history

Expression (press Shift+Enter for newlines)

Execute

- insert metric at cursor - ▾

Graph

Console

Element

Value

no data

Remove Graph

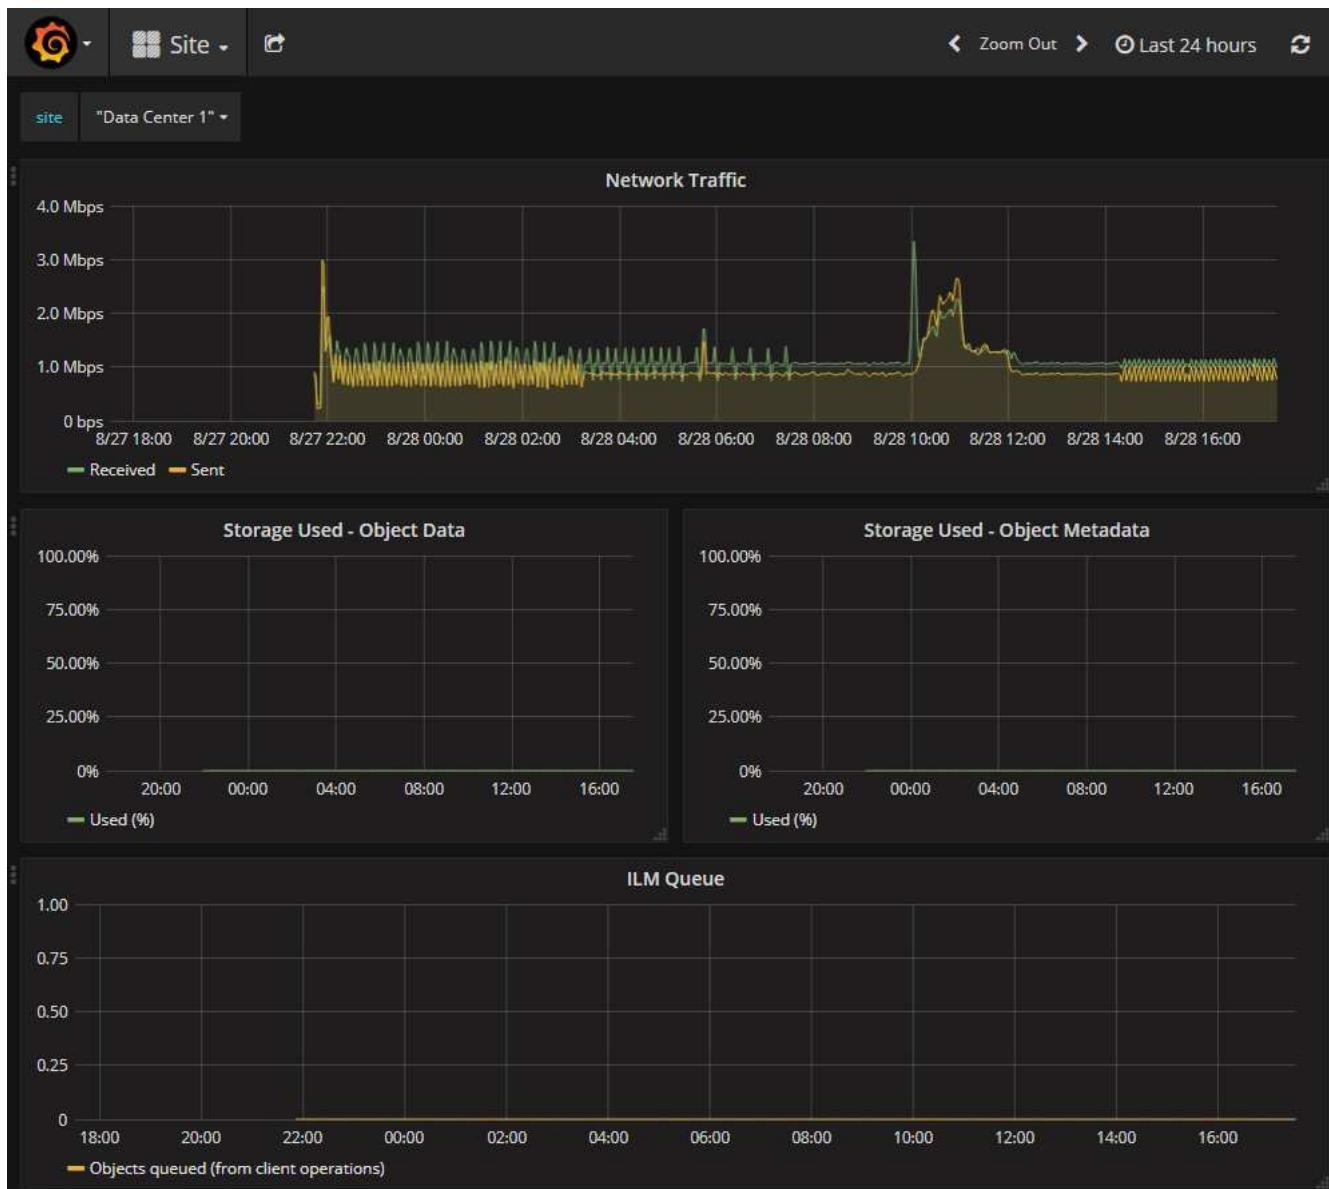
Add Graph



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

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 information

[Commonly used Prometheus metrics](#)

Running diagnostics

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

What you'll need

- 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 > Tools > 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.

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

Run Diagnostics

✓ Cassandra blocked task queue too large	▼
✓ Cassandra commit log latency	▼
✓ Cassandra commit log queue depth	▼
✓ Cassandra compaction queue too large	▼

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.)



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.

✓ [CPU utilization](#)

Checks the current CPU utilization on each node.

To view charts of CPU utilization and other per-node metrics, access the [Node Grafana dashboard](#).

Status	✓ Normal
Prometheus query	<code>sum by (instance) (sum by (instance, mode) (irate(node_cpu_seconds_total{mode!="idle"}[5m])) / count by (instance, mode)(node_cpu_seconds_total{mode!="idle"}))</code>
View in Prometheus	
Thresholds	⚠ Attention >= 75% ✖ Caution >= 95%

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

3. **Optional:** To see Grafana charts related to this diagnostic, click the [Grafana dashboard](#) link.

This link is not displayed for all diagnostics.

The related Grafana dashboard appears. In this example, the Node dashboard appears showing CPU Utilization over time for this node as well as other Grafana charts for the node.



You can also access the pre-constructed Grafana dashboards from the Grafana section of the [Support > Tools > Metrics](#) page.



4. **Optional:** To see a chart of the Prometheus expression over time, click **View in Prometheus**.

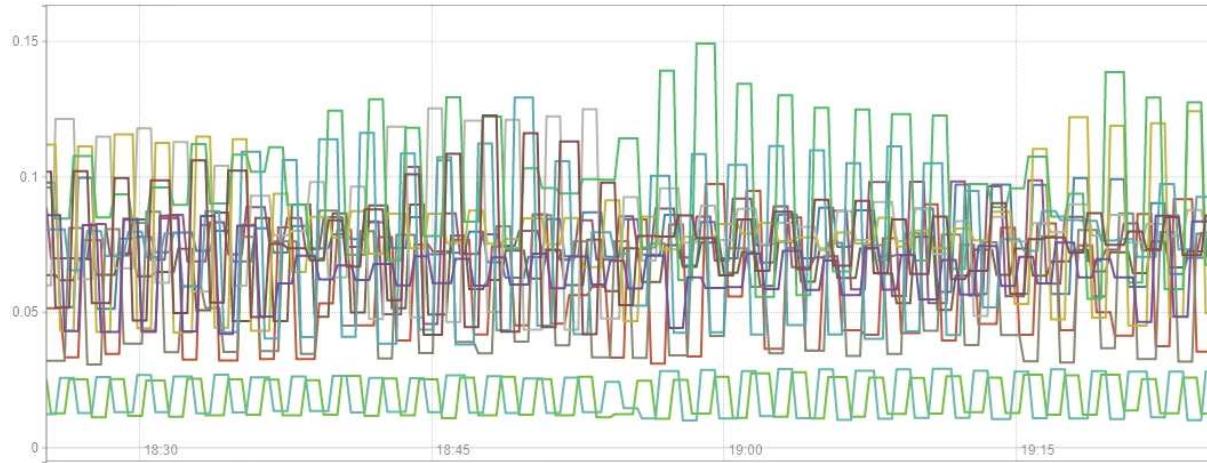
A Prometheus graph of the expression used in the diagnostic appears.

Enable query history

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

Load time: 547ms
Resolution: 14s
Total time series: 13**Execute**

- insert metric at cursor - ▾

Graph**Console**
- 1h + ◀ Until ▶ Res. (s) stacked


- ✓ {instance="DC3-S3"}
- ✓ {instance="DC3-S2"}
- ✓ {instance="DC3-S1"}
- ✓ {instance="DC2-S3"}
- ✓ {instance="DC2-S2"}
- ✓ {instance="DC2-S1"}
- ✓ {instance="DC2-ADM1"}
- ✓ {instance="DC1-S3"}
- ✓ {instance="DC1-S2"}
- ✓ {instance="DC1-S1"}
- ✓ {instance="DC1-G1"}
- ✓ {instance="DC1-ARC1"}
- ✓ {instance="DC1-ADM1"}

Remove Graph**Add Graph**

Related information

[Reviewing support metrics](#)

[Commonly used Prometheus metrics](#)

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.

You can also access Prometheus metrics directly with an external monitoring tool, such as Grafana. Using an external tool requires that you upload or generate an administrative client certificate to allow StorageGRID to authenticate the tool for security. See the instructions for administering StorageGRID.

To view the metrics API operations, including the complete list of the metrics that are available, go to the Grid Manager 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.

Related information

[Administer StorageGRID](#)

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 storage appliances◦ SG5700 storage appliances◦ SG5600 storage appliances2. If this alert persists, contact technical support.

Alert name	Description and recommended actions
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances 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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances 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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances 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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances 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 BMC communication error	<p>Communication with the baseboard management controller (BMC) has been lost.</p> <ol style="list-style-type: none"> 1. Confirm that the BMC is operating normally. Select Nodes, and then select the Hardware tab for the appliance node. Locate the Compute Controller BMC IP field, and browse to that IP. 2. Attempt to restore BMC communications by placing the node into maintenance mode and then powering the appliance off and back on. See the installation and maintenance instructions for your appliance. <ul style="list-style-type: none"> ◦ SG6000 storage appliances ◦ SG100 & SG1000 services appliances 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.

Alert name	Description and recommended actions
Appliance cache backup device insufficient capacity	There is insufficient cache backup device capacity. Contact technical support.
Appliance cache backup device write-protected	A cache backup device is write-protected. Contact technical support.
Appliance cache memory size mismatch	The two controllers in the appliance have different cache sizes. Contact technical support.
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"> <li data-bbox="838 650 1481 713">1. Check the hardware components for overheating conditions, and follow the recommended actions: <ul style="list-style-type: none"> <li data-bbox="904 734 1481 798">◦ If you have an SG100, SG1000, or SG6000, use the BMC. <li data-bbox="904 819 1426 882">◦ If you have an SG5600 or SG5700, use SANtricity System Manager. <li data-bbox="838 903 1481 1009">2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li data-bbox="904 1030 1269 1072">◦ SG6000 storage appliances <li data-bbox="904 1094 1269 1136">◦ SG5700 storage appliances <li data-bbox="904 1157 1269 1199">◦ SG5600 storage appliances <li data-bbox="904 1220 1400 1262">◦ SG100 & SG1000 services appliances

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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances ◦ SG100 & SG1000 services appliances
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances ◦ SG100 & SG1000 services appliances

Alert name	Description and recommended actions
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances ◦ SG100 & SG1000 services appliances
Appliance compute controller power supply B has a problem	<p>Power supply B 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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances ◦ SG100 & SG1000 services appliances

Alert name	Description and recommended actions
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 Fibre Channel fault detected	<p>There is a problem with the Fibre Channel connection between the storage and compute controllers in the appliance.</p> <ol style="list-style-type: none"> 1. Check the hardware components for errors (Nodes > appliance node > Hardware). If the status of any of the components is not “Nominal”, take these actions: <ol style="list-style-type: none"> a. Verify that the Fibre Channel cables between controllers are completely connected. b. Ensure that the Fibre Channel cables are free of excessive bends. c. Confirm that the SFP+ modules are properly seated. 2. If necessary, replace components. See the installation and maintenance instructions for your appliance. <p>Note: If this problem persists, the StorageGRID system might take the problematic connection offline automatically.</p>
Appliance Fibre Channel HBA port failure	<p>A Fibre Channel HBA port is failing or has failed. Contact technical support.</p>
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances 2. If this alert persists, contact technical support.

Alert name	Description and recommended actions
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances 2. If this alert persists, contact technical support.
Appliance LACP port missing	<p>A port on a StorageGRID appliance is not participating in the LACP bond.</p> <ol style="list-style-type: none"> 1. Check the configuration for the switch. Ensure the interface is configured in the correct link aggregation group. 2. If this alert persists, contact technical support.
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances ◦ SG100 & SG1000 services appliances

Alert name	Description and recommended actions
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances

Alert name	Description and recommended actions
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances
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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances

Alert name	Description and recommended actions
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.
Appliance storage shelves 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 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances
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>

Alert name	Description and recommended actions
Cassandra auto-compactor error	<p>The Cassandra auto-compactor has experienced an error. The Cassandra auto-compactor exists on all Storage Nodes and manages the size of the Cassandra database for overwrite and delete heavy workloads. While this condition persists, certain workloads will experience unexpectedly high metadata consumption.</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.
Cassandra auto-compactor metrics out of date	<p>The metrics that describe the Cassandra auto-compactor are out of date. The Cassandra auto-compactor exists on all Storage Nodes and manages the size of the Cassandra database for overwrite and delete heavy workloads. While this alert persists, certain workloads will experience unexpectedly high metadata consumption.</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.

Alert name	Description and recommended actions
Cassandra communication error	<p>The nodes that run the Cassandra service are having trouble communicating with each other. 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 > Tools > 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>Maintain & recover</p>
Cassandra compactions overloaded	<p>The Cassandra compaction process is overloaded. If the compaction process is overloaded, read performance might be degraded and RAM might be used up. The Cassandra service might also become unresponsive or crash.</p> <ol style="list-style-type: none"> 1. Restart the Cassandra service by following the steps for restarting a service in the recovery and maintenance instructions. 2. If this alert persists, contact technical support. <p>Maintain & recover</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. 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. Confirm that all Storage Nodes are online and there are no networking-related alerts. 2. Monitor this alert for up to 2 days to see if the issue resolves on its own. 3. If database repairs continue to proceed slowly, contact technical support.
Cassandra repair service not available	<p>The Cassandra repair service is not available. 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 > Tools > 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>Maintain & recover</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 managing objects with information lifecycle management. <p>Manage objects with ILM</p>

Alert name	Description and recommended actions
DHCP lease expired	<p>The DHCP lease on a network interface has expired. 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>Maintain & recover</p>
DHCP lease expiring soon	<p>The DHCP lease on a network interface is expiring soon. 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>Maintain & recover</p>

Alert name	Description and recommended actions
DHCP server unavailable	<p>The DHCP server is unavailable. 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>Maintain & recover</p>
Disk I/O is very slow	<p>Very slow disk I/O might be impacting StorageGRID performance.</p> <ol style="list-style-type: none"> 1. If the issue is related to a storage appliance node, use SANtricity System Manager to check for faulty drives, drives with predicted faults, or in-progress drive repairs. Also check the status of the Fibre Channel or SAS links between the appliance compute and storage controllers to see if any links are down or showing excessive error rates. 2. Examine the storage system that hosts this node's volumes to determine, and correct, the root cause of the slow I/O. 3. If this alert persists, contact technical support. <p> Affected nodes might disable services and reboot themselves to avoid impacting overall grid performance. When the underlying condition is cleared and these nodes detect normal I/O performance, they will return to full service automatically.</p>

Alert name	Description and recommended actions
Email notification failure	<p>The email notification for an alert could not be sent. 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, select Support > Tools > 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</code> (<code>/_/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 certificates configured on Client Certificates page	<p>One or more certificates configured on the Client Certificates page are about to expire.</p> <ol style="list-style-type: none"> 1. Select Configuration > Access Control > Client Certificates. 2. Select a certificate that will expire soon. 3. Select Edit to upload or generate a new certificate. 4. Repeat these steps for each certificate that will expire soon. <p>Administer StorageGRID</p>

Alert name	Description and recommended actions
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. Select Configuration > Network Settings > 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>Administer StorageGRID</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. Select Configuration > Network Settings > Server Certificates. 2. In the Management Interface Server Certificate section, upload a new certificate. <p>Administer StorageGRID</p>
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. Select Configuration > Network Settings > Server Certificates. 2. In the Object Storage API Service Endpoints Server Certificate section, upload a new certificate. <p>Administer StorageGRID</p>

Alert name	Description and recommended actions
Grid Network MTU mismatch	<p>The maximum transmission unit (MTU) setting for the Grid Network interface (eth0) differs significantly across nodes in the grid. 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>Troubleshooting the Grid Network MTU mismatch alert</p>
High Java heap use	<p>A high percentage of Java heap space is being used. If the Java heap becomes full, metadata services can become unavailable and client requests can fail.</p> <ol style="list-style-type: none"> <li data-bbox="840 720 1470 819">1. Review the ILM activity on the Dashboard. This alert might resolve on its own when the ILM workload decreases. <li data-bbox="840 836 1470 935">2. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert. <li data-bbox="840 952 1437 988">3. If this alert persists, contact technical support.
High latency for metadata queries	<p>The average time for Cassandra metadata queries is too long. 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"> <li data-bbox="840 1241 1503 1341">1. Determine if there were any hardware or workload changes around the time the query latency increased. <li data-bbox="840 1358 1470 1415">2. If you are unable to resolve the problem, contact technical support.

Alert name	Description and recommended actions
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>Administer StorageGRID</p>
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 managing objects with information lifecycle management. 3. As required, update rule settings and activate a new policy. <p> 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> 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>Manage objects with ILM</p>

Alert name	Description and recommended actions
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 > Network Settings > Traffic Classification, 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>Administer StorageGRID</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> Never change the ILM scan rate without contacting technical support.</p>

Alert name	Description and recommended actions
KMS CA certificate expiration	<p>The certificate authority (CA) certificate used to sign the key management server (KMS) certificate is about to expire.</p> <ol style="list-style-type: none"> 1. Using the KMS software, update the CA certificate for the key management server. 2. From the Grid Manager, select Configuration > System Settings > Key Management Server. 3. Select the KMS that has a certificate status warning. 4. Select Edit. 5. Select Next to go to Step 2 (Upload Server Certificate). 6. Select Browse to upload the new certificate. 7. Select Save. <p>Administer StorageGRID</p>
KMS client certificate expiration	<p>The client certificate for a key management server is about to expire.</p> <ol style="list-style-type: none"> 1. From the Grid Manager, select Configuration > System Settings > Key Management Server. 2. Select the KMS that has a certificate status warning. 3. Select Edit. 4. Select Next to go to Step 3 (Upload Client Certificates). 5. Select Browse to upload the new certificate. 6. Select Browse to upload the new private key. 7. Select Save. <p>Administer StorageGRID</p>
KMS configuration failed to load	<p>The configuration for the key management server exists but failed to load.</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. If this alert persists, contact technical support.

Alert name	Description and recommended actions
KMS connectivity error	<p>An appliance node could not connect to the key management server for its site.</p> <ol style="list-style-type: none"> 1. From the Grid Manager, select Configuration > System Settings > Key Management Server. 2. Confirm that the port and hostname entries are correct. 3. Confirm that the server certificate, client certificate, and the client certificate private key are correct and not expired. 4. Ensure that firewall settings allow the appliance node to communicate with the specified KMS. 5. Correct any networking or DNS issues. 6. If you need assistance or this alert persists, contact technical support.
KMS encryption key name not found	<p>The configured key management server does not have an encryption key that matches the name provided.</p> <ol style="list-style-type: none"> 1. Confirm that the KMS assigned to the site is using the correct name for the encryption key and any prior versions. 2. If you need assistance or this alert persists, contact technical support.
KMS encryption key rotation failed	<p>All appliance volumes were decrypted, but one or more volumes could not rotate to the latest key. Contact technical support.</p>
KMS is not configured	<p>No key management server exists for this site.</p> <ol style="list-style-type: none"> 1. From the Grid Manager, select Configuration > System Settings > Key Management Server. 2. Add a KMS for this site or add a default KMS. <p>Administer StorageGRID</p>

Alert name	Description and recommended actions
KMS key failed to decrypt an appliance volume	<p>One or more volumes on an appliance with node encryption enabled could not be decrypted with the current KMS key.</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. Ensure that the key management server (KMS) has the configured encryption key and any previous key versions. 3. If you need assistance or this alert persists, contact technical support.
KMS server certificate expiration	<p>The server certificate used by the key management server (KMS) is about to expire.</p> <ol style="list-style-type: none"> 1. Using the KMS software, update the server certificate for the key management server. 2. If you need assistance or this alert persists, contact technical support. <p>Administer StorageGRID</p>
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 > Monitoring > Audit). <p>Review audit logs</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.

Alert name	Description and recommended actions
Low available node memory	<p>The amount of RAM available on a node is low. 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 free space for storage pool	<p>The amount of space available to store object data in a storage pool is low.</p> <ol style="list-style-type: none"> 1. Select ILM > Storage Pools. 2. Select the storage pool listed in the alert, and select View details. 3. Determine where additional storage capacity is required. You can either add Storage Nodes to each site in the storage pool or add storage volumes (LUNs) to one or more existing Storage Nodes. 4. Perform an expansion procedure to increase storage capacity. <p>Expand your grid</p>
Low installed node memory	<p>The amount of installed memory on a node is low. 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"> • Install Red Hat Enterprise Linux or CentOS • Install Ubuntu or Debian • Install VMware

Alert name	Description and recommended actions
Low metadata storage	<p>The space available for storing object metadata is low. 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 Node > 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</p> <p>Expand your grid</p>
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. 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</p> <p>Expand your grid</p>

Alert name	Description and recommended actions
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 /var/local 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. 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</p>
Node network reception frame error	<p>A high percentage of the network frames received by a node had errors. This alert might indicate a hardware issue, such as a bad cable or a failed transceiver on either end of the Ethernet connection.</p> <ol style="list-style-type: none"> 1. If you are using an appliance, try replacing each SFP+ or SFP28 transceiver and cable, one at a time, to see if the alert clears. 2. If this alert persists, contact technical support.
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.

Alert name	Description and recommended actions
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. 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>
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. <p>Troubleshooting lost and missing object data</p> <ol style="list-style-type: none"> 2. When the underlying problem is resolved, reset the counter: <ol style="list-style-type: none"> a. Select Support > Tools > Grid Topology. b. For the Storage Node that raised the alert, select site > grid node > 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.</p> <p>See “Troubleshooting platform services” in the instructions for administering StorageGRID.</p> <p>Administer StorageGRID</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 Alerts > Alert Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> ◦ SG100 & SG1000 services appliances ◦ Disabling an alert rule
Services appliance link down on Admin Network (or Client Network)	<p>The appliance interface to the 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 Alerts > Alert Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> ◦ SG100 & SG1000 services appliances ◦ Disabling an alert rule
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 Alerts > Alert Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> ◦ SG100 & SG1000 services appliances ◦ Disabling an alert rule

Alert name	Description and recommended actions
Services appliance storage connectivity degraded	<p>One of the two SSDs in a services appliance has failed or is out of synchronization with the other. Appliance functionality is not impacted, but you should address the issue immediately. If both drives fail, the appliance will no longer function.</p> <ol style="list-style-type: none"> 1. From the Grid Manager, select Nodes > services appliance, and then select the Hardware tab. 2. Review the message in the Storage RAID Mode field. 3. If the message shows the progress of a resynchronization operation, wait for the operation to complete and then confirm that the alert is resolved. A resynchronization message means that SSD was replaced recently or that it is being resynchronized for another reason. 4. If the message indicates that one of the SSDs has failed, replace the failed drive as soon as possible. <p>For instructions on how to replace a drive in a services appliance, see the SG100 and SG1000 appliances installation and maintenance guide.</p> <p>SG100 & SG1000 services appliances</p>
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 Alerts > Alert Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> ◦ SG6000 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances ◦ Disabling an alert rule

Alert name	Description and recommended actions
Storage appliance link down on Admin Network (or Client Network)	<p>The appliance interface to the 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 Alerts > Alert Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> ◦ SG6000 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances ◦ Disabling an alert rule
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 Alerts > Alert Rules, select the rule, and click Edit rule. Then, uncheck the Enabled check box. <ul style="list-style-type: none"> ◦ SG6000 storage appliances ◦ SG5700 storage appliances ◦ SG5600 storage appliances ◦ Disabling an alert rule

Alert name	Description and recommended actions
Storage 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> 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> <ol style="list-style-type: none"> 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.

Alert name	Description and recommended actions
Storage device inaccessible	<p>A storage device cannot be accessed. 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"> 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"> Install Red Hat Enterprise Linux or CentOS Install Ubuntu or Debian Install VMware 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. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware. <ul style="list-style-type: none"> SG6000 storage appliances SG5700 storage appliances SG5600 storage appliances
Tenant quota usage high	<p>A high percentage of tenant quota space is being used. If a tenant exceeds its quota, new ingestions are rejected.</p> <p> This alert rule is disabled by default because it might generate a lot of notifications.</p> <ol style="list-style-type: none"> From the Grid Manager, select Tenants. Sort the table by Quota Utilization. Select a tenant whose quota utilization is close to 100%. Do either or both of the following: <ul style="list-style-type: none"> Select Edit to increase the storage quota for the tenant. Notify the tenant that their quota utilization is high.

Alert name	Description and recommended actions
Unable to communicate with node	<p>One or more services are unresponsive, or the node cannot be reached. 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> 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>Maintain & recover</p>
Unexpected node reboot	<p>A node rebooted unexpectedly within the last 24 hours.</p> <ol style="list-style-type: none"> 1. Monitor this alert. The alert will be cleared after 24 hours. However, if the node reboots unexpectedly again, this alert will be triggered again. 2. If you cannot resolve the alert, there might be a hardware failure. Contact technical support.

Alert name	Description and recommended actions
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>Running foreground verification</p>

Related information

[Commonly used Prometheus metrics](#)

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**.



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



Prometheus metrics are retained for 31 days.

Prometheus metric	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.

Prometheus metric	Description
node_network_carrier	Carrier value of /sys/class/net/<iface>.
node_network_receive_errs_total	Network device statistic receive_errs.
node_network_transmit_errs_total	Network device statistic transmit_errs.
storagegrid_administratively_down	The node is not connected to the grid for an expected reason. For example, the node, or services on the node, has been gracefully shut down, the node is rebooting, or the software is being upgraded.
storagegrid_appliance_compute_controller_hardware_status	The status of the compute controller hardware in an appliance.
storagegrid_appliance_failed_disks	For the storage controller in an appliance, the number of drives that 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	The total number of S3 buckets and Swift containers known by this Storage Node.
storagegrid_content_objects	The total number of S3 and Swift data objects known by this Storage Node. Count is valid only for data objects created by client applications that interface with the system through S3 or Swift.
storagegrid_content_objects_lost	<p>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.</p> <p>Troubleshooting lost and missing object data</p>
storagegrid_http_sessions_incoming_attempted	The total number of HTTP sessions that have been attempted to a Storage Node.
storagegrid_http_sessions_incoming_currently_established	The number of HTTP sessions that are currently active (open) on the Storage Node.
storagegrid_http_sessions_incoming_failed	The total number of HTTP sessions that failed to complete successfully, either due to a malformed HTTP request or a failure while processing an operation.

Prometheus metric	Description
storagegrid_http_sessions_incoming_successful	The total number of HTTP sessions that have completed successfully.
storagegrid_ilm_awaiting_background_objects	The total number of objects on this node awaiting ILM evaluation from the scan.
storagegrid_ilm_awaiting_client_evaluation_objects_per_second	The current rate at which objects are evaluated against the ILM policy on this node.
storagegrid_ilm_awaiting_client_objects	The total number of objects on this node awaiting ILM evaluation from client operations (for example, ingest).
storagegrid_ilm_awaiting_total_objects	The total number of objects awaiting ILM evaluation.
storagegrid_ilm_scan_objects_per_second	The rate at which objects owned by this node are scanned and queued for ILM.
storagegrid_ilm_scan_period_estimated_minutes	<p>The estimated time to complete a full ILM scan on this node.</p> <p>Note: A full scan does not guarantee that ILM has been applied to all objects owned by this node.</p>
storagegrid_load_balancer_endpoint_cert_expiry_time	The expiration time of the load balancer endpoint certificate in seconds since the epoch.
storagegrid_metadata_queries_average_latency_milliseconds	The average time required to run a query against the metadata store through this service.
storagegrid_network_received_bytes	The total amount of data received since installation.
storagegrid_network_transmitted_bytes	The total amount of data sent since installation.
storagegrid_ntp_chosen_time_source_offset_milliseconds	Systematic offset of time provided by a chosen time source. Offset is introduced when the delay to reach a time source is not equal to the time required for the time source to reach the NTP client.
storagegrid_ntp_locked	The node is not locked to a network time protocol (NTP) server.
storagegrid_s3_data_transfers_bytes_ingested	The total amount of data ingested from S3 clients to this Storage Node since the attribute was last reset.

Prometheus metric	Description
storagegrid_s3_data_transfers_bytes_retrieved	The total amount of data retrieved by S3 clients from this Storage Node since the attribute was last reset.
storagegrid_s3_operations_failed	The total number of failed S3 operations (HTTP status codes 4xx and 5xx), excluding those caused by S3 authorization failure.
storagegrid_s3_operations_successful	The total number of successful S3 operations (HTTP status code 2xx).
storagegrid_s3_operations_unauthorized	The total number of failed S3 operations that are the result of an authorization failure.
storagegrid_servercertificate_management_interface_cert_expiry_days	The number of days before the Management Interface certificate expires.
storagegrid_servercertificate_storage_api_endpoints_cert_expiry_days	The number of days before the Object Storage API certificate expires.
storagegrid_service_cpu_seconds	The cumulative amount of time that the CPU has been used by this service since installation.
storagegrid_service_load	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	The amount of memory (RAM) currently in use by this service. This value is identical to that displayed by the Linux top utility as RES.
storagegrid_service_network_received_bytes	The total amount of data received by this service since installation.
storagegrid_service_network_transmitted_bytes	The total amount of data sent by this service.
storagegrid_service_restarts	The total number of times the service has been restarted.
storagegrid_service_runtime_seconds	The total amount of time that the service has been running since installation.
storagegrid_service_uptime_seconds	The total amount of time the service has been running since it was last restarted.

Prometheus metric	Description
storagegrid_storage_state_current	<p>The current state of the storage services. Attribute values are:</p> <ul style="list-style-type: none"> • 10 = Offline • 15 = Maintenance • 20 = Read-only • 30 = Online
storagegrid_storage_status	<p>The current status of the storage services. Attribute values are:</p> <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_metadata_bytes	<p>An estimate of the total size of replicated and erasure coded object data on the Storage Node.</p>
storagegrid_storage_utilization_metadata_allowed_bytes	<p>The total space on volume 0 of each Storage Node that is allowed for object metadata. This value is always less than the actual space reserved for metadata on a node, because a portion of the reserved space is required for essential database operations (such as compaction and repair) and future hardware and software upgrades. The allowed space for object metadata controls overall object capacity.</p>
storagegrid_storage_utilization_metadata_bytes	<p>The amount of object metadata on storage volume 0, in bytes.</p>
storagegrid_storage_utilization_metadata_reserved_bytes	<p>The total space on volume 0 of each Storage Node that is actually reserved for object metadata. For any given Storage Node, the actual reserved space for metadata depends on the size of volume 0 for the node and the system-wide Metadata Reserved Space setting.</p>
storagegrid_storage_utilization_total_space_bytes	<p>The total amount of storage space allocated to all object stores.</p>

Prometheus metric	Description
storagegrid_storage_utilization_usable_space_bytes	The total amount of object storage space remaining. Calculated by adding together the amount of available space for all object stores on the Storage Node.
storagegrid_swift_data_transfers_bytes_ingested	The total amount of data ingested from Swift clients to this Storage Node since the attribute was last reset.
storagegrid_swift_data_transfers_bytes_retrieved	The total amount of data retrieved by Swift clients from this Storage Node since the attribute was last reset.
storagegrid_swift_operations_failed	The total number of failed Swift operations (HTTP status codes 4xx and 5xx), excluding those caused by Swift authorization failure.
storagegrid_swift_operations_successful	The total number of successful Swift operations (HTTP status code 2xx).
storagegrid_swift_operations_unauthorized	The total number of failed Swift operations that are the result of an authorization failure (HTTP status codes 401, 403, 405).
storagegrid_tenant_usage_data_bytes	The logical size of all objects for the tenant.
storagegrid_tenant_usage_object_count	The number of objects for the tenant.
storagegrid_tenant_usage_quota_bytes	The maximum amount of logical space available for the tenant's objects. If a quota metric is not provided, an unlimited amount of space is available.

Alarms reference (legacy system)

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



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

Code	Name	Service	Recommended action
ABRL	Available Attribute Relays	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BSSM, BDDS	<p>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.</p> <p>If the problem persists, contact technical support.</p>
ACMS	Available Metadata Services	BARC, BLDR, BCMN	<p>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.</p> <p>Check and restore connections to a DDS service to clear this alarm and return the service to full functionality.</p>

Code	Name	Service	Recommended action
ACTS	Cloud Tiering Service Status	ARC	<p>Only available for Archive Nodes with a Target Type of Cloud Tiering - Simple Storage Service (S3).</p> <p>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.</p> <p>If a major alarm is triggered due to an authentication failure, verify the credentials associated with destination bucket and update values, if necessary.</p> <p>If a major alarm is triggered due to any other reason, contact technical support.</p>
ADCA	ADC Status	ADC	<p>If an alarm is triggered, select Support > Tools > Grid Topology. Then select site > grid node > ADC > Overview > Main and ADC > Alarms > Main to determine the cause of the alarm.</p> <p>If the problem persists, contact technical support.</p>
ADCE	ADC State	ADC	<p>If the value of ADC State is Standby, continue monitoring the service and if the problem persists, contact technical support.</p> <p>If the value of ADC State is Offline, restart the service. If the problem persists, contact technical support.</p>

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 > Tools > Grid Topology. Then select site > grid node > ARC > Retrieve > Configuration > Main, select Archive Retrieve State > Online, and click Apply Changes.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
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	The ADC service has a high number of attribute sessions, and is becoming overloaded. If this alarm is triggered, contact technical support.

Code	Name	Service	Recommended action
ALUR	Unreachable Attribute Repositories	ADC	<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>

Code	Name	Service	Recommended action
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 > Tools > Grid Topology, then select site > grid node > SSM > Events), do the following:</p> <ol style="list-style-type: none"> <li data-bbox="1171 439 1498 572">1. Plug in or replace the cable that does not display any indicator lights. <li data-bbox="1171 587 1416 656">2. Wait one to five minutes. <p data-bbox="1209 692 1470 1058">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.</p> <ol style="list-style-type: none"> <li data-bbox="1171 1100 1498 1374">3. Return to the SSM > Resources page, and verify that the “Degraded” Multipath status has changed to “Nominal” in the Storage Hardware section.

Code	Name	Service	Recommended action
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. 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 > Tools > Grid Topology. Then select site > grid node > ARC > Retrieve > Configuration > Main, select Reset Request Failure Count and click Apply Changes.</p>

Code	Name	Service	Recommended action
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 > Tools > Grid Topology. Then select site > grid node > ARC > Retrieve > Configuration > Main, 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 > Tools > Grid Topology. Then select site > grid node > ARC > Retrieve > Configuration > Main, 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>

Code	Name	Service	Recommended action
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>
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	This alarm is triggered when one or more drives in a StorageGRID appliance has failed or is not optimal. Replace the drives as required.
BASF	Available Object Identifiers	CMN	<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>

Code	Name	Service	Recommended action
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>

Code	Name	Service	Recommended action
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>
BTOF	Offset	BADC, BLDR, BNMS, BAMS, BCLB, BCMN, BARC	<p>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.</p> <p>If the problem persists, contact technical support.</p>
BTSE	Clock State	BADC, BLDR, BNMS, BAMS, BCLB, BCMN, BARC	<p>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.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
CAHP	Java Heap Usage Percent	DDS	<p>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.</p> <p>Check the ILM Activity in the Dashboard, or select Support > Tools > Grid Topology, then select site > grid node > DDS > Resources > Overview > Main.</p> <p>If the problem persists, contact technical support.</p>
CAIH	Number Available Ingest Destinations	CLB	This alarm is deprecated.
CAQH	Number Available Destinations	CLB	<p>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.</p> <p>If the problem persists, contact technical support.</p>

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>Troubleshooting the Services: Status - Cassandra (SVST) alarm</p> <p>If the problem persists, contact technical support.</p>
CASE	Data Store State	DDS	This alarm is triggered during installation or expansion to indicate a new data store is joining the grid.
CCES	Incoming Sessions - Established	CLB	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.

Code	Name	Service	Recommended action
CCNA	Compute Hardware	SSM	This alarm is triggered if the status of the compute controller hardware in a StorageGRID appliance is Needs Attention.

Code	Name	Service	Recommended action
CDLP	Metadata Used Space (Percent)	DDS	<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</p> <p>Expand your grid</p>

Code	Name	Service	Recommended action
CLBA	CLB Status	CLB	<p>If an alarm is triggered, select Support > Tools > Grid Topology, then select site > grid node > 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 > Tools > Grid Topology, then select site > grid node > CMN > Overview > Main and CMN > Alarms > Main to determine the cause of the error and to troubleshoot the problem.</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>

Code	Name	Service	Recommended action
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. 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	FOPN can become large during peak activity. If it does not diminish during periods of slow activity, contact technical support.

Code	Name	Service	Recommended action
HSTE	HTTP State	BLDR	See recommended actions for HSTU.
HSTU	HTTP Status	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 > Tools > Grid Topology. Then select Storage Node > 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>
HTAS	Auto-Start HTTP	LDR	Specifies whether to start HTTP services automatically on start-up. This is a user-specified configuration option.

Code	Name	Service	Recommended action
IRSU	Inbound Replication Status	BLDR, BARC	<p>An alarm indicates that inbound replication has been disabled. Confirm configuration settings: Select Support > Tools > Grid Topology. Then select <i>site > grid node > LDR > Replication > Configuration > Main</i>.</p>
LATA	Average Latency	NMS	<p>Check for connectivity issues.</p> <p>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.</p> <p>Check for multiple alarms. An increase in average latency times can be indicated by an excessive number of triggered alarms.</p> <p>If the problem persists, contact technical support.</p>
LDRE	LDR State	LDR	<p>If the value of LDR State is Standby, continue monitoring the situation and if the problem persists, contact technical support.</p> <p>If the value of LDR State is Offline, restart the service. If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
LOST	Lost Objects	DDS, LDR	<p>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.</p> <p>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.</p> <p>Investigate LOST (LOST Object) alarms immediately. If the problem persists, contact technical support.</p> <p>Troubleshooting lost and missing object data</p>
MCEP	Management Interface Certificate Expiry	CMN	<p>Triggered when the certificate used for accessing 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>Administer StorageGRID</p>

Code	Name	Service	Recommended action
MINQ	E-mail Notifications Queued	NMS	<p>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.</p> <p>Configuring email server settings for alarms (legacy system)</p>
MINS	E-mail Notifications Status	BNMS	<p>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.</p> <p>Configuring email server settings for alarms (legacy system)</p>
MISS	NMS Interface Engine Status	BNMS	<p>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.</p>
NANG	Network Auto Negotiate Setting	SSM	<p>Check the network adapter configuration. The setting must match preferences of your network routers and switches.</p> <p>An incorrect setting can have a severe impact on system performance.</p>

Code	Name	Service	Recommended action
NDUP	Network Duplex Setting	SSM	<p>Check the network adapter configuration. The setting must match preferences of your network routers and switches.</p> <p>An incorrect setting can have a severe impact on system performance.</p>
NLNK	Network Link Detect	SSM	<p>Check the network cable connections on the port and at the switch.</p> <p>Check the network router, switch, and adapter configurations.</p> <p>Restart the server.</p> <p>If the problem persists, contact technical support.</p>
NRER	Receive Errors	SSM	<p>The following can be causes of NRER alarms:</p> <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 <p>Troubleshooting the Network Receive Error (NRER) alarm</p>

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	If the value of NMS Status is DB Connectivity Error, restart the service. If the problem persists, contact technical support.
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	This can be caused by network connectivity or driver compatibility issues. If the problem persists, contact technical support.

Code	Name	Service	Recommended action
NTBR	Free Tablespace	NMS	<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 > Tools > Grid Topology. Then select site > grid node > SSM > Resources > Configuration > Main, select Reset Transmit Error Count, and click Apply Changes.</p>

Code	Name	Service	Recommended action
NTFQ	NTP Frequency Offset	SSM	If the frequency offset exceeds the configured threshold, there is likely a hardware problem with the local clock. If the problem persists, contact technical support to arrange a replacement.
NTLK	NTP Lock	SSM	If the NTP daemon is not locked to an external time source, check network connectivity to the designated external time sources, their availability, and their stability.
NTOF	NTP Time Offset	SSM	If the time offset exceeds the configured threshold, there is likely a hardware problem with the oscillator of the local clock. If the problem persists, contact technical support to arrange a replacement.
NTSJ	Chosen Time Source Jitter	SSM	This value indicates the reliability and stability of the time source that NTP on the local server is using as its reference. If an alarm is triggered, it can be an indication that the time source's oscillator is defective, or that there is a problem with the WAN link to the time source.
NTSU	NTP Status	SSM	If the value of NTP Status is Not Running, contact technical support.

Code	Name	Service	Recommended action
OPST	Overall Power Status	SSM	<p>An alarm is triggered if the power of a StorageGRID appliance deviates from the recommended operating voltage.</p> <p>Check the status of Power Supply A or B to determine which power supply is operating abnormally.</p> <p>If necessary, replace the power supply.</p>
OQRT	Objects Quarantined	LDR	<p>After the objects are automatically restored by the StorageGRID system, the quarantined objects can be removed from the quarantine directory.</p> <ol style="list-style-type: none"> 1. Select Support > Tools > Grid Topology. 2. Select site > Storage Node > 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>

Code	Name	Service	Recommended action
ORSU	Outbound Replication Status	BLDR, BARC	<p>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 > Tools > Grid Topology. Then select site > grid node > LDR > Replication > Configuration.</p> <p>An alarm is triggered if the LDR service is unavailable for replication. Select Support > Tools > Grid Topology. Then select site > grid node > LDR > Storage.</p>
OSLF	Shelf Status	SSM	<p>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.</p>

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>

Code	Name	Service	Recommended action
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 > Tools > Grid Topology. Then select site > grid node > 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 251 errors.</p>

Code	Name	Service	Recommended action
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 > Tools > Grid Topology, then select site > grid node > LDR > Replication > Configuration > Main. 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>

Code	Name	Service	Recommended action
SCAS	Status	CMN	<p>If the value of Status for the active grid task is Error, look up the grid task message. Select Support > Tools > Grid Topology. Then select site > grid node > CMN > Grid Tasks > Overview > Main. The grid task message displays information about the error (for example, “check failed on node 12130011”).</p> <p>After you have investigated and corrected the problem, restart the grid task. Select Support > Tools > Grid Topology. Then select site > grid node > CMN > Grid Tasks > Configuration > Main, and select Actions > Run.</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"> 1. Go to Configuration > Server Certificates. 2. In the Object Storage API Service Endpoints Server Certificate section, upload a new certificate. <p>Administer StorageGRID</p>

Code	Name	Service	Recommended action
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>
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

Code	Name	Service	Recommended action
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 > Tools > Grid Topology. Then select site > grid node > SSM > Resources > Reports > Charts.</p> <p>If the load on the system is not heavy and the problem persists, contact technical support.</p>
SMST	Log Monitor State	SSM	If the value of Log Monitor State is not Connected for a persistent period of time, contact technical support.

Code	Name	Service	Recommended action
SMTT	Total Events	SSM	<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> 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>

Code	Name	Service	Recommended action
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 appliance Storage Node > Hardware. Scroll down to view the status of each component. In SANtricity software, check other appliance components to isolate the issue.</p>
SSMA	SSM Status	SSM	<p>If the value of SSM Status is Error, select Support > Tools > Grid Topology, then select site > grid node > SSM > Overview > Main and SSM > Overview > Alarms to determine the cause of the alarm.</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. If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
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 > Tools > Grid Topology. Then select site > grid node > 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</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 > Tools > Grid Topology. Then select <i>site > grid node > SSM > Services > Overview > Main</i>. 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.</p>

Code	Name	Service	Recommended action
TMEM	Installed Memory	SSM	Nodes running with less than 24 GiB of installed memory can lead to performance problems and system instability. The amount of memory installed on the system should be increased to at least 24 GiB.
TPOP	Pending Operations	ADC	A queue of messages can indicate that the ADC service is overloaded. Too few ADC services can be connected to the StorageGRID system. In a large deployment, the ADC service can require adding computational resources, or the system can require additional ADC services.
UMEM	Available Memory	SSM	If the available RAM gets low, determine whether this is a hardware or software issue. If it is not a hardware issue, or if available memory falls below 50 MB (the default alarm threshold), contact technical support.
VMFI	Entries Available	SSM	This is an indication that additional storage is required. Contact technical support.

Code	Name	Service	Recommended action
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>
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.

Code	Name	Service	Recommended action
VSTU	Object Verification Status	BLDR	<p>Select Support > Tools > Grid Topology. Then select <i>site > grid node > LDR > Storage > Overview > Main</i>.</p> <p>Check the operating system for any signs of block-device or file system errors.</p> <p>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.</p>
XAMS	Unreachable Audit Repositories	BADC, BARC, BCLB, BCMN, BLDR, BNMS	<p>Check network connectivity to the server hosting the Admin Node.</p> <p>If the problem persists, contact technical support.</p>

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.



While the legacy alarm system continues to be supported, the 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

Code	Name	Severity
AOTE	Store State	Minor
AOTU	Store Status	Major
AROQ	Objects Queued	Minor
ARRF	Request Failures	Major
ARRV	Verification Failures	Major
ARVF	Store Failures	Major
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

Code	Name	Severity
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
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
SAVP	Total Usable Space (Percent)	Notice
SHLH	Health	Notice
SLSA	CPU Load Average	Notice
SMTT	Total Events	Notice
SNST	Status	

Code	Name	Severity
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.

- [StorageGRID software logs](#)
- [Deployment and maintenance logs](#)
- [Logs for third-party software](#)
- [About the bycast.log](#)

 The tables in this section 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 > Tools > 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 information

[Collecting log files and system data](#)

StorageGRID software logs

You can use StorageGRID logs to troubleshoot issues.

General StorageGRID logs

File name	Notes	Found on
/var/local/log/bycast.log	<p>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 <code>ERROR</code> and <code>CRITICAL</code>). <code>CRITICAL</code> messages are also displayed in the system.</p> <p>Select Support > Tools > Grid Topology. Then select Site > Node > SSM > Events.</p>	All nodes
/var/local/log/bycast-err.log	<p>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 <code>ERROR</code> and <code>CRITICAL</code>). <code>CRITICAL</code> messages are also displayed in the system.</p> <p>Select Support > Tools > Grid Topology. Then select Site > Node > SSM > Events.</p>	All nodes
/var/local/core/	<p>Contains any core dump files created if the program terminates abnormally. Possible causes include assertion failures, violations, or thread timeouts.</p> <p>Note: The file <code>/var/local/core/kexec_cmd</code> usually exists on appliance nodes and does not indicate an error.</p>	All nodes

Server Manager logs

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

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

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

File name	Notes	Found on
/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
/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

File name	Notes	Found on
/var/local/log/miscd/*.log	Contains logs for the MISCD service (Information Service Control Daemon), which provides an interface for querying and managing services on other nodes and for managing environmental configurations on the node such as querying the state of services running on other nodes.	All nodes
/var/local/log/nginx/*.log	Contains logs for the nginx service, which acts as an authentication and secure communication mechanism for various grid services (such as Prometheus and Dynip) to be able to talk to services on other nodes over HTTPS APIs.	All nodes
/var/local/log/nginx-gw/*.log	Contains 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
/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

File name	Notes	Found on
/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

File name	Notes	Found on
/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>	Admin Nodes
/var/local/log/nms.request.log	Contains information about outgoing connections from the Management API to internal StorageGRID services.	Admin Nodes

Related information

[About the bycast.log](#)

[Use S3](#)

Deployment and maintenance logs

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

File name	Notes	Found on
/var/local/log/install.log	Created during software installation. Contains a record of the installation events.	All nodes
/var/local/log/expansion-progress.log	Created during expansion operations. Contains a record of the expansion events.	Storage Nodes
/var/local/log/gdu-server.log	Created by the GDU service. Contains events related to provisioning and maintenance procedures managed by the primary Admin Node.	Primary Admin Node
/var/local/log/send_admin_hw.log	Created during installation. Contains debugging information related to a node's communications with the primary Admin Node.	All nodes
/var/local/log/upgrade.log	Created during software upgrade. Contains a record of the software update events.	All nodes

Logs for third-party software

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

Category	File name	Notes	Found on
apache2 logs	/var/local/log/apache2/access.log /var/local/log/apache2/error.log /var/local/log/apache2/other_vhosts_access.log	Log files for apache2.	Admin Nodes
Archiving	/var/local/log/dsie rror.log	Error information for TSM Client APIs.	Archive Nodes
MySQL	/var/local/log/mysql.err` /var/local/log/mysql.err /var/local/log/mysql-slow.log	<p>Log files generated by MySQL.</p> <p>The file mysql.err captures database errors and events such as startups and shutdowns.</p> <p>The file mysql-slow.log (the slow query log) captures the SQL statements that took more than 10 seconds to execute.</p>	Admin Nodes
Operating system	/var/local/log/messages	This directory contains log files for the operating system. The errors contained in these logs are also displayed in the Grid Manager. Select Support > Tools > Grid Topology . Then select Topology > Site > Node > SSM > Events .	All nodes

Category	File name	Notes	Found on
NTP	/var/local/log/ntp.log /var/lib/ntp/var/log/ntpstats/	The /var/local/log/ntp.log contains the log file for NTP error messages. The /var/lib/ntp/var/log/ntpstats/ directory contains NTP timing statistics. loopstats records loop filter statistics information. peerstats records peer statistics information.	All nodes
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 `broadcast.log`

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

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

File rotation for `broadcast.log`

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

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

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

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



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

Related information

[Collecting log files and system data](#)

Messages in `broadcast.log`

Messages in `broadcast.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 `broadcast.log`

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

For example:

- **NOTICE** — An event that should be recorded has occurred. Most log messages are at this level.
- **WARNING** — An unexpected condition has occurred.
- **ERROR** — A major error has occurred that will impact operations.
- **CRITICAL** — An abnormal condition has occurred that has stopped normal operations. You should address

the underlying condition immediately. Critical messages are also displayed in the Grid Manager. Select **Support > Tools > Grid Topology**. Then select **Site > Node > SSM > Events**.

Error codes in `broadcast.log`

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

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

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

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

Error number	Error code	Meaning
001	EPERM	Operation not permitted
002	ENOENT	No such file or directory
003	ESRCH	No such process

Error number	Error code	Meaning
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
023	ENFILE	File table overflow
024	EMFILE	Too many open files
025	ENOTTY	Not a typewriter

Error number	Error code	Meaning
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

Error number	Error code	Meaning
047	EL3RST	Level 3 reset
048	ELNRNG	Link number out of range
049	EUNATCH	Protocol driver not attached
050	ENOCSI	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
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

Error number	Error code	Meaning
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	EOVERFLOW	Value too large for defined data type
076	ENOTUNIQ	Name not unique on network
077	EBADFD	File descriptor in bad state
078	EREMCHG	Remote address changed
079	ELIBACC	Cannot access a needed shared library
080	ELIBBAD	Accessing a corrupted shared library
081	ELIBSCN	
082	ELIBMAX	Attempting to link in too many shared libraries
083	ELIBEXEC	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

Error number	Error code	Meaning
088	ENOTSOCK	Socket operation on non-socket
089	EDESTADDRREQ	Destination address required
090	EMSGSIZE	Message too long
091	EPROTOTYPE	Protocol wrong type for socket
092	ENOPROTOOPT	Protocol not available
093	EPROTONOSUPPORT	Protocol not supported
094	ESOCKTNOSUPPORT	Socket type not supported
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

Error number	Error code	Meaning
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
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

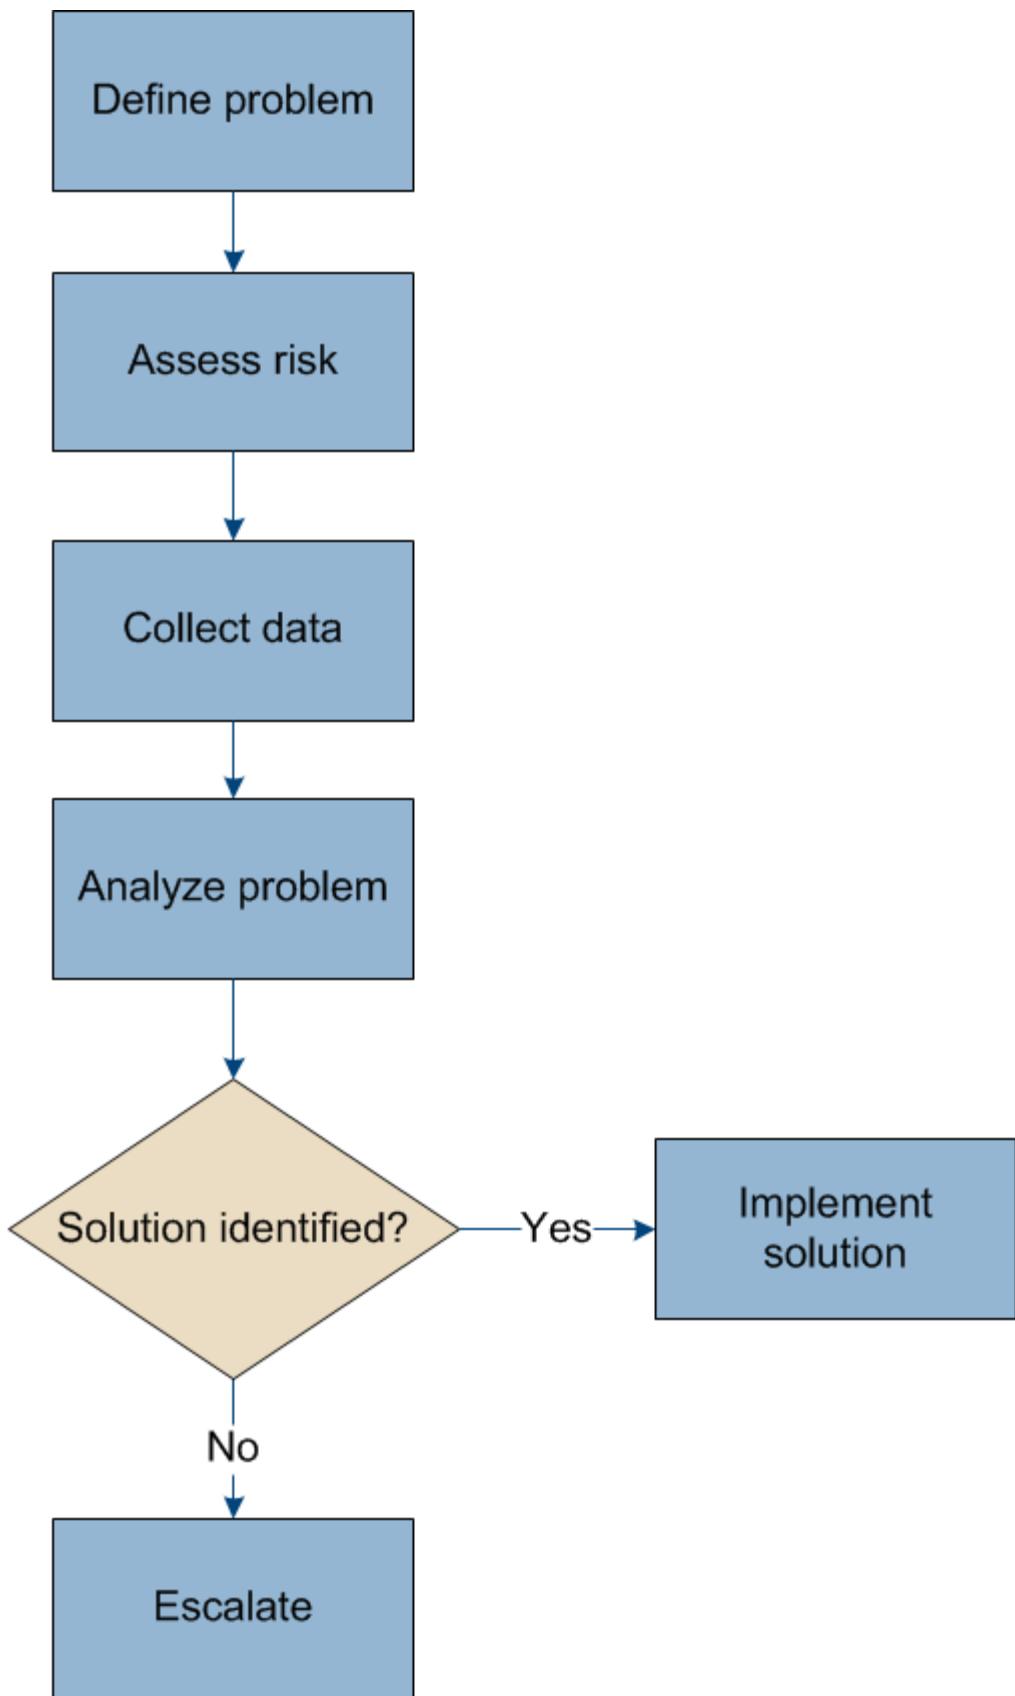
Error number	Error code	Meaning
129	EKEYREJECTED	Key was rejected by service
130	EOWNERDEAD	For robust mutexes: Owner died
131	ENOTRECOVERABLE	For robust mutexes: State not recoverable

Troubleshoot 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
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 Viewing legacy alarms Viewing resolved alerts Reviewing historical alarms and alarm frequency (legacy system)
Monitor events	Events include any system error or fault events for a node, including errors such as network errors. Monitor events to learn more about issues or to help with troubleshooting.	<ul style="list-style-type: none"> Viewing the Events tab Monitoring events
Identify trends, using chart and text reports	Trends can provide valuable clues about when issues first appeared, and can help you understand how quickly things are changing.	<ul style="list-style-type: none"> Using charts and reports
Establish baselines	Collect information about the normal levels of various operational values. These baseline values, and deviations from these baselines, can provide valuable clues.	<ul style="list-style-type: none"> Establishing baselines
Perform ingest and retrieval tests	To troubleshoot performance issues with ingest and retrieval, use a workstation to store and retrieve objects. Compare results against those seen when using the client application.	<ul style="list-style-type: none"> Monitoring PUT and GET performance
Review audit messages	Review audit messages to follow StorageGRID operations in detail. The details in audit messages can be useful for troubleshooting many types of issues, including performance issues.	<ul style="list-style-type: none"> Reviewing audit messages
Check object locations and storage integrity	If you are having storage problems, verify that objects are being placed where you expect. Check the integrity of object data on a Storage Node.	Monitoring object verification operations.

Type of data to collect	Why collect this data	Instructions
Collect data for technical support	Technical support might ask you to collect data or review specific information to help troubleshoot issues.	<ul style="list-style-type: none"> • Collecting log files and system data • Manually triggering an AutoSupport message • Reviewing support metrics

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? 	<p>What happened? What did you do?</p> <p>Document any relevant details about the change. For example:</p> <ul style="list-style-type: none"> • Details of the network changes. • Which hotfix was installed. • How client workloads changed. <p>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?</p>

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 Percent consumed/day	<p>Go to the Grid Manager. On the Nodes page, select the entire grid or a site and go to the Storage tab.</p> <p>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</p> <p>You can collect this information for the entire system or for a specific data center.</p>
Average metadata consumption	GB consumed/day Percent consumed/day	<p>Go to the Grid Manager. On the Nodes page, select the entire grid or a site and go to the Storage tab.</p> <p>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</p> <p>You can collect this information for the entire system or for a specific data center.</p>

Property	Value	How to obtain
Rate of S3/Swift operations	Operations/second	<p>Go to the Dashboard in the Grid Manager. In the Protocol Operations section, view the values for S3 rate and the Swift rate.</p> <p>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.</p>
Failed S3/Swift operations	Operations	Select Support > Tools > Grid Topology . On the Overview tab in the API Operations section, view the value for S3 Operations - Failed or Swift Operations - Failed.
ILM evaluation rate	Objects/second	<p>From the Nodes page, select grid > ILM.</p> <p>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.</p>
ILM scan rate	Objects/second	<p>Select Nodes > grid > ILM.</p> <p>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.</p>
Objects queued from client operations	Objects/second	<p>Select Nodes > grid > ILM.</p> <p>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.</p>
Average query latency	Milliseconds	Select Nodes > Storage Node > 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	<p>What are the problem symptoms? When did the problem start? Does it happen consistently or intermittently? If intermittently, what times has it occurred?</p> <p>Defining the problem</p>
	Impact assessment	<p>What is the severity of the problem? What is the impact to the client application?</p> <ul style="list-style-type: none"> • Has the client connected successfully before? • Can the client ingest, retrieve, and delete data?
	StorageGRID System ID	<p>Select Maintenance > System > License. The StorageGRID System ID is shown as part of the current license.</p>
	Software version	<p>Click Help > About to see the StorageGRID version.</p>
	Customization	<p>Summarize how your StorageGRID system is configured. For example, list the following:</p> <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?
	Log files and system data	<p>Collect log files and system data for your system. Select Support > Tools > 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</p>
	Baseline information	<p>Collect baseline information regarding ingest operations, retrieval operations, and storage consumption.</p> <p>Establishing baselines</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</p>

Item	Notes
History of efforts to diagnose the issue	If you have taken steps to diagnose or troubleshoot the issue yourself, make sure to record the steps you took and the outcome.

Related information

[Administer 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.

What you'll need

- 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 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",  
  "PHID": "FEA0AE51-534A-11EA-9FCD-31FF00C36056",  
  "PPTH": "source",  
  "META": {  
    "BASE": {  
      "PAWS": "2",  
      "PAWS2": "2",  
      "PAWS3": "2",  
      "PAWS4": "2",  
      "PAWS5": "2",  
      "PAWS6": "2",  
      "PAWS7": "2",  
      "PAWS8": "2",  
      "PAWS9": "2",  
      "PAWS10": "2",  
      "PAWS11": "2",  
      "PAWS12": "2",  
      "PAWS13": "2",  
      "PAWS14": "2",  
      "PAWS15": "2",  
      "PAWS16": "2",  
      "PAWS17": "2",  
      "PAWS18": "2",  
      "PAWS19": "2",  
      "PAWS20": "2",  
      "PAWS21": "2",  
      "PAWS22": "2",  
      "PAWS23": "2",  
      "PAWS24": "2",  
      "PAWS25": "2",  
      "PAWS26": "2",  
      "PAWS27": "2",  
      "PAWS28": "2",  
      "PAWS29": "2",  
      "PAWS30": "2",  
      "PAWS31": "2",  
      "PAWS32": "2",  
      "PAWS33": "2",  
      "PAWS34": "2",  
      "PAWS35": "2",  
      "PAWS36": "2",  
      "PAWS37": "2",  
      "PAWS38": "2",  
      "PAWS39": "2",  
      "PAWS40": "2",  
      "PAWS41": "2",  
      "PAWS42": "2",  
      "PAWS43": "2",  
      "PAWS44": "2",  
      "PAWS45": "2",  
      "PAWS46": "2",  
      "PAWS47": "2",  
      "PAWS48": "2",  
      "PAWS49": "2",  
      "PAWS50": "2",  
      "PAWS51": "2",  
      "PAWS52": "2",  
      "PAWS53": "2",  
      "PAWS54": "2",  
      "PAWS55": "2",  
      "PAWS56": "2",  
      "PAWS57": "2",  
      "PAWS58": "2",  
      "PAWS59": "2",  
      "PAWS60": "2",  
      "PAWS61": "2",  
      "PAWS62": "2",  
      "PAWS63": "2",  
      "PAWS64": "2",  
      "PAWS65": "2",  
      "PAWS66": "2",  
      "PAWS67": "2",  
      "PAWS68": "2",  
      "PAWS69": "2",  
      "PAWS70": "2",  
      "PAWS71": "2",  
      "PAWS72": "2",  
      "PAWS73": "2",  
      "PAWS74": "2",  
      "PAWS75": "2",  
      "PAWS76": "2",  
      "PAWS77": "2",  
      "PAWS78": "2",  
      "PAWS79": "2",  
      "PAWS80": "2",  
      "PAWS81": "2",  
      "PAWS82": "2",  
      "PAWS83": "2",  
      "PAWS84": "2",  
      "PAWS85": "2",  
      "PAWS86": "2",  
      "PAWS87": "2",  
      "PAWS88": "2",  
      "PAWS89": "2",  
      "PAWS90": "2",  
      "PAWS91": "2",  
      "PAWS92": "2",  
      "PAWS93": "2",  
      "PAWS94": "2",  
      "PAWS95": "2",  
      "PAWS96": "2",  
      "PAWS97": "2",  
      "PAWS98": "2",  
      "PAWS99": "2",  
      "PAWS100": "2",  
      "PAWS101": "2",  
      "PAWS102": "2",  
      "PAWS103": "2",  
      "PAWS104": "2",  
      "PAWS105": "2",  
      "PAWS106": "2",  
      "PAWS107": "2",  
      "PAWS108": "2",  
      "PAWS109": "2",  
      "PAWS110": "2",  
      "PAWS111": "2",  
      "PAWS112": "2",  
      "PAWS113": "2",  
      "PAWS114": "2",  
      "PAWS115": "2",  
      "PAWS116": "2",  
      "PAWS117": "2",  
      "PAWS118": "2",  
      "PAWS119": "2",  
      "PAWS120": "2",  
      "PAWS121": "2",  
      "PAWS122": "2",  
      "PAWS123": "2",  
      "PAWS124": "2",  
      "PAWS125": "2",  
      "PAWS126": "2",  
      "PAWS127": "2",  
      "PAWS128": "2",  
      "PAWS129": "2",  
      "PAWS130": "2",  
      "PAWS131": "2",  
      "PAWS132": "2",  
      "PAWS133": "2",  
      "PAWS134": "2",  
      "PAWS135": "2",  
      "PAWS136": "2",  
      "PAWS137": "2",  
      "PAWS138": "2",  
      "PAWS139": "2",  
      "PAWS140": "2",  
      "PAWS141": "2",  
      "PAWS142": "2",  
      "PAWS143": "2",  
      "PAWS144": "2",  
      "PAWS145": "2",  
      "PAWS146": "2",  
      "PAWS147": "2",  
      "PAWS148": "2",  
      "PAWS149": "2",  
      "PAWS150": "2",  
      "PAWS151": "2",  
      "PAWS152": "2",  
      "PAWS153": "2",  
      "PAWS154": "2",  
      "PAWS155": "2",  
      "PAWS156": "2",  
      "PAWS157": "2",  
      "PAWS158": "2",  
      "PAWS159": "2",  
      "PAWS160": "2",  
      "PAWS161": "2",  
      "PAWS162": "2",  
      "PAWS163": "2",  
      "PAWS164": "2",  
      "PAWS165": "2",  
      "PAWS166": "2",  
      "PAWS167": "2",  
      "PAWS168": "2",  
      "PAWS169": "2",  
      "PAWS170": "2",  
      "PAWS171": "2",  
      "PAWS172": "2",  
      "PAWS173": "2",  
      "PAWS174": "2",  
      "PAWS175": "2",  
      "PAWS176": "2",  
      "PAWS177": "2",  
      "PAWS178": "2",  
      "PAWS179": "2",  
      "PAWS180": "2",  
      "PAWS181": "2",  
      "PAWS182": "2",  
      "PAWS183": "2",  
      "PAWS184": "2",  
      "PAWS185": "2",  
      "PAWS186": "2",  
      "PAWS187": "2",  
      "PAWS188": "2",  
      "PAWS189": "2",  
      "PAWS190": "2",  
      "PAWS191": "2",  
      "PAWS192": "2",  
      "PAWS193": "2",  
      "PAWS194": "2",  
      "PAWS195": "2",  
      "PAWS196": "2",  
      "PAWS197": "2",  
      "PAWS198": "2",  
      "PAWS199": "2",  
      "PAWS200": "2",  
      "PAWS201": "2",  
      "PAWS202": "2",  
      "PAWS203": "2",  
      "PAWS204": "2",  
      "PAWS205": "2",  
      "PAWS206": "2",  
      "PAWS207": "2",  
      "PAWS208": "2",  
      "PAWS209": "2",  
      "PAWS210": "2",  
      "PAWS211": "2",  
      "PAWS212": "2",  
      "PAWS213": "2",  
      "PAWS214": "2",  
      "PAWS215": "2",  
      "PAWS216": "2",  
      "PAWS217": "2",  
      "PAWS218": "2",  
      "PAWS219": "2",  
      "PAWS220": "2",  
      "PAWS221": "2",  
      "PAWS222": "2",  
      "PAWS223": "2",  
      "PAWS224": "2",  
      "PAWS225": "2",  
      "PAWS226": "2",  
      "PAWS227": "2",  
      "PAWS228": "2",  
      "PAWS229": "2",  
      "PAWS230": "2",  
      "PAWS231": "2",  
      "PAWS232": "2",  
      "PAWS233": "2",  
      "PAWS234": "2",  
      "PAWS235": "2",  
      "PAWS236": "2",  
      "PAWS237": "2",  
      "PAWS238": "2",  
      "PAWS239": "2",  
      "PAWS240": "2",  
      "PAWS241": "2",  
      "PAWS242": "2",  
      "PAWS243": "2",  
      "PAWS244": "2",  
      "PAWS245": "2",  
      "PAWS246": "2",  
      "PAWS247": "2",  
      "PAWS248": "2",  
      "PAWS249": "2",  
      "PAWS250": "2",  
      "PAWS251": "2",  
      "PAWS252": "2",  
      "PAWS253": "2",  
      "PAWS254": "2",  
      "PAWS255": "2",  
      "PAWS256": "2",  
      "PAWS257": "2",  
      "PAWS258": "2",  
      "PAWS259": "2",  
      "PAWS260": "2",  
      "PAWS261": "2",  
      "PAWS262": "2",  
      "PAWS263": "2",  
      "PAWS264": "2",  
      "PAWS265": "2",  
      "PAWS266": "2",  
      "PAWS267": "2",  
      "PAWS268": "2",  
      "PAWS269": "2",  
      "PAWS270": "2",  
      "PAWS271": "2",  
      "PAWS272": "2",  
      "PAWS273": "2",  
      "PAWS274": "2",  
      "PAWS275": "2",  
      "PAWS276": "2",  
      "PAWS277": "2",  
      "PAWS278": "2",  
      "PAWS279": "2",  
      "PAWS280": "2",  
      "PAWS281": "2",  
      "PAWS282": "2",  
      "PAWS283": "2",  
      "PAWS284": "2",  
      "PAWS285": "2",  
      "PAWS286": "2",  
      "PAWS287": "2",  
      "PAWS288": "2",  
      "PAWS289": "2",  
      "PAWS290": "2",  
      "PAWS291": "2",  
      "PAWS292": "2",  
      "PAWS293": "2",  
      "PAWS294": "2",  
      "PAWS295": "2",  
      "PAWS296": "2",  
      "PAWS297": "2",  
      "PAWS298": "2",  
      "PAWS299": "2",  
      "PAWS300": "2",  
      "PAWS301": "2",  
      "PAWS302": "2",  
      "PAWS303": "2",  
      "PAWS304": "2",  
      "PAWS305": "2",  
      "PAWS306": "2",  
      "PAWS307": "2",  
      "PAWS308": "2",  
      "PAWS309": "2",  
      "PAWS310": "2",  
      "PAWS311": "2",  
      "PAWS312": "2",  
      "PAWS313": "2",  
      "PAWS314": "2",  
      "PAWS315": "2",  
      "PAWS316": "2",  
      "PAWS317": "2",  
      "PAWS318": "2",  
      "PAWS319": "2",  
      "PAWS320": "2",  
      "PAWS321": "2",  
      "PAWS322": "2",  
      "PAWS323": "2",  
      "PAWS324": "2",  
      "PAWS325": "2",  
      "PAWS326": "2",  
      "PAWS327": "2",  
      "PAWS328": "2",  
      "PAWS329": "2",  
      "PAWS330": "2",  
      "PAWS331": "2",  
      "PAWS332": "2",  
      "PAWS333": "2",  
      "PAWS334": "2",  
      "PAWS335": "2",  
      "PAWS336": "2",  
      "PAWS337": "2",  
      "PAWS338": "2",  
      "PAWS339": "2",  
      "PAWS340": "2",  
      "PAWS341": "2",  
      "PAWS342": "2",  
      "PAWS343": "2",  
      "PAWS344": "2",  
      "PAWS345": "2",  
      "PAWS346": "2",  
      "PAWS347": "2",  
      "PAWS348": "2",  
      "PAWS349": "2",  
      "PAWS350": "2",  
      "PAWS351": "2",  
      "PAWS352": "2",  
      "PAWS353": "2",  
      "PAWS354": "2",  
      "PAWS355": "2",  
      "PAWS356": "2",  
      "PAWS357": "2",  
      "PAWS358": "2",  
      "PAWS359": "2",  
      "PAWS360": "2",  
      "PAWS361": "2",  
      "PAWS362": "2",  
      "PAWS363": "2",  
      "PAWS364": "2",  
      "PAWS365": "2",  
      "PAWS366": "2",  
      "PAWS367": "2",  
      "PAWS368": "2",  
      "PAWS369": "2",  
      "PAWS370": "2",  
      "PAWS371": "2",  
      "PAWS372": "2",  
      "PAWS373": "2",  
      "PAWS374": "2",  
      "PAWS375": "2",  
      "PAWS376": "2",  
      "PAWS377": "2",  
      "PAWS378": "2",  
      "PAWS379": "2",  
      "PAWS380": "2",  
      "PAWS381": "2",  
      "PAWS382": "2",  
      "PAWS383": "2",  
      "PAWS384": "2",  
      "PAWS385": "2",  
      "PAWS386": "2",  
      "PAWS387": "2",  
      "PAWS388": "2",  
      "PAWS389": "2",  
      "PAWS390": "2",  
      "PAWS391": "2",  
      "PAWS392": "2",  
      "PAWS393": "2",  
      "PAWS394": "2",  
      "PAWS395": "2",  
      "PAWS396": "2",  
      "PAWS397": "2",  
      "PAWS398": "2",  
      "PAWS399": "2",  
      "PAWS400": "2",  
      "PAWS401": "2",  
      "PAWS402": "2",  
      "PAWS403": "2",  
      "PAWS404": "2",  
      "PAWS405": "2",  
      "PAWS406": "2",  
      "PAWS407": "2",  
      "PAWS408": "2",  
      "PAWS409": "2",  
      "PAWS410": "2",  
      "PAWS411": "2",  
      "PAWS412": "2",  
      "PAWS413": "2",  
      "PAWS414": "2",  
      "PAWS415": "2",  
      "PAWS416": "2",  
      "PAWS417": "2",  
      "PAWS418": "2",  
      "PAWS419": "2",  
      "PAWS420": "2",  
      "PAWS421": "2",  
      "PAWS422": "2",  
      "PAWS423": "2",  
      "PAWS424": "2",  
      "PAWS425": "2",  
      "PAWS426": "2",  
      "PAWS427": "2",  
      "PAWS428": "2",  
      "PAWS429": "2",  
      "PAWS430": "2",  
      "PAWS431": "2",  
      "PAWS432": "2",  
      "PAWS433": "2",  
      "PAWS434": "2",  
      "PAWS435": "2",  
      "PAWS436": "2",  
      "PAWS437": "2",  
      "PAWS438": "2",  
      "PAWS439": "2",  
      "PAWS440": "2",  
      "PAWS441": "2",  
      "PAWS442": "2",  
      "PAWS443": "2",  
      "PAWS444": "2",  
      "PAWS445": "2",  
      "PAWS446": "2",  
      "PAWS447": "2",  
      "PAWS448": "2",  
      "PAWS449": "2",  
      "PAWS450": "2",  
      "PAWS451": "2",  
      "PAWS452": "2",  
      "PAWS453": "2",  
      "PAWS454": "2",  
      "PAWS455": "2",  
      "PAWS456": "2",  
      "PAWS457": "2",  
      "PAWS458": "2",  
      "PAWS459": "2",  
      "PAWS460": "2",  
      "PAWS461": "2",  
      "PAWS462": "2",  
      "PAWS463": "2",  
      "PAWS464": "2",  
      "PAWS465": "2",  
      "PAWS466": "2",  
      "PAWS467": "2",  
      "PAWS468": "2",  
      "PAWS469": "2",  
      "PAWS470": "2",  
      "PAWS471": "2",  
      "PAWS472": "2",  
      "PAWS473": "2",  
      "PAWS474": "2",  
      "PAWS475": "2",  
      "PAWS476": "2",  
      "PAWS477": "2",  
      "PAWS478": "2",  
      "PAWS479": "2",  
      "PAWS480": "2",  
      "PAWS481": "2",  
      "PAWS482": "2",  
      "PAWS483": "2",  
      "PAWS484": "2",  
      "PAWS485": "2",  
      "PAWS486": "2",  
      "PAWS487": "2",  
      "PAWS488": "2",  
      "PAWS489": "2",  
      "PAWS490": "2",  
      "PAWS491": "2",  
      "PAWS492": "2",  
      "PAWS493": "2",  
      "PAWS494": "2",  
      "PAWS495": "2",  
      "PAWS496": "2",  
      "PAWS497": "2",  
      "PAWS498": "2",  
      "PAWS499": "2",  
      "PAWS500": "2",  
      "PAWS501": "2",  
      "PAWS502": "2",  
      "PAWS503": "2",  
      "PAWS504": "2",  
      "PAWS505": "2",  
      "PAWS506": "2",  
      "PAWS507": "2",  
      "PAWS508": "2",  
      "PAWS509": "2",  
      "PAWS510": "2",  
      "PAWS511": "2",  
      "PAWS512": "2",  
      "PAWS513": "2",  
      "PAWS514": "2",  
      "PAWS515": "2",  
      "PAWS516": "2",  
      "PAWS517": "2",  
      "PAWS518": "2",  
      "PAWS519": "2",  
      "PAWS520": "2",  
      "PAWS521": "2",  
      "PAWS522": "2",  
      "PAWS523": "2",  
      "PAWS524": "2",  
      "PAWS525": "2",  
      "PAWS526": "2",  
      "PAWS527": "2",  
      "PAWS528": "2",  
      "PAWS529": "2",  
      "PAWS530": "2",  
      "PAWS531": "2",  
      "PAWS532": "2",  
      "PAWS533": "2",  
      "PAWS534": "2",  
      "PAWS535": "2",  
      "PAWS536": "2",  
      "PAWS537": "2",  
      "PAWS538": "2",  
      "PAWS539": "2",  
      "PAWS540": "2",  
      "PAWS541": "2",  
      "PAWS542": "2",  
      "PAWS543": "2",  
      "PAWS544": "2",  
      "PAWS545": "2",  
      "PAWS546": "2",  
      "PAWS547": "2",  
      "PAWS548": "2",  
      "PAWS549": "2",  
      "PAWS550": "2",  
      "PAWS551": "2",  
      "PAWS552": "2",  
      "PAWS553": "2",  
      "PAWS554": "2",  
      "PAWS555": "2",  
      "PAWS556": "2",  
      "PAWS557": "2",  
      "PAWS558": "2",  
      "PAWS559": "2",  
      "PAWS560": "2",  
      "PAWS561": "2",  
      "PAWS562": "2",  
      "PAWS563": "2",  
      "PAWS564": "2",  
      "PAWS565": "2",  
      "PAWS566": "2",  
      "PAWS567": "2",  
      "PAWS568": "2",  
      "PAWS569": "2",  
      "PAWS570": "2",  
      "PAWS571": "2",  
      "PAWS572": "2",  
      "PAWS573": "2",  
      "PAWS574": "2",  
      "PAWS575": "2",  
      "PAWS576": "2",  
      "PAWS577": "2",  
      "PAWS578": "2",  
      "PAWS579": "2",  
      "PAWS580": "2",  
      "PAWS581": "2",  
      "PAWS582": "2",  
      "PAWS583": "2",  
      "PAWS584": "2",  
      "PAWS585": "2",  
      "PAWS586": "2",  
      "PAWS587": "2",  
      "PAWS588": "2",  
      "PAWS589": "2",  
      "PAWS590": "2",  
      "PAWS591": "2",  
      "PAWS592": "2",  
      "PAWS593": "2",  
      "PAWS594": "2",  
      "PAWS595": "2",  
      "PAWS596": "2",  
      "PAWS597": "2",  
      "PAWS598": "2",  
      "PAWS599": "2",  
      "PAWS600": "2",  
      "PAWS601": "2",  
      "PAWS602": "2",  
      "PAWS603": "2",  
      "PAWS604": "2",  
      "PAWS605": "2",  
      "PAWS606": "2",  
      "PAWS607": "2",  
      "PAWS608": "2",  
      "PAWS609": "2",  
      "PAWS610": "2",  
      "PAWS611": "2",  
      "PAWS612": "2",  
      "PAWS613": "2",  
      "PAWS614": "2",  
      "PAWS615": "2",  
      "PAWS616": "2",  
      "PAWS617": "2",  
      "PAWS618": "2",  
      "PAWS619": "2",  
      "PAWS620": "2",  
      "PAWS621": "2",  
      "PAWS622": "2",  
      "PAWS623": "2",  
      "PAWS624": "2",  
      "PAWS625": "2",  
      "PAWS626": "2",  
      "PAWS627": "2",  
      "PAWS628": "2",  
      "PAWS629": "2",  
      "PAWS630": "2",  
      "PAWS631": "2",  
      "PAWS632": "2",  
      "PAWS633": "2",  
      "PAWS634": "2",  
      "PAWS635": "2",  
      "PAWS636": "2",  
      "PAWS637": "2",  
      "PAWS638": "2",  
      "PAWS639": "2",  
      "PAWS640": "2",  
      "PAWS641": "2",  
      "PAWS642": "2",  
      "PAWS643": "2",  
      "PAWS644": "2",  
      "PAWS645": "2",  
      "PAWS646": "2",  
      "PAWS647": "2",  
      "PAWS648": "2",  
      "PAWS649": "2",  
      "PAWS650": "2",  
      "PAWS651": "2",  
      "PAWS652": "2",  
      "PAWS653": "2",  
      "PAWS654": "2",  
      "PAWS655": "2",  
      "PAWS656": "2",  
      "PAWS657": "2",  
      "PAWS658": "2",  
      "PAWS659": "2",  
      "PAWS660": "2",  
      "PAWS661": "2",  
      "PAWS662": "2",  
      "PAWS663": "2",  
      "PAWS664": "2",  
      "PAWS665": "2",  
      "PAWS666": "2",  
      "PAWS667": "2",  
      "PAWS668": "2",  
      "PAWS669": "2",  
      "PAWS670": "2",  
      "PAWS671": "2",  
      "PAWS672": "2",  
      "PAWS673": "2",  
      "PAWS674": "2",  
      "PAWS675": "2",  
      "PAWS676": "2",  
      "PAWS677": "2",  
      "PAWS678": "2",  
      "PAWS679": "2",  
      "PAWS680": "2",  
      "PAWS681": "2",  
      "PAWS682": "2",  
      "PAWS683": "2",  
      "PAWS684": "2",  
      "PAWS685": "2",  
      "PAWS686": "2",  
      "PAWS687": "2",  
      "PAWS688": "2",  
      "PAWS689": "2",  
      "PAWS690": "2",  
      "PAWS691": "2",  
      "PAWS692": "2",  
      "PAWS693": "2",  
      "PAWS694": "2",  
      "PAWS695": "2",  
      "PAWS696": "2",  
      "PAWS697": "2",  
      "PAWS698": "2",  
      "PAWS699": "2",  
      "PAWS700": "2",  
      "PAWS701": "2",  
      "PAWS702": "2",  
      "PAWS703": "2",  
      "PAWS704": "2",  
      "PAWS705": "2",  
      "PAWS706": "2",  
      "PAWS707": "2",  
      "PAWS708": "2",  
      "PAWS709": "2",  
      "PAWS710": "2",  
      "PAWS711": "2",  
      "PAWS712": "2",  
      "PAWS713": "2",  
      "PAWS714": "2",  
      "PAWS715": "2",  
      "PAWS716": "2",  
      "PAWS717": "2",  
      "PAWS718": "2",  
      "PAWS719": "2",  
      "PAWS720": "2",  
      "PAWS721": "2",  
      "PAWS722":
```

Related information

[Manage objects with ILM](#)

[Use S3](#)

[Use Swift](#)

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 > Tools > 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

[Maintain & recover](#)

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.

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.
- **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.

 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.

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.

Changing 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.

What you'll need

- 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 > Tools > Grid Topology**.
2. Select **Storage Node > LDR > Verification**.
3. Select **Configuration > Main**.
4. Go to **LDR > Verification > Configuration > Main**.
5. Under Background Verification, select **Verification Rate > High** or **Verification Rate > Adaptive**.

Configuration: LDR (DC2-S1-106-147) - Verification
Updated: 2019-04-24 16:13:44 PDT

Reset Missing Objects Count

Foreground Verification

ID	Verify
0	<input type="checkbox"/>
1	<input type="checkbox"/>
2	<input type="checkbox"/>

Background Verification

Verification Rate: Adaptive

Reset Corrupt Objects Count

Quarantined Objects

Delete Quarantined Objects

[Apply Changes](#)



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 > Tools > Grid Topology**.
- b. Select **Storage Node > LDR > Erasure Coding**.
- c. In the Verification Results table, monitor the Corrupt Fragments Detected (ECCD) attribute.

9. After corrupt objects have been automatically restored by the StorageGRID system, reset the count of corrupt objects.

 - a. Select **Support > Tools > Grid Topology**.
 - b. Select **Storage Node > LDR > Verification > Configuration**.
 - c. Select **Reset Corrupt Object Count**.
 - d. Click **Apply Changes**.

10. If you are confident that quarantined objects are not required, you can delete them.



If the **Objects lost** alert or the LOST (Lost Objects) legacy alarm was triggered, technical support might want to access quarantined objects to help debug the underlying issue or to attempt data recovery.

- a. Select **Support > Tools > Grid Topology**.
- b. Select **Storage Node > LDR > Verification > Configuration**.
- c. Select **Delete Quarantined Objects**.
- d. 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.

What you'll need

- 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 > Tools > 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.

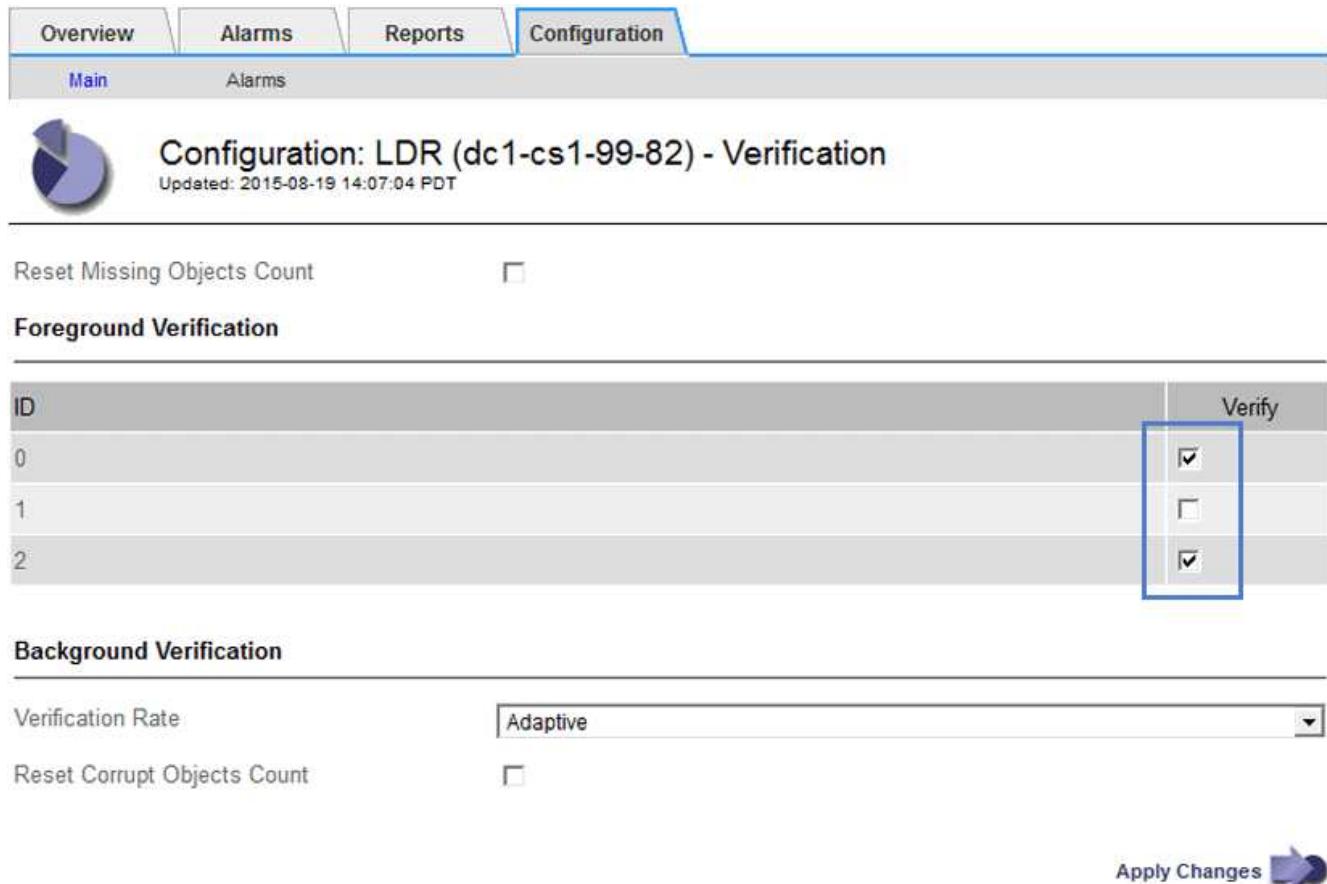
Foreground verification generates an LDR Foreground Verification grid task that, depending on the number of objects stored on a Storage Node, can take days or weeks to complete. It is possible to select multiple Storage Nodes at the same time; however, these grid tasks are not run simultaneously. Instead, 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. Select **Configuration > Main**.
3. Under **Foreground Verification**, select the check box for each storage volume ID you want to verify.



ID	Verify
0	<input checked="" type="checkbox"/>
1	<input type="checkbox"/>
2	<input checked="" type="checkbox"/>

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 > Tools > 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

[Maintain & recover](#)

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 information

[Investigating lost objects](#)

[Resetting lost and missing object counts](#)

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.

What you'll need

- 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.

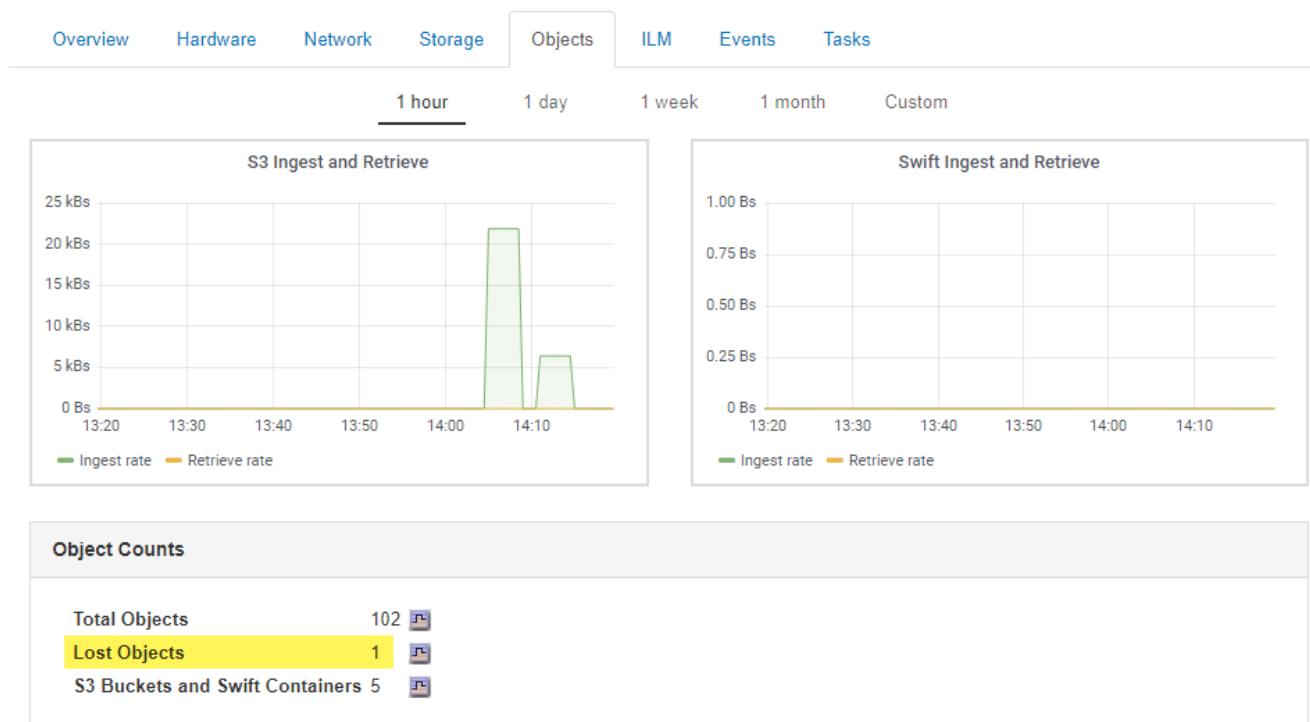
Investigate lost object alarms or alerts immediately. You might need to take action to prevent further data loss. In some cases, you might be able to restore a lost object if you take prompt action.

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] [A
MID(FC32):ILMX] [ATID(UI64):7729403978647354233]
```

5. Use the `ObjectByUUID` command to find the object by its identifier (UUID), and then determine if data is

at risk.

- a. Telnet to localhost 1402 to access the LDR console.
- b. Enter: /proc/OBRP/ObjectByUUID UUID_value

In this first example, the object with UUID 926026C4-00A4-449B-AC72-BCCA72DD1311 has two locations listed.

```
ade 12448208: /proc/OBRP > ObjectByUUID 926026C4-00A4-449B-AC72-BCCA72DD1311

{
    "TYPE(Object Type)": "Data object",
    "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
    "NAME": "cats",
    "CBID": "0x38186FE53E3C49A5",
    "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-ACBB00BB82DD",
    "PPTH(Parent path)": "source",
    "META": {
        "BASE(Protocol metadata)": {
            "PAWS(S3 protocol version)": "2",
            "ACCT(S3 account ID)": "44084621669730638018",
            "*ctp(HTTP content MIME type)": "binary/octet-stream"
        },
        "BYCB(System metadata)": {
            "CSIZ(Plaintext object size)": "5242880",
            "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
            "BSIZ(Content block size)": "5252084",
            "CVER(Content block version)": "196612",
            "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
            "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
            "ITME": "1581534970983000"
        },
        "CMSSM": {
            "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%DkRt78Il1a\#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"
        },
        "CMSSM": {
            "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": "")	<p>If the object is not found, the message "ERROR": "" is returned.</p> <p>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.</p>
Locations > 0	<p>If there are locations listed in the output, the Lost Objects alarm might be a false positive.</p> <p>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.</p> <p>(The procedure for finding potentially lost objects explains how to use the Node ID to find the correct Storage Node.)</p> <p>Searching for and restoring potentially lost objects</p> <p>If the objects exist, you can reset the count of Lost Objects to clear the alarm and the alert.</p>
Locations = 0	<p>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.</p> <p>Searching for and restoring potentially lost objects</p> <p>Technical support might ask you to determine if there is a storage recovery procedure in progress. That is, has a <i>repair-data</i> command been issued on any Storage Node, and is the recovery still in progress? See the information about restoring object data to a storage volume in the recovery and maintenance instructions.</p>

Related information

[Maintain & recover](#)

[Review audit logs](#)

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.

What you'll need

- 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.



Contact technical support for assistance with this procedure.

Steps

1. From an Admin Node, search the audit logs for possible object locations:
 - a. Log in to the grid node:
 - i. Enter the following command: `ssh admin@grid_node_IP`
 - ii. Enter the password listed in the `Passwords.txt` file.
 - iii. Enter the following command to switch to root: `su -`
 - iv. Enter the password listed in the `Passwords.txt` file. When you are logged in as root, the prompt changes from `$` to `#`.
 - b. Change to the directory where the audit logs are located: `cd /var/local/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-valueaudit_file_name > output_file_name`

For example:

```
Admin: # grep 926026C4-00A4-449B-AC72-BCCA72DD1311 audit.log > messages_about_lost_object.txt
```

- d. Use grep to extract the Location Lost (LLST) audit messages from this output file. Enter: `grep LLST output_file_name`

For example:

```
Admin: # grep LLST messages_about_lost_objects.txt
```

An LLST audit message looks like this 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\(\FC32\):\CL  
SM]  
[ATID\(\UI64\):7086871083190743409]]
```

- e. Find the PCLD field and the NOID field in the LLST message.

If present, the value of PCLD is the complete path on disk to the missing replicated object copy. The value of NOID is the node id of the LDR where a copy of the object might be found.

If you find an object location, you might be able to restore the object.

f. Find the Storage Node 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 > Tools > 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:

- Enter the following command: `ssh admin@grid_node_IP`
- Enter the password listed in the `Passwords.txt` file.
- Enter the following command to switch to root: `su -`
- 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 [Restore the object to StorageGRID](#). 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:

- 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'`
- Telnet to localhost 1402 to access the LDR console. Enter: `telnet 0 1402`
- Enter: `cd /proc/STOR`
- 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%DkRt78I1a#3udu'
```

Continue with step [Verify that new locations were created](#)

4. If the object was successfully restored to StorageGRID, verify that new locations were created.
 - a. Enter: `cd /proc/OBRP`
 - b. Enter: `ObjectByUUID UUID_value`

The following example shows that there are two locations for the object with UUID `926026C4-00A4-449B-AC72-BCCA72DD1311`.

```
ade 12448208: /proc/OBRP > ObjectByUUID 926026C4-00A4-449B-AC72-
BCCA72DD1311

{
    "TYPE(Object Type)": "Data object",
    "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
    "NAME": "cats",
    "CBID": "0x38186FE53E3C49A5",
    "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-
ACBB00BB82DD",
    "PPTH(Parent path)": "source",
    "META": {
        "BASE(Protocol metadata)": {
            "PAWS(S3 protocol version)": "2",
            "ACCT(S3 account ID)": "44084621669730638018",
        }
    }
}
```

```

    "*ctp(HTTP content MIME type)": "binary/octet-stream"
},
"BYCB(System metadata)": {
    "CSIZ(Plaintext object size)": "5242880",
    "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
    "BSIZ(Content block size)": "5252084",
    "CVER(Content block version)": "196612",
    "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
    "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
    "ITME": "1581534970983000"
},
"CMSM": {
    "LATM(Object last access time)": "2020-02-
12T19:16:10.983000"
},
"AWS3": {
    "LOCC": "us-east-1"
},
"CLCO\ (Locations\)": \[
    \{
        "Location Type": "CLDI\ (Location online\)",
        "NOID\ (Node ID)": "12448208",
        "VOLI\ (Volume ID)": "3222345473",
        "Object File Path":
        "/var/local/rangedb/1/p/17/11/00rH0%DkRt78Il1a\#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-valueaudit_file_name > output_file_name`

For example:

```
Admin: # grep 926026C4-00A4-449B-AC72-BCCA72DD1311 audit.log >
messages_about_restored_object.txt
```

- d. Use grep to extract the Object Rules Met (ORLM) audit messages from this output file. Enter: `grep ORLM output_file_name`

For example:

```
Admin: # grep ORLM messages_about_restored_object.txt
```

An ORLM audit message looks like this 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):15633982
30669]
[ATID(UI64):15494889725796157557][ANID(UI32):13100453][AMID(FC32):BCM
S]]
```

- e. Find the LOCS field in the audit message.

If present, the value of CLDI in LOCS is the node ID and the volume ID where an object copy has been created. This message shows that the ILM has been applied and that two object copies have been created in two locations in the grid.

- f. Reset the count of lost objects in the Grid Manager.

Related information

[Investigating lost objects](#)

[Confirming object data locations](#)

[Resetting lost and missing object counts](#)

[Review audit logs](#)

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.

What you'll need

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

About this task

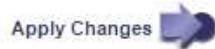
You can reset the Lost Objects counter from either of the following pages:

- **Support > Tools > Grid Topology > site > Storage Node > LDR > Data Store > Overview > Main**
- **Support > Tools > Grid Topology > site > Storage Node > DDS > Data Store > Overview > Main**

These instructions show resetting the counter from the **LDR > Data Store** page.

Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Site > Storage Node > LDR > Data Store > Configuration** for the Storage Node that has the **Objects lost** alert or the **LOST** alarm.
3. Select **Reset Lost Objects Count**.



4. Click **Apply Changes**.

The Lost Objects attribute is reset to 0 and the **Objects lost** alert and the **LOST** alarm clear, which can take a few minutes.

5. Optionally, reset other related attribute values that might have been incremented in the process of identifying the lost object.

- a. Select **Site > Storage Node > LDR > Erasure Coding > Configuration**.
- b. Select **Reset Reads Failure Count** and **Reset Corrupt Copies Detected Count**.
- c. Click **Apply Changes**.
- d. Select **Site > Storage Node > LDR > Verification > Configuration**.
- e. Select **Reset Missing Objects Count** and **Reset Corrupt Objects Count**.
- f. If you are confident that quarantined objects 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

[Administer 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.

What you'll need

- 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.



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

3. Review the details in the dialog box, and note the following:

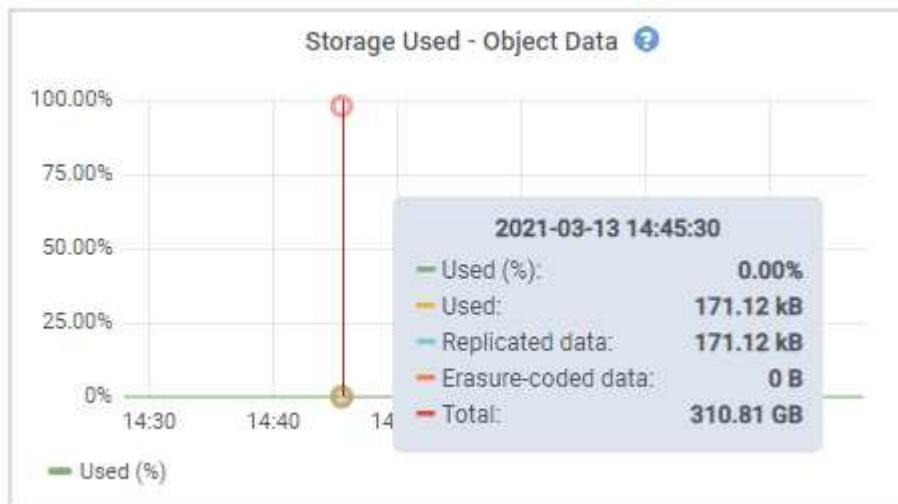
- Time triggered
- The name of the site and node
- The current values of the metrics for this alert

4. Select **Nodes > Storage Node or Site > Storage**.

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

The following values are shown:

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



6. Select the time controls above the graph to view storage use over different time periods.

Looking at storage use over time can help you understand how much storage was used before and after the alert was triggered and can help you estimate how long it might take for the node's remaining space to become full.

7. As soon as possible, 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.



To manage a full Storage Node, see the instructions for administering StorageGRID.

Related information

[Troubleshooting the Storage Status \(SSTS\) alarm](#)

[Expand your grid](#)

[Administer StorageGRID](#)

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.

What you'll need

- 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 > Alarms (legacy) > Current Alarms**.
2. From the Service column, select the data center, node, and service that are associated with the SSTS alarm.

The Grid Topology page appears. The Alarms tab shows the active alarms for the node and service you selected.

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 checked="" type="checkbox"/>

In this example, both the SSTS (Storage Status) and SAVP (Total Usable Space (Percent)) alarms have been triggered at the Notice level.



Typically, both the SSTS alarm and the SAVP alarm are triggered at about the same time; however, whether both alarms are triggered depends on the the watermark setting in GB and the SAVP alarm setting in percent.

3. To determine how much usable space is actually available, select **LDR > Storage > Overview**, and find the Total Usable Space (STAS) attribute.

Overview Alarms Reports Configuration

Main



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	

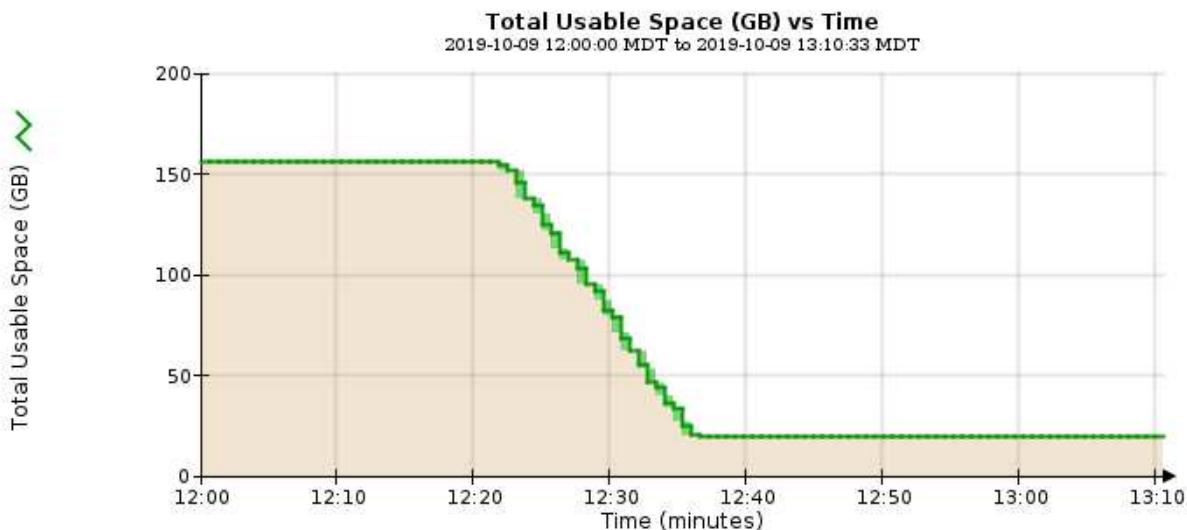
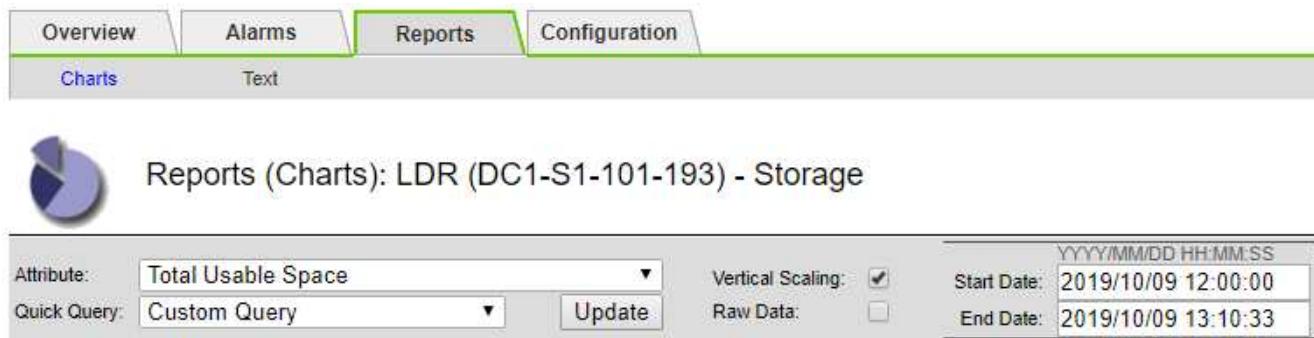
Object Store Volumes

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.

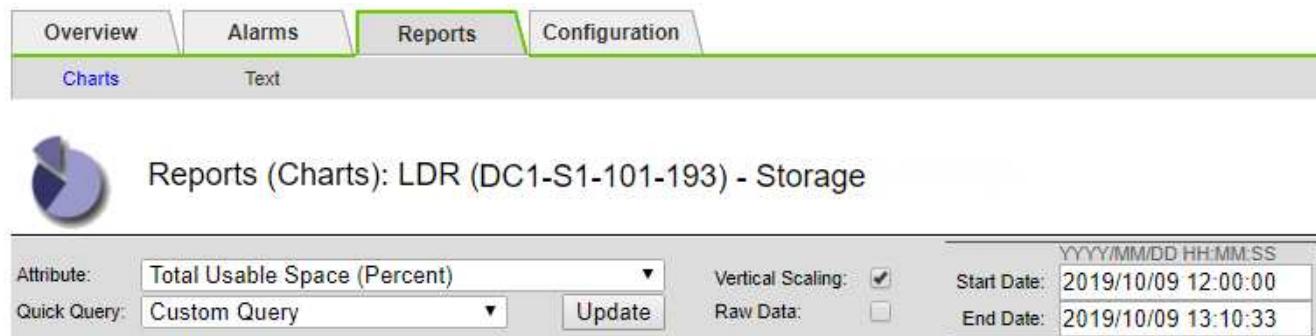
4. To understand how storage has been used over time, select the **Reports** tab, and plot Total Usable Space over the last few hours.

In this example, Total Usable Space dropped from roughly 155 GB at 12:00 to 20 GB at 12:35, which corresponds to the time at which the SSTS alarm was triggered.



5. To understand how storage is being used as a percent of the total, plot Total Usable Space (Percent) over the last few hours.

In this example, the total usable space dropped from 95% to just over 10% at approximately the same time.



6. As required, add storage capacity by expanding the StorageGRID system.

For procedures on how to manage a full Storage Node, see the instructions for administering StorageGRID.

Related information

[Expand your grid](#)

[Administer 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 information

[Administer StorageGRID](#)

[Use a tenant account](#)

[Log files reference](#)

[Resetting event counts](#)

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 is triggered, you must add new Storage Nodes.

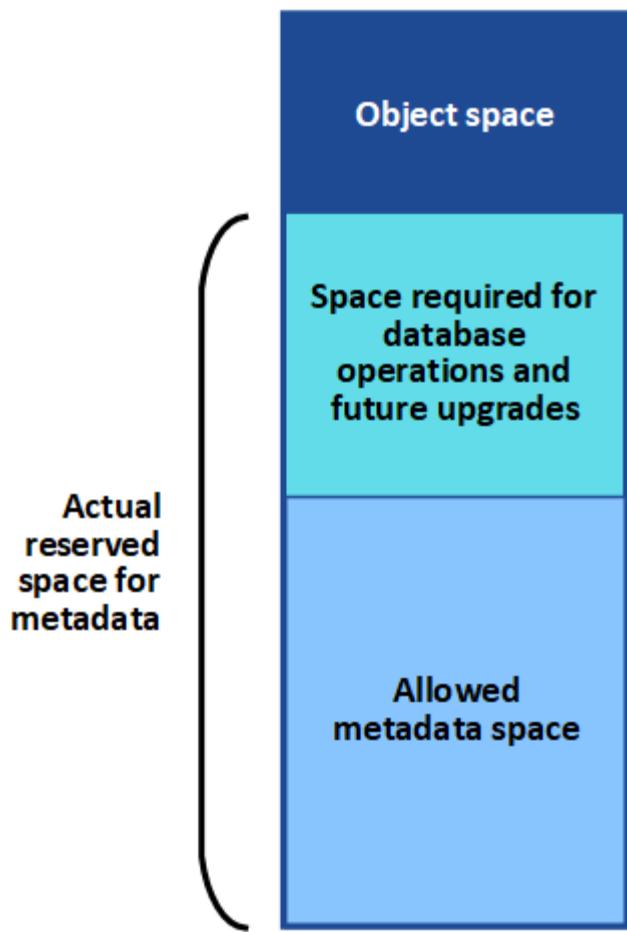
What you'll need

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

About this task

StorageGRID reserves a certain amount of space on volume 0 of each Storage Node for object metadata. This space is known as the actual reserved space, and it is subdivided into the space allowed for object metadata (the allowed metadata space) and the space required for essential database operations, such as compaction and repair. The allowed metadata space governs overall object capacity.

Volume 0



If object metadata consumes more than 100% of the space allowed for metadata, database operations cannot run efficiently and errors will occur.

StorageGRID uses the following Prometheus metric to measure how full the allowed metadata 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.

- **Minor:** Object metadata is using 70% or more of the allowed metadata space. You should add new Storage Nodes as soon as possible.
- **Major:** Object metadata is using 90% or more of the allowed metadata space. You must add new Storage Nodes immediately.



When object metadata is using 90% or more of the allowed metadata space, a warning appears on the Dashboard. If this warning appears, you must add new Storage Nodes immediately. You must never allow object metadata to use more than 100% of the allowed space.

- **Critical:** Object metadata is using 100% or more of the allowed metadata space and is starting to consume the space required for essential database operations. You must stop the ingest of new objects, and you

must add new Storage Nodes immediately.

In the following example, object metadata is using more than 100% of the allowed metadata space. This is a critical situation, which will result in inefficient database operation and errors.

The following Storage Nodes are using more than 90% of the space allowed for object metadata:

Node	% Used	Used	Allowed
DC1-S2-227	104.51%	6.73 GB	6.44 GB
DC1-S3-228	104.36%	6.72 GB	6.44 GB
DC2-S2-233	104.20%	6.71 GB	6.44 GB
DC1-S1-226	104.20%	6.71 GB	6.44 GB
DC2-S3-234	103.43%	6.66 GB	6.44 GB

Undesirable results can occur if object metadata uses more than 100% of the allowed space. You must add new Storage Nodes immediately or contact support.



If the size of volume 0 is smaller than the Metadata Reserved Space storage option (for example, in a non-production environment), the calculation for the **Low metadata storage** alert might be inaccurate.

Steps

1. Select **Alerts > Current**.
2. From the table of alerts, expand the **Low metadata storage** alert group, if required, and select the specific alert you want to view.
3. Review the details in the alert dialog box.
4. If a major or critical **Low metadata storage** alert has been triggered, perform an expansion to add Storage Nodes immediately.



Because StorageGRID keeps complete copies of all object metadata at each site, the metadata capacity of the entire grid is limited by the metadata capacity of the smallest site. If you need to add metadata capacity to one site, you should also expand any other sites by the same number of Storage Nodes.

After you perform the expansion, StorageGRID redistributes the existing object metadata to the new nodes, which increases the overall metadata capacity of the grid. No user action is required. The **Low metadata storage** alert is cleared.

Related information

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

[Expand your grid](#)

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.

What you'll need

- 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.

You can use the Diagnostics page to obtain additional information on the current state of your grid.

Running diagnostics



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

1. Select **Support > Tools > Grid Topology**.
2. Select **site > Storage Node > SSM > Services > Alarms > Main** to display alarms.

This example shows that the SVST alarm was triggered.

Severity Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge	Acknowledge Time
SVST (Services: Status - Cassandra)	Not Running	2014-08-14 14:56:28 PDT	Not Running	Not Running		

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.



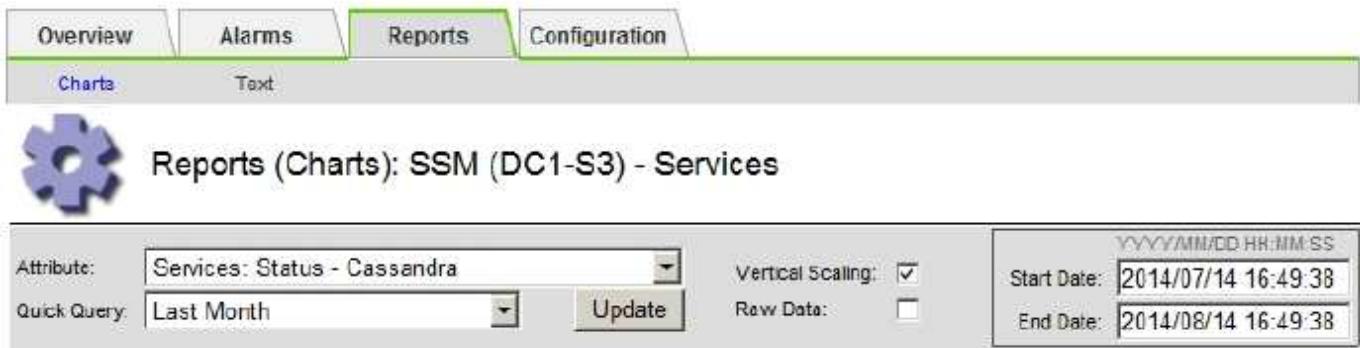
If two or more of the Cassandra database services are down, contact technical support, and do not proceed with the steps below.

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 > Tools > Grid Topology**. Then select **site > Storage Node > SSM > Services > Reports > Charts**.
- b. Select **Attribute > Service: Status - Cassandra**.
- c. For **Start Date**, enter a date that is at least 16 days before the current date. For **End Date**, enter the current date.
- d. Click **Update**.
- e. If the chart shows Cassandra as being down for more than 15 days, rebuild the Cassandra database.

The following chart example shows that Cassandra has been down for at least 17 days.



6. To review the servermanager.log file on the Storage Node:

- Log in to the grid node:
 - Enter the following command: `ssh admin@grid_node_IP`
 - Enter the password listed in the `Passwords.txt` file.
 - Enter the following command to switch to root: `su -`
 - 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
```

- Make sure the timestamp of this message is the time when you attempted restarting Cassandra as instructed in step [Restart Cassandra from the Storage Node](#).

There can be more than one entry for Cassandra; you must locate the most recent entry.

d. If Cassandra has been down for longer than 15 days, you must rebuild the Cassandra database.

For instructions, see “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

[Maintain & recover](#)

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.

You can use the Diagnostics page to obtain additional information on the current state of your grid.

[Running diagnostics](#)

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**.



To reset event counts, you must have the Grid Topology Page Configuration permission.

Related information

[Resetting event counts](#)

Troubleshooting certificate errors

If you see a security or certificate issue when you try to connect to StorageGRID using a web browser, an S3 or Swift client, or an external monitoring tool, you should check the certificate.

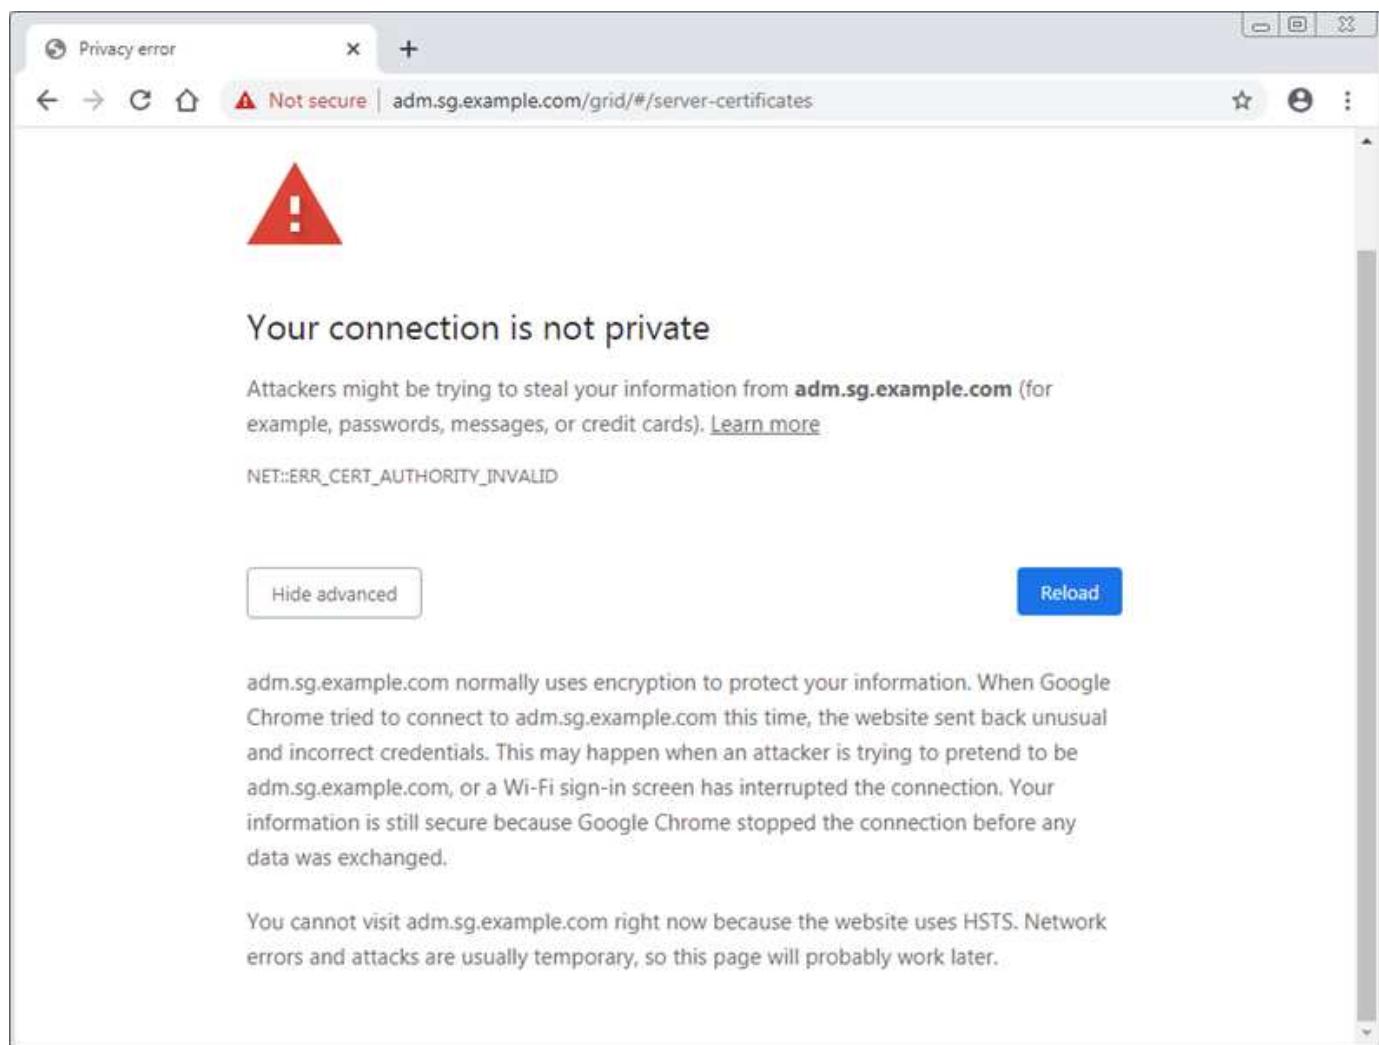
About this task

Certificate errors can cause problems when you try to connect to StorageGRID using the Grid Manager, Grid Management API, Tenant Manager, or the Tenant Management API. Certificate errors can also occur when you try to connect with an S3 or Swift client or external monitoring tool.

If you are accessing the Grid Manager or Tenant Manager using a domain name instead of an IP address, the browser shows a certificate error without an option to bypass if either of the following occurs:

- Your custom management interface 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:



To ensure that operations are not disrupted by a failed server certificate, the **Expiration of server certificate for Management Interface** alert is triggered when the server certificate is about to expire.

When you are using client certificates for external Prometheus integration, certificate errors can be caused by the StorageGRID management interface server certificate or by client certificates. The **Expiration of certificates configured on Client Certificates** page alert is triggered when a client certificate is about to expire.

Steps

1. If you received an alert notification about an expired certificate, access the certificate details:
 - For a server certificate, select **Configuration > Network Settings > Server Certificates**.
 - For a client certificate, select **Configuration > Access Control > Client Certificates**.

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 398 days.

3. If the certificate has expired or will expire soon, upload or generate a new certificate.

- For a server certificate, see the steps for configuring a custom server certificate for the Grid Manager and the Tenant Manager in the instructions for administering StorageGRID.
- For a client certificate, see the steps for configuring a client certificate in the instructions for administering StorageGRID.

4. For server certificate errors, try either or both of the following options:

- Ensure that the Subject Alternative Name (SAN) of the certificate is populated, and that the SAN matches the IP address or host name of the node that you are connecting to.
- If you are attempting to connect to StorageGRID using a domain name:
 - i. Enter the IP address of the Admin Node instead of the domain name to bypass the connection error and access the Grid Manager.
 - ii. From the Grid Manager, select **Configuration > Network Settings > Server Certificates** to install a new custom certificate or continue with the default certificate.
 - iii. In the instructions for administering StorageGRID, see the steps for configuring a custom server certificate for the Grid Manager and the Tenant Manager.

Related information

[Administer 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.

What you'll need

- 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** 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.



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:

- i. Review any displayed alarms.

- ii. Select **Configuration > Identity Federation**.

- iii. Click **Test Connection** to validate your connection settings for the LDAP server.

- iv. If the test fails, resolve any configuration errors.

- If the local user 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.

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

When you are logged in as root, the prompt changes from `$` to `#`.

13. View the status of all services running on the grid node: `storagegrid-status`

Make sure the idnt, acct, nginx, and cassandra services are all running.

14. Repeat steps [Use Lumberjack to collect logs](#) and [Review logs](#) to review the logs on the Storage Nodes.
15. If you are unable to resolve the issue, contact technical support.

Provide the logs you collected to technical support.

Related information

[Administer StorageGRID](#)

[Log files reference](#)

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.

If you experience issues with the web interface:

- Make sure you are using a supported browser.



Browser support has changed for StorageGRID 11.5. Confirm you are using a supported version.

- 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 information

[Web browser requirements](#)

[Administer 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.

What you'll need

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 > Tools > Grid Topology**.
3. Select **Site > unavailable Admin Node > SSM > Services > Overview > Main**.
4. Look for services that have a status of Not Running and that might also be displayed in blue.



Overview: SSM (MM-10-224-4-81-ADM1) - Services

Updated: 2017-01-27 11:52:51 EST

Operating System:

Linux
3.16.0-4-amd64

Services

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 Downampler 1	10.4.0-20161224.0333.803cd91	Running	22	0.049 %	195 MB
NMS Data Downampler 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

- Determine if alarms have been triggered.
- Take the appropriate actions to resolve the issue.

Related information

[Administer 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.

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.

If you see one of the listed error messages, take the recommended action.

Error message	Root cause and corrective action
<p>422: Unprocessable Entity</p> <p>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:</p> <p>LDAP Result Code 8 "Strong Auth Required": 00002028:</p> <p>LdapErr: DSID-0C090256,</p> <p>comment: The server requires binds to turn on integrity checking if SSL\TLS are not already active on the connection, data 0, v3839</p>	<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>
<p>422: Unprocessable Entity</p> <p>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</p> <p>(EOF)</p>	<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 identity 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

[Administer 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
<code>mtu</code>	The MTU to set. Must be in the range 1280 to 9216.
<code>network</code>	The networks to apply the MTU to. Include one or more of the following network types: <ul style="list-style-type: none">• <code>grid</code>• <code>admin</code>• <code>client</code>

Optional arguments	Description
<code>-h, - help</code>	Show the help message and exit.
<code>-n node, --node node</code>	The node. The default is the local node.

Related information

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

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:
 - i. Check the FEC status of the port in the switch attached to your StorageGRID appliance.
 - ii. Check the physical integrity of the cables from the appliance to the switch.
 - iii. 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.



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.



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:
 - i. Enable FEC, if not already enabled.
 - ii. Verify that your network cabling is of good quality and is not damaged or improperly connected.
 - iii. If the cables do not appear to be the problem, contact technical support.



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 > Tools > Grid Topology**.
 - b. Select **site > grid node > SSM > Resources > Configuration > Main**.
 - c. Select **Reset Receive Error Count** and click **Apply Changes**.

Related information

[Troubleshooting the Grid Network MTU mismatch alert](#)

[Alarms reference \(legacy system\)](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[SG100 & SG1000 services appliances](#)

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.



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

[Maintain & recover](#)

Linux: Network connectivity issues

You might see issues with network connectivity for StorageGRID grid nodes hosted on Linux hosts.

MAC address cloning

In some cases, network issues can be resolved by using MAC address cloning. If you are using virtual hosts, set the value of the MAC address cloning key for each of your networks to "true" in your node configuration file. This setting causes the MAC address of the StorageGRID container to use the MAC address of the host. To create node configuration files, see the instructions in the installation guide for your platform.



Create separate virtual network interfaces for use by the Linux host OS. Using the same network interfaces for the Linux host OS and the StorageGRID container might cause the host OS to become unreachable if promiscuous mode has not been enabled on the hypervisor.

For more information on enabling MAC cloning, see the instructions in the installation guide for your platform.

Promiscuous mode

If you do not want to use MAC address cloning and would rather allow all interfaces to receive and transmit data for MAC addresses other than the ones assigned by the hypervisor, ensure that the security properties at the virtual switch and port group levels are set to **Accept** for Promiscuous Mode, MAC Address Changes, and Forged Transmits. The values set on the virtual switch can be overridden by the values at the port group level, so ensure that settings are the same in both places.

Related information

[Install Red Hat Enterprise Linux or CentOS](#)

[Install Ubuntu or Debian](#)

Linux: Node status is “orphaned”

A Linux node in an orphaned state usually indicates that either the storagegrid service or the StorageGRID node daemon controlling the node's container died unexpectedly.

About this task

If a Linux node reports that it is in an orphaned state, you should:

- Check logs for errors and messages.
- Attempt to start the node again.
- If necessary, use 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

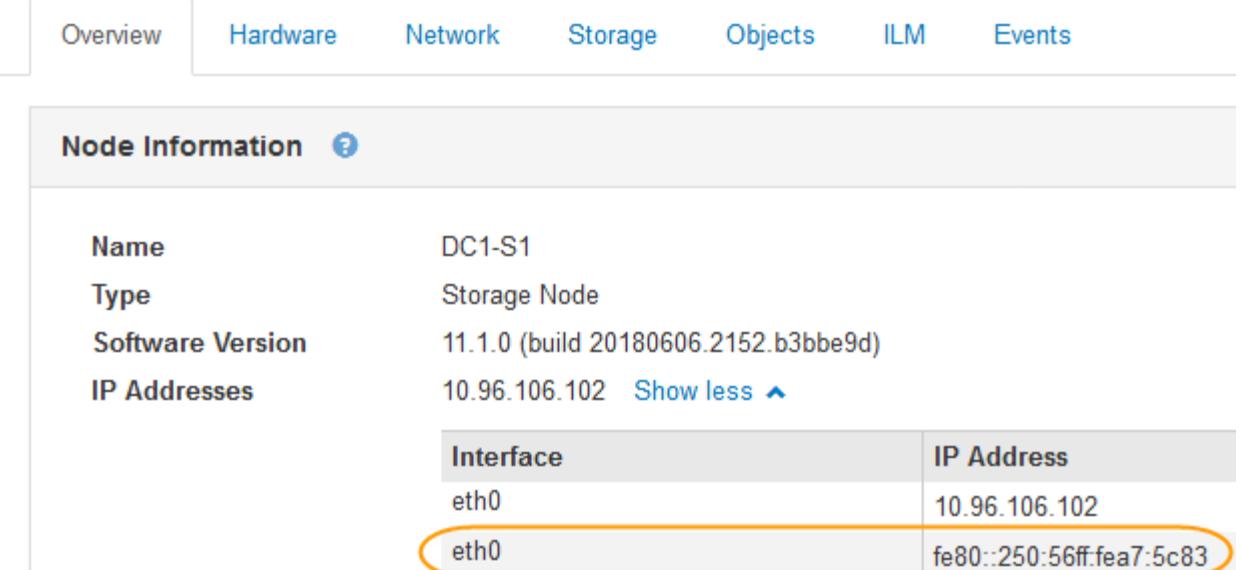
You might need to enable IPv6 support in the kernel if you have installed StorageGRID nodes on Linux hosts and you notice that IPv6 addresses have not been assigned to the node containers as expected.

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)



Interface	IP Address
eth0	10.96.106.102
eth0	fe80::250:56ff:fea7:5c83

- Select **Support > Tools > Grid Topology**. Then, select **node > SSM > Resources**. If an IPv6 address has been assigned, it is listed below the IPv4 address in the **Network Addresses** section.

If the IPv6 address is not shown and the node is installed on a Linux host, follow these steps to enable IPv6 support in the kernel.

Steps

1. Log in to the host as root or using an account with sudo permission.
2. Run the following command: `sysctl net.ipv6.conf.all.disable_ipv6`

```
root@SG:~ # sysctl net.ipv6.conf.all.disable_ipv6
```

The result should be 0.

```
net.ipv6.conf.all.disable_ipv6 = 0
```



If the result is not 0, see the documentation for your operating system for changing `sysctl` settings. Then, change the value to 0 before continuing.

3. Enter the StorageGRID node container: `storagegrid node enter node-name`
4. Run the following command: `sysctl net.ipv6.conf.all.disable_ipv6`

```
root@DC1-S1:~ # sysctl net.ipv6.conf.all.disable_ipv6
```

The result should be 1.

```
net.ipv6.conf.all.disable_ipv6 = 1
```



If the result is not 1, this procedure does not apply. Contact technical support.

5. Exit the container: `exit`

```
root@DC1-S1:~ # exit
```

6. As root, edit the following file: `/var/lib/storagegrid/settings/sysctl.d/net.conf`.

```
sudo vi /var/lib/storagegrid/settings/sysctl.d/net.conf
```

7. Locate the following two lines and remove the comment tags. Then, save and close the file.

```
net.ipv6.conf.all.disable_ipv6 = 0
```

```
net.ipv6.conf.default.disable_ipv6 = 0
```

8. Run these commands to restart the StorageGRID container:

```
storagegrid node stop node-name
```

```
storagegrid node start node-name
```

Review audit logs

Learn the StorageGRID system audit logs and see a list of all audit messages.

- [Audit message overview](#)
- [Audit log file and message formats](#)
- [Audit messages and the object lifecycle](#)
- [Audit messages](#)

Audit message overview

These instructions contain information about the structure and content of StorageGRID audit messages and audit logs. You can use this information to read and analyze the

audit trail of system activity.

These instructions are for administrators responsible for producing reports of system activity and usage that require analysis of the StorageGRID system's audit messages.

You are assumed to have a sound understanding of the nature of audited activities within the StorageGRID system. To use the text log file, you must have access to the configured audit share on the Admin Node.

Related information

[Administer StorageGRID](#)

Audit message flow and retention

All StorageGRID services generate audit messages during normal system operation. You should understand how these audit messages move through the StorageGRID system to the `audit.log` file.

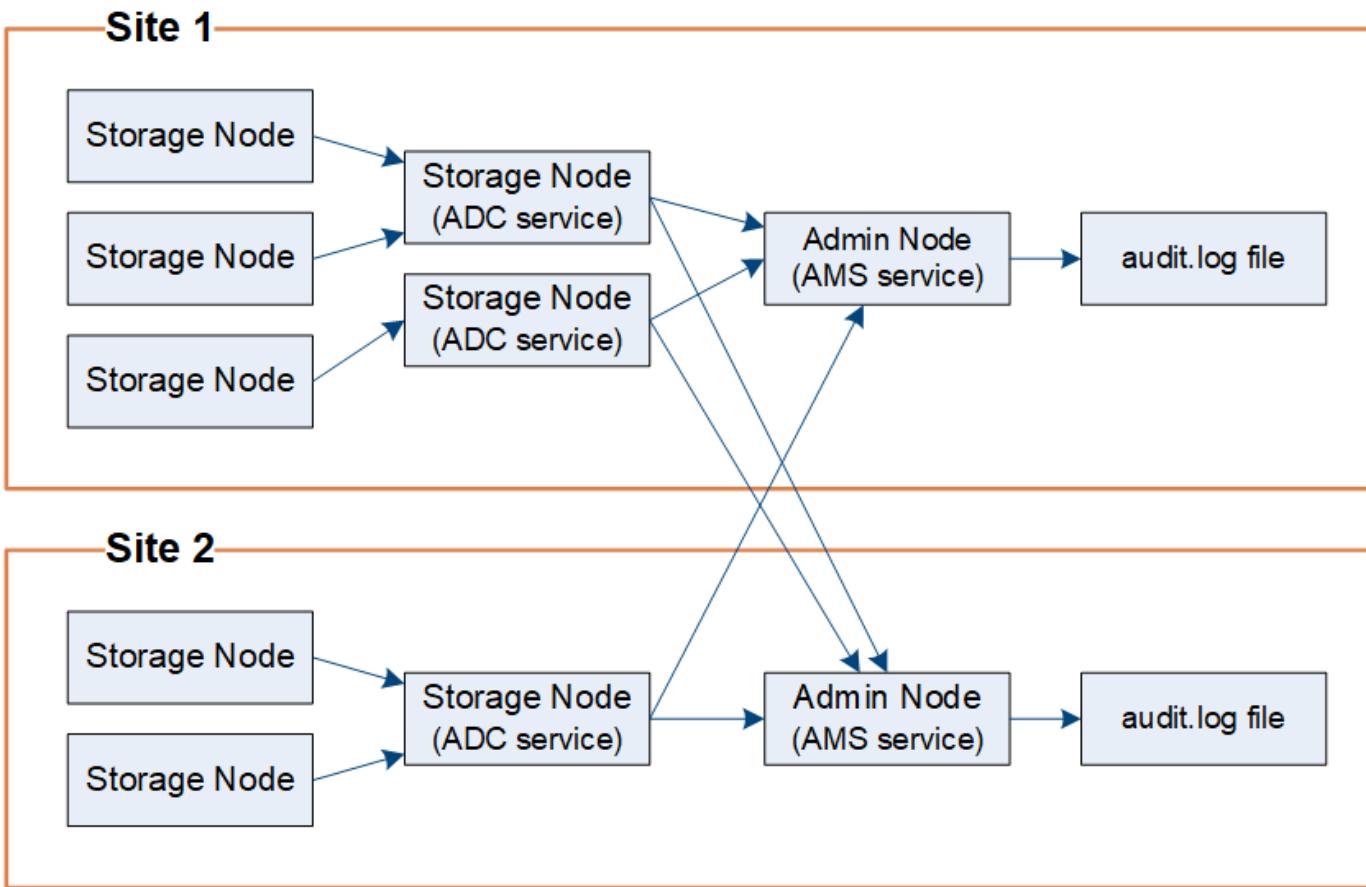
Audit message flow

Audit messages are processed by Admin Nodes and by those Storage Nodes that have an Administrative Domain Controller (ADC) service.

As shown in the audit message flow diagram, each StorageGRID node sends its audit messages to one of the ADC services at the data center site. The ADC service is automatically enabled for the first three Storage Nodes installed at each site.

In turn, each ADC service acts as a relay and sends its collection of audit messages to every Admin Node in the StorageGRID system, which gives each Admin Node a complete record of system activity.

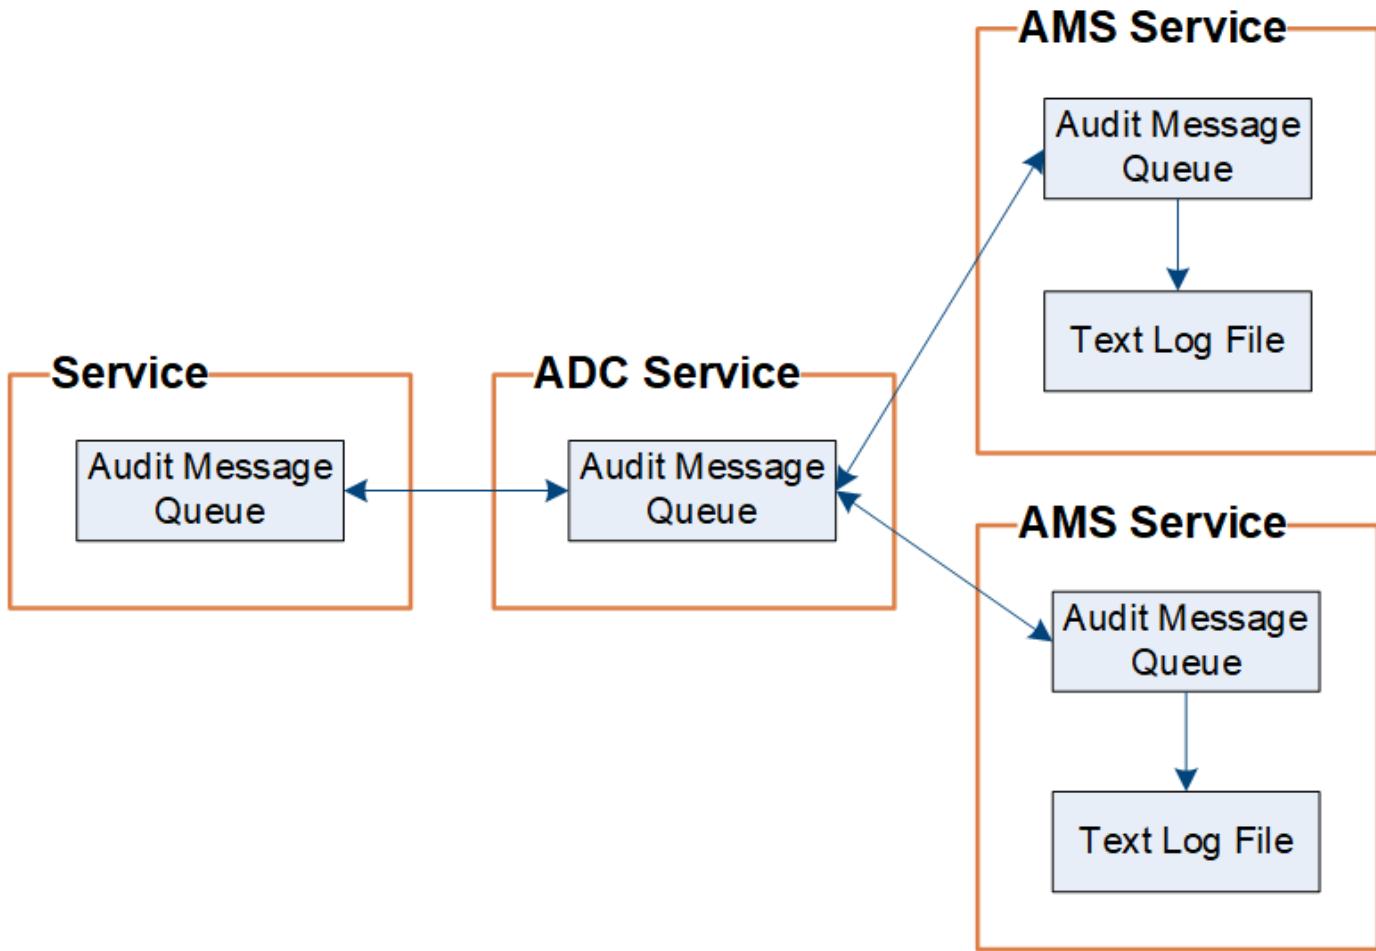
Each Admin Node stores audit messages in text log files; the active log file is named `audit.log`.



Audit message retention

StorageGRID uses a copy-and-delete process to ensure that no audit messages are lost before they can be written to the audit log.

When a node generates or relays an audit message, the message is stored in an audit message queue on the system disk of the grid node. A copy of the message is always held in an audit message queue until the message is written to the audit log file in the Admin Node's `/var/local/audit/export` directory. This helps prevent loss of an audit message during transport.



The audit message queue can temporarily increase due to network connectivity issues or insufficient audit capacity. As the queues increase, they consume more of the available space in each node's `/var/local/` directory. If the issue persists and a node's audit message directory becomes too full, the individual nodes will prioritize processing their backlog and become temporarily unavailable for new messages.

Specifically, you might see the following behaviors:

- If the `/var/local/audit/export` directory used by an Admin Node becomes full, the Admin Node will be flagged as unavailable to new audit messages until the directory is no longer full. S3 and Swift client requests are not affected. The XAMS (Unreachable Audit Repositories) alarm is triggered when an audit repository is unreachable.
- If the `/var/local/` directory used by a Storage Node with the ADC service becomes 92% full, the node will be flagged as unavailable to audit messages until the directory is only 87% full. S3 and Swift client requests to other nodes are not affected. The NRY (Available Audit Relays) alarm is triggered when audit relays are unreachable.



If there are no available Storage Nodes with the ADC service, the Storage Nodes store the audit messages locally.

- If the `/var/local/` directory used by a Storage Node becomes 85% full, the node will start refusing S3 and Swift client requests with 503 Service Unavailable.

The following types of issues can cause audit message queues to grow very large:

- The outage of an Admin Node or a Storage Node with the ADC service. If one of the system's nodes is down, the remaining nodes might become backlogged.
- A sustained activity rate that exceeds the audit capacity of the system.
- The `/var/local/` space on an ADC Storage Node becoming full for reasons unrelated to audit messages. When this happens, the node stops accepting new audit messages and prioritizes its current backlog, which can cause backlogs on other nodes.

Large audit queue alert and Audit Messages Queued (AMQS) alarm

To help you monitor the size of audit message queues over time, the **Large audit queue** alert and the legacy AMQS alarm are triggered when the number of messages in a Storage Node queue or Admin Node queue reaches certain thresholds.

If the **Large audit queue** alert or the legacy AMQS alarm is triggered, start by checking the load on the system—if there have been a significant number of recent transactions, the alert and the alarm should resolve over time and can be ignored.

If the alert or alarm persists and increases in severity, view a chart of the queue size. If the number is steadily increasing over hours or days, the audit load has likely exceeded the audit capacity of the system. Reduce the client operation rate or decrease the number of audit messages logged by changing the audit level for Client Writes and Client Reads to Error or Off. See "[Changing audit message levels](#)."

Duplicate messages

The StorageGRID system takes a conservative approach if a network or node failure occurs. For this reason, duplicate messages might exist in the audit log.

Changing audit message levels

You can adjust audit levels to increase or decrease the number of audit messages recorded in the audit log for each audit message category.

What you'll need

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

About this task

The audit messages recorded in the audit log are filtered based on the settings on the **Configuration > Monitoring > Audit** page.

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

- **System:** By default, this level is set to Normal.
- **Storage:** By default, this level is set to Error.
- **Management:** By default, this level is set to Normal.
- **Client Reads:** By default, this level is set to Normal.
- **Client Writes:** By default, this level is set to Normal.



These defaults apply if you initially installed StorageGRID using version 10.3 or later. If you have upgraded from an earlier version of StorageGRID, the default for all categories is set to Normal.



During upgrades, audit level configurations will not be effective immediately.

Steps

1. Select **Configuration > Monitoring > Audit**.

Audit

Audit Levels

System	Normal	▼
Storage	Error	▼
Management	Normal	▼
Client Reads	Normal	▼
Client Writes	Normal	▼

Audit Protocol Headers

Header Name 1	X-Forwarded-For	✖
Header Name 2	x-amz-*	✖

Save

2. For each category of audit message, select an audit level from the drop-down list:

Audit level	Description
Off	No audit messages from the category are logged.
Error	Only error messages are logged—audit messages for which the result code was not "successful" (SUCS).
Normal	Standard transactional messages are logged—the messages listed in these instructions for the category.
Debug	Deprecated. This level behaves the same as the Normal audit level.

The messages included for any particular level include those that would be logged at the higher levels. For example, the Normal level includes all of the Error messages.

3. Under **Audit Protocol Headers**, enter the name of the HTTP request headers to be included in Client Read and Client Write audit messages. Use an asterisk (*) as a wildcard, or use the escape sequence (*) as a literal asterisk. Click the plus sign to create a list of header name fields.



Audit protocol headers apply to S3 and Swift requests only.

When such HTTP headers are found in a request, they are included in the audit message under the field HTRH.



Audit protocol request headers are logged only if the audit level for **Client Reads** or **Client Writes** is not **Off**.

4. Click **Save**.

Related information

[System audit messages](#)

[Object storage audit messages](#)

[Management audit message](#)

[Client read audit messages](#)

[Administer StorageGRID](#)

Accessing the audit log file

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

What you'll need

- You must have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of an Admin Node.

Steps

1. Log in to an Admin Node:

- a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
- b. Enter the password listed in the `Passwords.txt` file.

2. Go to the directory containing the audit log files:

```
cd /var/local/audit/export
```

3. View the current or a saved audit log file, as required.

Related information

[Administer StorageGRID](#)

Audit log file rotation

Audit logs files are saved to an Admin Node's `/var/local/audit/export` directory. The active audit log files are named `audit.log`.

Once a day, the active `audit.log` file is saved, and a new `audit.log` file is started. The name of the saved file indicates when it was saved, in the format `yyyy-mm-dd.txt`. If more than one audit log is created in a single day, the file names use the date the file was saved, appended by a number, in the format `yyyy-mm-dd.txt.n`. For example, `2018-04-15.txt` and `2018-04-15.txt.1` are the first and second log files created and saved on 15 April 2018.

After a day, the saved file is compressed and renamed, in the format `yyyy-mm-dd.txt.gz`, which preserves the original date. Over time, this results in the consumption of storage allocated for audit logs on the Admin Node. A script monitors the audit log space consumption and deletes log files as necessary to free space in the `/var/local/audit/export` directory. Audit logs are deleted based on the date they were created, with the oldest being deleted first. You can monitor the script's actions in the following file:

`/var/local/log/manage-audit.log`.

This example shows the active `audit.log` file, the previous day's file (`2018-04-15.txt`), and the compressed file for the prior day (`2018-04-14.txt.gz`).

```
audit.log
2018-04-15.txt
2018-04-14.txt.gz
```

Audit log file and message formats

You can use audit logs to gather information about your system and troubleshoot issues. You should understand the format of the audit log file and the general format used for audit messages.

Audit log file format

The audit log files are found on every Admin Node and contain a collection of individual audit messages.

Each audit message contains the following:

- The Coordinated Universal Time (UTC) of the event that triggered the audit message (ATIM) in ISO 8601 format, followed by a space:

`YYYY-MM-DDTHH:MM:SS.UUUUUU`, where `UUUUUU` are microseconds.

- The audit message itself, enclosed within square brackets and beginning with `AUDT`.

The following example shows three audit messages in an audit log file (line breaks added for readability). These messages were generated when a tenant created an S3 bucket and added two objects to that bucket.

2019-08-07T18:43:30.247711

```
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991681][TIME(UI64):73520][SAI
P(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCKNB8M3MTWNT-
PhoTDwB9J0k7PtyLkQmA=="][SUSR(CSTR):"urn:sgws:identity::175300642415970547
18:root"]
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"buc
ket1"][AVER(UI32):10][ATIM(UI64):1565203410247711]
[ATYP(FC32):PUT][ANID(UI32):12454421][AMID(FC32):S3RQ][ATID(UI64):7074142
142472611085]]
```

2019-08-07T18:43:30.783597

```
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991696][TIME(UI64):120713][SA
IP(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCKNB8M3MTWNT-
PhoTDwB9J0k7PtyLkQmA=="][SUSR(CSTR):"urn:sgws:identity::175300642415970547
18:root"]
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"buc
ket1"][S3KY(CSTR):"fh-small-0"]
[CBID(UI64):0x779557A069B2C037][UUID(CSTR):"94BA6949-38E1-4B0C-BC80-
EB44FB4FCC7F"][CSIZ(UI64):1024][AVER(UI32):10]
[ATIM(UI64):1565203410783597][ATYP(FC32):PUT][ANID(UI32):12454421][AMID(F
C32):S3RQ][ATID(UI64):8439606722108456022]]
```

2019-08-07T18:43:30.784558

```
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991693][TIME(UI64):121666][SA
IP(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCKNB8M3MTWNT-
PhoTDwB9J0k7PtyLkQmA=="][SUSR(CSTR):"urn:sgws:identity::175300642415970547
18:root"]
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"buc
ket1"][S3KY(CSTR):"fh-small-2000"]
[CBID(UI64):0x180CBD8E678EED17][UUID(CSTR):"19CE06D0-D2CF-4B03-9C38-
E578D66F7ADD"][CSIZ(UI64):1024][AVER(UI32):10]
[ATIM(UI64):1565203410784558][ATYP(FC32):PUT][ANID(UI32):12454421][AMID(F
C32):S3RQ][ATID(UI64):13489590586043706682]]
```

In their default format, the audit messages in the audit log files are not easy to read or interpret. You can use the `audit-explain` tool to obtain simplified summaries of the audit messages in the audit log. You can use the `audit-sum` tool to summarize how many write, read, and delete operations were logged and how long these operations took.

Related information

[Using the audit-explain tool](#)

[Using the audit-sum tool](#)

Using the audit-explain tool

You can use the `audit-explain` tool to translate the audit messages in the audit log into an easy-to-read format.

What you'll need

- You must have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of the primary Admin Node.

About this task

The `audit-explain` tool, available on the primary Admin Node, provides simplified summaries of the audit messages in an audit log.



The `audit-explain` tool is primarily intended for use by technical support during troubleshooting operations. Processing `audit-explain` queries can consume a large amount of CPU power, which might impact StorageGRID operations.

This example shows typical output from the `audit-explain` tool. These four SPUT audit messages were generated when the S3 tenant with account ID 92484777680322627870 used S3 PUT requests to create a bucket named "bucket1" and add three objects to that bucket.

```
SPUT S3 PUT bucket bucket1 account:92484777680322627870 usec:124673
SPUT S3 PUT object bucket1/part1.txt tenant:92484777680322627870
cbid:9DCB157394F99FE5 usec:101485
SPUT S3 PUT object bucket1/part2.txt tenant:92484777680322627870
cbid:3CFBB07AB3D32CA9 usec:102804
SPUT S3 PUT object bucket1/part3.txt tenant:92484777680322627870
cbid:5373D73831ECC743 usec:93874
```

The `audit-explain` tool can process plain or compressed audit logs. For example:

```
audit-explain audit.log
```

```
audit-explain 2019-08-12.txt.gz
```

The `audit-explain` tool can also process multiple files at once. For example:

```
audit-explain audit.log 2019-08-12.txt.gz 2019-08-13.txt.gz
```

```
audit-explain /var/local/audit/export/*
```

Finally, the `audit-explain` tool can accept input from a pipe, which allows you to filter and preprocess the input using the `grep` command or other means. For example:

```
grep SPUT audit.log | audit-explain
```

```
grep bucket-name audit.log | audit-explain
```

Since audit logs can be very large and slow to parse, you can save time by filtering parts that you want to look at and running `audit-explain` on the parts, instead of the entire file.

 The `audit-explain` tool does not accept compressed files as piped input. To process compressed files, provide their file names as command-line arguments, or use the `zcat` tool to decompress the files first. For example:

```
zcat audit.log.gz | audit-explain
```

Use the `help` (`-h`) option to see the available options. For example:

```
$ audit-explain -h
```

Steps

1. Log in to the primary Admin Node:
 - a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
2. Enter the following command, where `/var/local/audit/export/audit.log` represents the name and the location of the file or files you want to analyze:

```
$ audit-explain /var/local/audit/export/audit.log
```

The `audit-explain` tool prints human-readable interpretations of all messages in the specified file or files.



To reduce line lengths and to aid readability, timestamps are not shown by default. If you want to see the timestamps, use the `timestamp` (`-t`) option.

Related information

[SPUT: S3 PUT](#)

Using the `audit-sum` tool

You can use the `audit-sum` tool to count the write, read, head, and delete audit messages and to see the minimum, maximum, and average time (or size) for each

operation type.

What you'll need

- You must have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of the primary Admin Node.

About this task

The `audit-sum` tool, available on the primary Admin Node, summarizes how many write, read, and delete operations were logged and how long these operations took.



The `audit-sum` tool is primarily intended for use by technical support during troubleshooting operations. Processing `audit-sum` queries can consume a large amount of CPU power, which might impact StorageGRID operations.

This example shows typical output from the `audit-sum` tool. This example shows how long protocol operations took.

message group average (sec)	count	min (sec)	max (sec)
=====	=====	=====	=====
=====			
IDEL	274		
SDEL	213371	0.004	20.934
0.352			
SGET	201906	0.010	1740.290
1.132			
SHEA	22716	0.005	2.349
0.272			
SPUT	1771398	0.011	1770.563
0.487			

The `audit-sum` tool provides counts and times for the following S3, Swift, and ILM audit messages in an audit log:

Code	Description	Refer to
ARCT	Archive Retrieve from Cloud-Tier	ARCT: Archive Retrieve from Cloud-Tier
ASCT	Archive Store Cloud-Tier	ASCT: Archive Store Cloud-Tier
IDEI	ILM Initiated Delete: Logs when ILM starts the process of deleting an object.	IDEI: ILM Initiated Delete
SDEL	S3 DELETE: Logs a successful transaction to delete an object or bucket.	SDEL: S3 DELETE

Code	Description	Refer to
SGET	S3 GET: Logs a successful transaction to retrieve an object or list the objects in a bucket.	SGET: S3 GET
SHEA	S3 HEAD: Logs a successful transaction to check for the existence of an object or bucket.	SHEA: S3 HEAD
SPUT	S3 PUT: Logs a successful transaction to create a new object or bucket.	SPUT: S3 PUT
WDEL	Swift DELETE: Logs a successful transaction to delete an object or container.	WDEL: Swift DELETE
WGET	Swift GET: Logs a successful transaction to retrieve an object or list the objects in a container.	WGET: Swift GET
WHEA	Swift HEAD: Logs a successful transaction to check for the existence of an object or container.	WHEA: Swift HEAD
WPUT	Swift PUT: Logs a successful transaction to create a new object or container.	WPUT: Swift PUT

The `audit-sum` tool can process plain or compressed audit logs. For example:

```
audit-sum audit.log
```

```
audit-sum 2019-08-12.txt.gz
```

The `audit-sum` tool can also process multiple files at once. For example:

```
audit-sum audit.log 2019-08-12.txt.gz 2019-08-13.txt.gz
```

```
audit-sum /var/local/audit/export/*
```

Finally, the `audit-sum` tool can also accept input from a pipe, which allows you to filter and preprocess the input using the `grep` command or other means. For example:

```
grep WGET audit.log | audit-sum
```

```
grep bucket1 audit.log | audit-sum
```

```
grep SPUT audit.log | grep bucket1 | audit-sum
```



This tool does not accept compressed files as piped input. To process compressed files, provide their file names as command-line arguments, or use the `zcat` tool to decompress the files first. For example:

```
audit-sum audit.log.gz
```

```
zcat audit.log.gz | audit-sum
```

You can use command-line options to summarize operations on buckets separately from operations on objects or to group message summaries by bucket name, by time period, or by target type. By default, the summaries show the minimum, maximum, and average operation time, but you can use the `size` (`-s`) option to look at object size instead.

Use the `help` (`-h`) option to see the available options. For example:

```
$ audit-sum -h
```

Steps

1. Log in to the primary Admin Node:
 - a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
2. If you want to analyze all messages related to write, read, head, and delete operations, follow these steps:
 - a. Enter the following command, where `/var/local/audit/export/audit.log` represents the name and the location of the file or files you want to analyze:

```
$ audit-sum /var/local/audit/export/audit.log
```

This example shows typical output from the `audit-sum` tool. This example shows how long protocol operations took.

message group average (sec)	count	min (sec)	max (sec)
=====	=====	=====	=====
=====			
IDEL	274		
SDEL	213371	0.004	20.934
0.352			
SGET	201906	0.010	1740.290
1.132			
SHEA	22716	0.005	2.349
0.272			
SPUT	1771398	0.011	1770.563
0.487			

In this example, SGET (S3 GET) operations are the slowest on average at 1.13 seconds, but SGET and SPUT (S3 PUT) operations both show long worst-case times of about 1,770 seconds.

- b. To show the slowest 10 retrieval operations, use the grep command to select only SGET messages and add the long output option (-l) to include object paths: `grep SGET audit.log | audit-sum -l`

The results include the type (object or bucket) and path, which allows you to grep the audit log for other messages relating to these particular objects.

```

Total:          201906 operations
Slowest:        1740.290 sec
Average:        1.132 sec
Fastest:        0.010 sec
Slowest operations:
  time(usec)      source ip      type      size(B)  path
  ======  ======  ======  ======  =====
  1740289662    10.96.101.125  object    5663711385
backup/r901OaQ8JB-1566861764-4519.iso
  1624414429    10.96.101.125  object    5375001556
backup/r901OaQ8JB-1566861764-6618.iso
  1533143793    10.96.101.125  object    5183661466
backup/r901OaQ8JB-1566861764-4518.iso
  70839        10.96.101.125  object    28338
bucket3/dat.1566861764-6619
  68487        10.96.101.125  object    27890
bucket3/dat.1566861764-6615
  67798        10.96.101.125  object    27671
bucket5/dat.1566861764-6617
  67027        10.96.101.125  object    27230
bucket5/dat.1566861764-4517
  60922        10.96.101.125  object    26118
bucket3/dat.1566861764-4520
  35588        10.96.101.125  object    11311
bucket3/dat.1566861764-6616
  23897        10.96.101.125  object    10692
bucket3/dat.1566861764-4516

```

From this example output, you can see that the three slowest S3 GET requests were for objects about 5 GB in size, which is much larger than the other objects. The large size accounts for the slow worst-case retrieval times.

3. If you want to determine what sizes of objects are being ingested into and retrieved from your grid, use the `size` option (-s):

```
audit-sum -s audit.log
```

message group average (MB)	count	min (MB)	max (MB)
=====	=====	=====	=====
IDEL 1654.502	274	0.004	5000.000
SDEL 1.695	213371	0.000	10.504
SGET 14.920	201906	0.000	5000.000
SHEA 2.967	22716	0.001	10.504
SPUT 2.495	1771398	0.000	5000.000

In this example, the average object size for SPUT is under 2.5 MB, but the average size for SGET is much larger. The number of SPUT messages is much higher than the number of SGET messages, indicating that most objects are never retrieved.

4. If you want to determine if retrievals were slow yesterday:

a. Issue the command on the appropriate audit log and use the group-by-time option (-gt), followed by the time period (for example, 15M, 1H, 10S):

```
grep SGET audit.log | audit-sum -gt 1H
```

message group average(sec)	count	min(sec)	max(sec)
=====	=====	=====	=====
=====			
2019-09-05T00 1.254	7591	0.010	1481.867
2019-09-05T01 1.115	4173	0.011	1740.290
2019-09-05T02 1.562	20142	0.011	1274.961
2019-09-05T03 1.254	57591	0.010	1383.867
2019-09-05T04 1.405	124171	0.013	1740.290
2019-09-05T05 1.562	420182	0.021	1274.511
2019-09-05T06 5.562	1220371	0.015	6274.961
2019-09-05T07 2.002	527142	0.011	1974.228
2019-09-05T08 1.105	384173	0.012	1740.290
2019-09-05T09 1.354	27591	0.010	1481.867

These results show that S3 GET traffic spiked between 06:00 and 07:00. The max and average times are both considerably higher at these times as well, and they did not ramp up gradually as the count increased. This suggests that capacity was exceeded somewhere, perhaps in the network or in the grid's ability to process requests.

b. To determine what size objects were being retrieved each hour yesterday, add the size option (`-s`) to the command:

```
grep SGET audit.log | audit-sum -gt 1H -s
```

message group average (B)	count	min (B)	max (B)
=====	=====	=====	=====
2019-09-05T00 1.976	7591	0.040	1481.867
2019-09-05T01 2.062	4173	0.043	1740.290
2019-09-05T02 2.303	20142	0.083	1274.961
2019-09-05T03 1.182	57591	0.912	1383.867
2019-09-05T04 1.528	124171	0.730	1740.290
2019-09-05T05 2.398	420182	0.875	4274.511
2019-09-05T06 51.328	1220371	0.691	5663711385.961
2019-09-05T07 2.147	527142	0.130	1974.228
2019-09-05T08 1.878	384173	0.625	1740.290
2019-09-05T09 1.354	27591	0.689	1481.867

These results indicate that some very large retrievals occurred when the overall retrieval traffic was at its maximum.

- To see more detail, use the audit-explain tool to review all the SGET operations during that hour:

```
grep 2019-09-05T06 audit.log | grep SGET | audit-explain | less
```

If the output of the grep command is expected to be many lines, add the less command to show the contents of the audit log file one page (one screen) at a time.

- If you want to determine if SPUT operations on buckets are slower than SPUT operations for objects:

- Start by using the -go option, which groups messages for object and bucket operations separately:

```
grep SPUT sample.log | audit-sum -go
```

message group	count	min(sec)	max(sec)
average(sec)			
=====	=====	=====	=====
=====			
SPUT.bucket	1	0.125	0.125
0.125			
SPUT.object	12	0.025	1.019
0.236			

The results show that SPUT operations for buckets have different performance characteristics than SPUT operations for objects.

b. To determine which buckets have the slowest SPUT operations, use the `-gb` option, which groups messages by bucket:

```
grep SPUT audit.log | audit-sum -gb
```

message group	count	min(sec)	max(sec)
average(sec)			
=====	=====	=====	=====
=====			
SPUT.cho-non-versioning	71943	0.046	1770.563
1.571			
SPUT.cho-versioning	54277	0.047	1736.633
1.415			
SPUT.cho-west-region	80615	0.040	55.557
1.329			
SPUT.1dt002	1564563	0.011	51.569
0.361			

c. To determine which buckets have the largest SPUT object size, use both the `-gb` and the `-s` options:

```
grep SPUT audit.log | audit-sum -gb -s
```

message group average (B)	count	min (B)	max (B)
=====	=====	=====	=====
=====			
SPUT.cho-non-versioning 21.672	71943	2.097	5000.000
SPUT.cho-versioning 21.120	54277	2.097	5000.000
SPUT.cho-west-region 14.433	80615	2.097	800.000
SPUT.1dt002 0.352	1564563	0.000	999.972

Related information

[Using the audit-explain tool](#)

Audit message format

Audit messages exchanged within the StorageGRID system include standard information common to all messages and specific content describing the event or activity being reported.

If the summary information provided by the `audit-explain` and `audit-sum` tools is insufficient, refer to this section to understand the general format of all audit messages.

The following is an example audit message as it might appear in the audit log file:

```
2014-07-17T03:50:47.484627
[AUDT:[RSLT(FC32):VRGN][AVER(UI32):10][ATIM(UI64):1405569047484627][ATYP(F
C32):SYSU][ANID(UI32):11627225][AMID(FC32):ARNI][ATID(UI64):94457363265006
03516]]
```

Each audit message contains a string of attribute elements. The entire string is enclosed in brackets ([]), and each attribute element in the string has the following characteristics:

- Enclosed in brackets []
- Introduced by the string AUDT, which indicates an audit message
- Without delimiters (no commas or spaces) before or after
- Terminated by a line feed character \n

Each element includes an attribute code, a data type, and a value that are reported in this format:

```
[ATTR(type):value] [ATTR(type):value]...  
[ATTR(type):value] \n
```

The number of attribute elements in the message depends on the event type of the message. The attribute elements are not listed in any particular order.

The following list describes the attribute elements:

- ATTR is a four-character code for the attribute being reported. There are some attributes that are common to all audit messages and others that are event-specific.
- type is a four-character identifier of the programming data type of the value, such as UI64, FC32, and so on. The type is enclosed in parentheses ().
- value is the content of the attribute, typically a numeric or text value. Values always follow a colon (:). Values of data type CSTR are surrounded by double quotes " ".

Related information

[Using the audit-explain tool](#)

[Using the audit-sum tool](#)

[Audit messages](#)

[Common elements in audit messages](#)

[Data types](#)

[Audit message examples](#)

Data types

Different data types are used to store information in audit messages.

Type	Description
UI32	Unsigned long integer (32 bits); it can store the numbers 0 to 4,294,967,295.
UI64	Unsigned double long integer (64 bits); it can store the numbers 0 to 18,446,744,073,709,551,615.
FC32	Four-character constant; a 32-bit unsigned integer value represented as four ASCII characters such as "ABCD."
IPAD	Used for IP addresses.

Type	Description
CSTR	<p>A variable-length array of UTF-8 characters. Characters can be escaped with the following conventions:</p> <ul style="list-style-type: none"> • Backslash is \\. • Carriage return is \\r. • Double quotes is \\". • Line feed (new line) is \\n. • Characters can be replaced by their hexadecimal equivalents (in the format \\xHH, where HH is the hexadecimal value representing the character).

Event-specific data

Each audit message in the audit log records data specific to a system event.

Following the opening **[AUDT:** container that identifies the message itself, the next set of attributes provide information about the event or action described by the audit message. These attributes are highlighted in the following example:

```
2018-12-05T08:24:45.921845 [AUDT:[RSLT(FC32):SUCS]
[TIME(UI64):11454] [SAIP(IPAD):"10.224.0.100"]
[S3AI(CSTR):"60025621595611246499"] [SACC(CSTR):"account"]
[S3AK(CSTR):"SGKH4_Nc8SO1H6w3w0nCOFCGgk_E6dYzKlumRsKJA=="]
[SUSR(CSTR):"urn:sgws:identity::60025621595611246499:root"]
[SBAI(CSTR):"60025621595611246499"] [SBAC(CSTR):"account"] [S3BK(CSTR):"bucket"]
[S3KY(CSTR):"object"] [CBID(UI64):0xCC128B9B9E428347]
[UUID(CSTR):"B975D2CE-E4DA-4D14-8A23-1CB4B83F2CD8"] [CSIZ(UI64):30720]
[AVER(UI32):10] [ATIM(UI64):1543998285921845] [ATYP(FC32):SHEA]
[ANID(UI32):12281045] [AMID(FC32):S3RQ] [ATID(UI64):15552417629170647261]
```

The **ATYP** element (underlined in the example) identifies which event generated the message. This example message includes the **SHEA** message code ([ATYP(FC32):SHEA]), indicating it was generated by a successful S3 HEAD request.

Related information

[Common elements in audit messages](#)

[Audit messages](#)

[Common elements in audit messages](#)

All audit messages contain the common elements.

Code	Type	Description
AMID	FC32	Module ID: A four-character identifier of the module ID that generated the message. This indicates the code segment within which the audit message was generated.
ANID	UI32	Node ID: The grid node ID assigned to the service that generated the message. Each service is allocated a unique identifier at the time the StorageGRID system is configured and installed. This ID cannot be changed.
ASES	UI64	Audit Session Identifier: In previous releases, this element indicated the time at which the audit system was initialized after the service started up. This time value was measured in microseconds since the operating system epoch (00:00:00 UTC on 1 January, 1970). Note: This element is obsolete and no longer appears in audit messages.
ASQN	UI64	Sequence Count: In previous releases, this counter was incremented for each generated audit message on the grid node (ANID) and reset to zero at service restart. Note: This element is obsolete and no longer appears in audit messages.
ATID	UI64	Trace ID: An identifier that is shared by the set of messages that were triggered by a single event.
ATIM	UI64	Timestamp: The time the event was generated that triggered the audit message, measured in microseconds since the operating system epoch (00:00:00 UTC on 1 January, 1970). Note that most available tools for converting the timestamp to local date and time are based on milliseconds. Rounding or truncation of the logged timestamp might be required. The human-readable time that appears at the beginning of the audit message in the audit.log file is the ATIM attribute in ISO 8601 format. The date and time are represented as <i>YYYY-MMDDTHH:MM:SS.uuuuuu</i> , where the <i>T</i> is a literal string character indicating the beginning of the time segment of the date. <i>uuuuuu</i> are microseconds.
ATYP	FC32	Event Type: A four-character identifier of the event being logged. This governs the "payload" content of the message: the attributes that are included.
AVER	UI32	Version: The version of the audit message. As the StorageGRID software evolves, new versions of services might incorporate new features in audit reporting. This field enables backward compatibility in the AMS service to process messages from older versions of services.
RSLT	FC32	Result: The result of event, process, or transaction. If is not relevant for a message, NONE is used rather than SUCS so that the message is not accidentally filtered.

Audit message examples

You can find detailed information in each audit message. All audit messages use the same format.

The following is a sample audit message as it might appear in the `audit.log` file:

```
2014-07-17T21:17:58.959669
[AUDT: [RSLT(FC32):SUCS] [TIME(UI64):246979] [S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"] [
S3AK(CSTR):"UJXDKQOXB7YARDS71Q2"] [S3BK(CSTR):"s3small1"] [S3K
Y(CSTR):"hello1"] [CBID(UI64):0x50C4F7AC2BC8EDF7] [CSIZ(UI64):0
] [AVER(UI32):10] [ATIM(UI64):1405631878959669] [ATYP(FC32):PUT
] [ANID(UI32):12872812] [AMID(FC32):S3RQ] [ATID(UI64):1579224144
102530435]
```

The audit message contains information about the event being recorded, as well as information about the audit message itself.

To identify which event is recorded by the audit message, look for the `ATYP` attribute (highlighted below):

```
2014-07-17T21:17:58.959669
[AUDT: [RSLT(FC32):SUCS] [TIME(UI64):246979] [S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"] [
S3AK(CSTR):"UJXDKQOXB7YARDS71Q2"] [S3BK(CSTR):"s3small1"] [S3K
Y(CSTR):"hello1"] [CBID(UI64):0x50C4F7AC2BC8EDF7] [CSIZ(UI64):0
] [AVER(UI32):10] [ATIM(UI64):1405631878959669] [ATYP(FC32):SP
UT] [ANID(UI32):12872812] [AMID(FC32):S3RQ] [ATID(UI64):1579224
144102530435]
```

The value of the `ATYP` attribute is `PUT`. `PUT` represents an S3 PUT transaction, which logs the ingest of an object to a bucket.

The following audit message also shows the bucket to which the object is associated:

```
2014-07-17T21:17:58.959669
[AUDT: [RSLT(FC32):SUCS] [TIME(UI64):246979] [S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"] [
S3AK(CSTR):"UJXDKQOXB7YARDS71Q2"] [S3BK(CSTR):"s3small1"] [S3
KY(CSTR):"hello1"] [CBID(UI64):0x50C4F7AC2BC8EDF7] [CSIZ(UI64):0
] [AVER(UI32):10] [ATIM(UI64):1405631878959669] [ATYP(FC32):PUT
] [ANID(UI32):12872812] [AMID(FC32):S3RQ] [ATID(UI64):157922414
4102530435]
```

To discover when the `PUT` event occurred, note the Universal Coordinated Time (UTC) timestamp at the

beginning of the audit message. This value is a human-readable version of the ATIM attribute of the audit message itself:

2014-07-17T21:17:58.959669

```
[AUDT: [RSLT (FC32) : SUCS] [TIME (UI64) : 246979] [S3AI (CSTR) : "bc644d381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"] [S3AK (CSTR) : "UJXDKKQOXB7YARDS71Q2"] [S3BK (CSTR) : "s3small1"] [S3KY (CSTR) : "hello1"] [CBID (UI64) : 0x50C4F7AC2BC8EDF7] [CSIZ (UI64) : 0] [AVER (UI32) : 10] [ATIM (UI64) : 1405631878959669] [ATYP (FC32) : SPUT] [ANID (UI32) : 12872812] [AMID (FC32) : S3RQ] [ATID (UI64) : 1579224144102530435]]
```

ATIM records the time, in microseconds, since the beginning of the UNIX epoch. In the example, the value 1405631878959669 translates to Thursday, 17-Jul-2014 21:17:59 UTC.

Related information

[SPUT: S3 PUT](#)

[Common elements in audit messages](#)

Audit messages and the object lifecycle

Audit messages are generated each time an object is ingested, retrieved, or deleted. You can identify these transactions in the audit log by locating API-specific (S3 or Swift) audit messages.

Audit messages are linked through identifiers specific to each protocol.

Protocol	Code
Linking S3 operations	S3BK (S3 Bucket) and/or S3KY (S3 Key)
Linking Swift operations	WCON (Swift Container) and/or WOBJ (Swift Object)
Linking internal operations	CBID (Object's Internal Identifier)

Timing of audit messages

Because of factors such as timing differences between grid nodes, object size, and network delays, the order of audit messages generated by the different services can vary from that shown in the examples in this section.

Information lifecycle management policy configuration

With the default ILM policy (Baseline 2 Copy), object data is copied once for a total of two copies. If the ILM policy requires more than two copies, there will be an additional set of CBRE, CBSE, and SCMT messages for each extra copy. For more information about ILM policies, see information about managing objects with information lifecycle management.

Archive Nodes

The series of audit messages generated when an Archive Node sends object data to an external archival storage system is similar to that for Storage Nodes except that there is no SCMT (Store Object Commit) message, and the ATCE (Archive Object Store Begin) and ASCE (Archive Object Store End) messages are generated for each archived copy of object data.

The series of audit messages generated when an Archive Node retrieves object data from an external archival storage system is similar to that for Storage Nodes except that the ARCB (Archive Object Retrieve Begin) and ARCE (Archive Object Retrieve End) messages are generated for each retrieved copy of object data.

The series of audit messages generated when an Archive Node deletes object data from an external archival storage system is similar to that for Storage Nodes except that there is no SREM (Object Store Remove) message, and there is an AREM (Archive Object Remove) message for each delete request.

Related information

[Manage objects with ILM](#)

Object ingest transactions

You can identify client ingest transactions in the audit log by locating API-specific (S3 or Swift) audit messages.

Not all audit messages generated during an ingest transaction are listed in the following tables. Only the messages required to trace the ingest transaction are included.

S3 ingest audit messages

Code	Name	Description	Trace	See
SPUT	S3 PUT transaction	An S3 PUT ingest transaction has completed successfully.	CBID, S3BK, S3KY	SPUT: S3 PUT
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	ORLM: Object Rules Met

Swift ingest audit messages

Code	Name	Description	Trace	See
WPUT	Swift PUT transaction	A Swift PUT ingest transaction has successfully completed.	CBID, WCON, WOBJ	WPUT: Swift PUT
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	ORLM: Object Rules Met

Example: S3 object ingest

The series of audit messages below is an example of the audit messages generated and saved to the audit log when an S3 client ingests an object to a Storage Node (LDR service).

In this example, the active ILM policy includes the stock ILM rule, Make 2 Copies.



Not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 ingest transaction (SPUT) are listed.

This example assumes that an S3 bucket has been previously created.

SPUT: S3 PUT

The SPUT message is generated to indicate that an S3 PUT transaction has been issued to create an object in a specific bucket.

```
2017-07-
17T21:17:58.959669 [AUDT: [RSLT (FC32) :SUCS] [TIME (UI64) :25771] [SAIP (IPAD) :"10
.96.112.29"] [S3AI (CSTR) :"70899244468554783528"] [SACC (CSTR) :"test"] [S3AK (CS
TR) :"SGKHyalRU_5cLf1qajtaFmxJn9461AWRJfBF33gAOg=="] [SUSR (CSTR) :"urn:sgws:i
dentity::70899244468554783528:root"] [SBAI (CSTR) :"70899244468554783528"] [SB
AC (CSTR) :"test"] [S3BK (CSTR) :"example"]<strong
class="S3KY (CSTR) :"testobject-0-
3"">[CBID (UI64) :0x8EF52DF8025E63A8]</strong>[CSIZ (UI64) :30720] [AVER (UI32) :
10]<strong
class="ATIM (UI64) :150032627859669">[ATYP (FC32) :PUT]</strong>[ANID (UI32) :1
2086324] [AMID (FC32) :S3RQ] [ATID (UI64) :14399932238768197038] ]
```

ORLM: Object Rules Met

The ORLM message indicates that the ILM policy has been satisfied for this object. The message includes the object's CBID and the name of the ILM rule that was applied.

For replicated objects, the LOCS field includes the LDR node ID and volume ID of the object locations.

```
2019-07-17T21:18:31.230669 [AUDT:
<strong>[CBID (UI64) :0x50C4F7AC2BC8EDF7]</strong> [RULE (CSTR) :"Make 2
Copies"] [STAT (FC32) :DONE] [CSIZ (UI64) :0] [UUID (CSTR) :"0B344E18-98ED-4F22-
A6C8-A93ED68F8D3F"]<strong class="LOCS (CSTR) : *"CLDI 12828634
2148730112">[RSLT (FC32) :SUCS] [AVER (UI32) :10] [ATYP (FC32) :ORLM]</strong>
[ATIM (UI64) :1563398230669] [ATID (UI64) :15494889725796157557] [ANID (UI32) :131
00453] [AMID (FC32) :BCMS] ]
```

For erasure-coded objects, the LOCS field includes the Erasure Coding profile ID and the Erasure Coding group ID

```
2019-02-23T01:52:54.647537
```

```
[AUDT:[CBID(UI64):0xFA8ABE5B5001F7E2][RULE(CSTR):"EC_2_plus_1"][STAT(FC32):DONE][CSIZ(UI64):10000][UUID(CSTR):"E291E456-D11A-4701-8F51-D2F7CC9AFECA"][LOCS(CSTR):"CLEC 1 A471E45D-A400-47C7-86AC-12E77F229831"][RSLT(FC32):SUCS][AVER(UI32):10][ATYP(FC32):ORLM][ANID(UI32):12355278][AMID(FC32):ILMX][ATID(UI64):4168559046473725560]]
```

The PATH field includes S3 bucket and key information or Swift container and object information, depending on which API was used.

```
2019-09-15.txt:2018-01-24T13:52:54.131559
```

```
[AUDT:[CBID(UI64):0x82704DFA4C9674F4][RULE(CSTR):"Make 2 Copies"][STAT(FC32):DONE][CSIZ(UI64):3145729][UUID(CSTR):"8C1C9CAC-22BB-4880-9115-CE604F8CE687"][PATH(CSTR):"frisbee_Bucket1/GridDataTests151683676324774_1_1vf9d"][LOCS(CSTR):"CLDI 12525468, CLDI 12222978"][RSLT(FC32):SUCS][AVER(UI32):10][ATIM(UI64):1568555574559][ATYP(FC32):ORLM][ANID(UI32):12525468][AMID(FC32):OBDI][ATID(UI64):344833886538369336]]
```

Object delete transactions

You can identify object delete transactions in the audit log by locating API-specific (S3 and Swift) audit messages.

Not all audit messages generated during a delete transaction are listed in the following tables. Only messages required to trace the delete transaction are included.

S3 delete audit messages

Code	Name	Description	Trace	See
SDEL	S3 Delete	Request made to delete the object from a bucket.	CBID, S3KY	SDEL: S3 DELETE

Swift delete audit messages

Code	Name	Description	Trace	See
WDEL	Swift Delete	Request made to delete the object from a container, or the container.	CBID, WOBJ	WDEL: Swift DELETE

Example: S3 object deletion

When an S3 client deletes an object from a Storage Node (LDR service), an audit message is generated and saved to the audit log.



Not all audit messages generated during a delete transaction are listed in the example below. Only those related to the S3 delete transaction (SDEL) are listed.

SDEL: S3 Delete

Object deletion begins when the client sends a DELETE Object request to an LDR service. The message contains the bucket from which to delete the object and the object's S3 Key, which is used to identify the object.

```
2017-07-
17T21:17:58.959669 [AUDT:[RSLT(FC32):SUCS] [TIME(UI64):14316] [SAIP(IPAD):"10
.96.112.29"] [S3AI(CSTR):"70899244468554783528"] [SACC(CSTR):"test"] [S3AK(CS
TR):"SGKHyalRU_5cLflqajtaFmxJn9461AWRJfBF33gAOg=="] [SUSR(CSTR):"urn:sgws:i
dentity::70899244468554783528:root"] [SBAI(CSTR):"70899244468554783528"] [SB
AC(CSTR):"test"] <strong>[S3BK(CSTR):"example"] [S3KY(CSTR):"testobject-0-
7"] [CBID(UI64):0x339F21C5A6964D89]</strong>
[CSIZ(UI64):30720] [AVER(UI32):10] [ATIM(UI64):150032627859669]
<strong>[ATYP(FC32):SDEL]</strong> [ANID(UI32):12086324] [AMID(FC32):S3RQ] [A
TID(UI64):4727861330952970593]]
```

Object retrieve transactions

You can identify object retrieve transactions in the audit log by locating API-specific (S3 and Swift) audit messages.

Not all audit messages generated during a retrieve transaction are listed in the following tables. Only messages required to trace the retrieve transaction are included.

S3 retrieval audit messages

Code	Name	Description	Trace	See
SGET	S3 GET	Request made to retrieve an object from a bucket.	CBID, S3BK, S3KY	SGET: S3 GET

Swift retrieval audit messages

Code	Name	Description	Trace	See
WGET	Swift GET	Request made to retrieve an object from a container.	CBID, WCON, WOBJ	WGET: Swift GET

Example: S3 object retrieval

When an S3 client retrieves an object from a Storage Node (LDR service), an audit message is generated and saved to the audit log.

Note that not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 retrieval transaction (SGET) are listed.

SGET: S3 GET

Object retrieval begins when the client sends a GET Object request to an LDR service. The message contains the bucket from which to retrieve the object and the object's S3 Key, which is used to identify the object.

```
2017-09-20T22:53:08.782605
[AUDT:[RSLT(FC32):SUFS][TIME(UI64):47807][SAIP(IPAD):"10.96.112.26"][S3AI(
CSTR):"43979298178977966408"][SACC(CSTR):"s3-account-
a"][S3AK(CSTR):"SGKht7GzEcu0yXhFhT_rL5mep4nJt1w75GBh-
O_Few=="][SUSR(CSTR):"urn:sgws:identity::43979298178977966408:root"][SBAI(
CSTR):"43979298178977966408"][SBAC(CSTR):"s3-account-a"]
[S3BK(CSTR):"bucket-
anonymous"][S3KY(CSTR):"Hello.txt"][CBID(UI64):0x83D70C6F1F662B02][CSIZ(UI
64):12][AVER(UI32):10][ATIM(UI64):1505947988782605][ATYP(FC32):SGET][ANID(
UI32):12272050][AMID(FC32):S3RQ][ATID(UI64):17742374343649889669]]
```

If the bucket policy allows, a client can anonymously retrieve objects, or can retrieve objects from a bucket that is owned by a different tenant account. The audit message contains information about the bucket owner's tenant account so that you can track these anonymous and cross-account requests.

In the following example message, the client sends a GET Object request for an object stored in a bucket that they do not own. The values for SBAI and SBAC record the bucket owner's tenant account ID and name, which differs from the tenant account ID and name of the client recorded in S3AI and SACC.

```
2017-09-20T22:53:15.876415
[AUDT:[RSLT(FC32):SUFS][TIME(UI64):53244][SAIP(IPAD):"10.96.112.26"]
<strong>[S3AI(CSTR):"17915054115450519830"][SACC(CSTR):"s3-account-
b"]</strong>[S3AK(CSTR):"SGKHpoblWlP_kBkqSCbTi754Ls81BUog67I2LlSiUg=="]<st
rong
class="SUSR(CSTR):"urn:sgws:identity::17915054115450519830:root"">[SBAI(CS
TR):"43979298178977966408"][SBAC(CSTR):"s3-account-
a"]</strong>[S3BK(CSTR):"bucket-
anonymous"][S3KY(CSTR):"Hello.txt"][CBID(UI64):0x83D70C6F1F662B02][CSIZ(UI
64):12][AVER(UI32):10][ATIM(UI64):1505947995876415][ATYP(FC32):SGET][ANID(
UI32):12272050][AMID(FC32):S3RQ][ATID(UI64):6888780247515624902]]
```

Metadata update messages

Audit messages are generated when an S3 client updates an object's metadata.

S3 metadata update audit messages

Code	Name	Description	Trace	See
SUPD	S3 Metadata Updated	Generated when an S3 client updates the metadata for an ingested object.	CBID, S3KY, HTRH	SUPD: S3 Metadata Updated

Example: S3 metadata update

The example shows a successful transaction to update the metadata for an existing S3 object.

SUPD: S3 Metadata Update

The S3 client makes a request (SUPD) to update the specified metadata (`x-amz-meta-*`) for the S3 object (S3KY). In this example, request headers are included in the field HTRH because it has been configured as an audit protocol header (**Configuration > Monitoring > Audit**).

```
2017-07-11T21:54:03.157462
[AUDT:[RSLT(FC32):SUFS][TIME(UI64):17631][SAIP(IPAD):"10.96.100.254"]
[HTRH(CSTR):>{"accept-encoding": "identity", "authorization": "AWS
LIUF17FGJARQHPY2E761:jul/hnzs/uNY+aVvV01TSYhEGts=", "content-length": "0", "date": "Tue, 11 Jul 2017 21:54:03
GMT", "host": "10.96.99.163:18082", "user-agent": "aws-cli/1.9.20 Python/2.7.6 Linux/3.13.0-119-generic
botocore/1.3.20", "x-amz-copy-source": "/testbkt1/testobj1", "x-amz-metadata-
directive": "REPLACE", "x-amz-meta-city": "Vancouver"}]
[S3AI(CSTR):"20956855414285633225"] [SACC(CSTR):"acct1"] [S3AK(CSTR):"SGKHyy
v9ZQqWRbJSQc5vI7mgioJwrdplShE02AUaww=="]
[SUSR(CSTR):"urn:sgws:identity::20956855414285633225:root"]
[SBAI(CSTR):"20956855414285633225"] [SBAC(CSTR):"acct1"] [S3BK(CSTR):"testbk
t1"]
[S3KY(CSTR):"testobj1"] [CBID(UI64):0xCB1D5C213434DD48] [CSIZ(UI64):10] [AVER
(UI32):10]
[ATIM(UI64):1499810043157462] [ATYP(FC32):SUPD] [ANID(UI32):12258396] [AMID(F
C32):S3RQ]
[ATID(UI64):8987436599021955788]]
```

Related information

[Changing audit message levels](#)

Audit messages

Detailed descriptions of audit messages returned by the system are listed in the following

sections. Each audit message is first listed in a table that groups related messages by the class of activity that the message represents. These groupings are useful both for understanding the types of activities that are audited, and for selecting the desired type of audit message filtering.

The audit messages are also listed alphabetically by their four-character codes. This alphabetic listing enables you to find information about specific messages.

The four-character codes used throughout this chapter are the ATYP values found in the audit messages as shown in the following sample message:

```
2014-07-17T03:50:47.484627
\[AUDT:[RSLT(FC32):VRGN][AVER(UI32):10][ATIM(UI64):1405569047484627][<strong>ATYP\ (FC32\):SYSU</strong>][ANID(UI32):11627225][AMID(FC32):ARNI][ATID(UI64):9445736326500603516]]
```

Related information

[Audit messages](#)

[Changing audit message levels](#)

Audit message categories

You should be familiar with the various categories within which audit messages are grouped. These groups are organized based on the class of activity that the message represents.

System audit messages

You should be familiar with audit messages belonging to the system audit category. These are events related to the auditing system itself, grid node states, system-wide task activity (grid tasks), and service backup operations, so that you can address potential issues.

Code	Message title and description	See
ECOC	Corrupt Erasure Coded Data Fragment: Indicates that a corrupt erasure coded data fragment has been detected.	ECOC: Corrupt Erasure Coded Data Fragment
ETAF	Security Authentication Failed: A connection attempt using Transport Layer Security (TLS) failed.	ETAF: Security Authentication Failed

Code	Message title and description	See
GNRG	GNDS Registration: A service updated or registered information about itself in the StorageGRID system.	GNRG: GNDS Registration
GNUR	GNDS Unregistration: A service has unregistered itself from the StorageGRID system.	GNUR: GNDS Unregistration
GTED	Grid Task Ended: The CMN service finished processing the grid task.	GTED: Grid Task Ended
GTST	Grid Task Started: The CMN service started to process the grid task.	GTST: Grid Task Started
GTSU	Grid Task Submitted: A grid task was submitted to the CMN service.	GTSU: Grid Task Submitted
IDEL	ILM Initiated Delete: This audit message is generated when ILM starts the process of deleting an object.	IDEL: ILM Initiated Delete
LKCU	Overwritten Object Cleanup. This audit message is generated when an overwritten object is automatically removed to free up storage space.	LKCU: Overwritten Object Cleanup
LLST	Location Lost: This audit message is generated when a location is lost.	LLST: Location Lost
OLST	Object Lost: A requested object cannot be located within the StorageGRID system.	OLST: System Detected Lost Object
ORLM	Object Rules Met: Object data is stored as specified by the ILM rules.	ORLM: Object Rules Met
SADD	Security Audit Disable: Audit message logging was turned off.	SADD: Security Audit Disable

Code	Message title and description	See
SADE	Security Audit Enable: Audit message logging has been restored.	SADE: Security Audit Enable
SVRF	Object Store Verify Fail: A content block failed verification checks.	SVRF: Object Store Verify Fail
SVRU	Object Store Verify Unknown: Unexpected object data detected in the object store.	SVRU: Object Store Verify Unknown
SYSD	Node Stop: A shutdown was requested.	SYSD: Node Stop
SYST	Node Stopping: A service initiated a graceful stop.	SYST: Node Stopping
SYSU	Node Start: A service started; the nature of the previous shutdown is indicated in the message.	SYSU: Node Start
VLST	User Initiated Volume Lost: The /proc/cmsi/Volume_Lost command was run.	VLST: User Initiated Volume Lost

Related information

[LKCU: Overwritten Object Cleanup](#)

Object storage audit messages

You should be familiar with audit messages belonging to the object storage audit category. These are events related to the storage and management of objects within the StorageGRID system. These include object storage and retrievals, grid-node to grid-node transfers, and verifications.

Code	Description	See
APCT	Archive Purge from Cloud-Tier: Archived object data is deleted from an external archival storage system, which connects to the StorageGRID through the S3 API.	APCT: Archive Purge from Cloud-Tier
ARCB	Archive Object Retrieve Begin: The ARC service begins the retrieval of object data from the external archival storage system.	ARCB: Archive Object Retrieve Begin

Code	Description	See
ARCE	Archive Object Retrieve End: Object data has been retrieved from an external archival storage system, and the ARC service reports the status of the retrieval operation.	ARCE: Archive Object Retrieve End
ARCT	Archive Retrieve from Cloud-Tier: Archived object data is retrieved from an external archival storage system, which connects to the StorageGRID through the S3 API.	ARCT: Archive Retrieve from Cloud-Tier
AREM	Archive Object Remove: A content block was successfully or unsuccessfully deleted from the external archival storage system.	AREM: Archive Object Remove
ASCE	Archive Object Store End: A content block has been written to the external archival storage system, and the ARC service reports the status of the write operation.	ASCE: Archive Object Store End
ASCT	Archive Store Cloud-Tier: Object data is stored to an external archival storage system, which connects to the StorageGRID through the S3 API.	ASCT: Archive Store Cloud-Tier
ATCE	Archive Object Store Begin: Writing a content block to an external archival storage has started.	ATCE: Archive Object Store Begin
AVCC	Archive Validate Cloud-Tier Configuration: The account and bucket settings provided were successfully or unsuccessfully validated.	AVCC: Archive Validate Cloud-Tier Configuration
CBSE	Object Send End: The source entity completed a grid-node to grid-node data transfer operation.	CBSE: Object Send End

Code	Description	See
CBRE	Object Receive End: The destination entity completed a grid-node to grid-node data transfer operation.	CBRE: Object Receive End
SCMT	Object Store Commit: A content block was completely stored and verified, and can now be requested.	SCMT: Object Store Commit
SREM	Object Store Remove: A content block was deleted from a grid node, and can no longer be requested directly.	SREM: Object Store Remove

Client read audit messages

Client read audit messages are logged when an S3 or Swift client application makes a request to retrieve an object.

Code	Description	Used by	See
SGET	S3 GET: Logs a successful transaction to retrieve an object or list the objects in a bucket. Note: If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	SGET: S3 GET
SHEA	S3 HEAD: Logs a successful transaction to check for the existence of an object or bucket.	S3 client	SHEA: S3 HEAD
WGET	Swift GET: Logs a successful transaction to retrieve an object or list the objects in a container.	Swift client	WGET: Swift GET
WHEA	Swift HEAD: Logs a successful transaction to check for the existence of an object or container.	Swift client	WHEA: Swift HEAD

Client write audit messages

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

Code	Description	Used by	See
OVWR	Object Overwrite: Logs a transaction to overwrite one object with another object.	S3 clients Swift clients	OVWR: Object Overwrite
SDEL	S3 DELETE: Logs a successful transaction to delete an object or bucket. Note: If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	SDEL: S3 DELETE
SPOS	S3 POST: Logs a successful transaction to restore an object from AWS Glacier storage to a Cloud Storage Pool.	S3 client	SPOS: S3 POST
SPUT	S3 PUT: Logs a successful transaction to create a new object or bucket. Note: If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	SPUT: S3 PUT
SUPD	S3 Metadata Updated: Logs a successful transaction to update the metadata for an existing object or bucket.	S3 client	SUPD: S3 Metadata Updated
WDEL	Swift DELETE: Logs a successful transaction to delete an object or container.	Swift client	WDEL: Swift DELETE

Code	Description	Used by	See
WPUT	Swift PUT: Logs a successful transaction to create a new object or container.	Swift client	WPUT: Swift PUT

Management audit message

The Management category logs user requests to the Management API.

Code	Message title and description	See
MGAU	Management API audit message: A log of user requests.	MGAU: Management audit message

Audit messages

When system events occur, the StorageGRID system generates audit messages and records them in the audit log.

APCT: Archive Purge from Cloud-Tier

This message is generated when archived object data is deleted from an external archival storage system, which connects to the StorageGRID through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was deleted.
CSIZ	Content Size	The size of the object in bytes. Always returns 0.
RSLT	Result Code	Returns successful (SUCC) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the cloud-tier from which the object was deleted.

ARCB: Archive Object Retrieve Begin

This message is generated when a request is made to retrieve archived object data and the retrieval process begins. Retrieval requests are processed immediately, but can be reordered to improve efficiency of retrieval from linear media such as tape.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival storage system.
RSLT	Result	Indicates the result of starting the archive retrieval process. Currently defined value is: SUCS: The content request was received and queued for retrieval.

This audit message marks the time of an archive retrieval. It allows you to match the message with a corresponding ARCE end message to determine the duration of archive retrieval, and whether the operation was successful.

ARCE: Archive Object Retrieve End

This message is generated when an attempt by the Archive Node to retrieve object data from an external archival storage system completes. If successful, the message indicates that the requested object data has been completely read from the archive location, and was successfully verified. After the object data has been retrieved and verified, it is delivered to the requesting service.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival storage system.
VLID	Volume Identifier	The identifier of the volume on which the data was archived. If an archive location for the content is not found, a Volume ID of 0 is returned.
RSLT	Retrieval Result	The completion status of the archive retrieval process: <ul style="list-style-type: none"> • SUCS: successful • VRFL: failed (object verification failure) • ARUN: failed (external archival storage system unavailable) • CANC: failed (retrieval operation canceled) • GERR: failed (general error)

Matching this message with the corresponding ARCB message can indicate the time taken to perform the

archive retrieval. This message indicates whether the retrieval was successful, and in the case of failure, the cause of the failure to retrieve the content block.

ARCT: Archive Retrieve from Cloud-Tier

This message is generated when archived object data is retrieved from an external archival storage system, which connects to the StorageGRID through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was retrieved.
CSIZ	Content Size	The size of the object in bytes. The value is only accurate for successful retrieves.
RSLT	Result Code	Returns successful (SUCC) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the external archival storage system.
TIME	Time	Total processing time for the request in microseconds.

AREM: Archive Object Remove

The Archive Object Remove audit message indicates that a content block was successfully or unsuccessfully deleted from an Archive Node. If the result is successful, the Archive Node has successfully informed the external archival storage system that StorageGRID has released an object location. Whether the object is removed from the external archive storage system depends on the type of system and its configuration.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival media system.
VLID	Volume Identifier	The identifier of the volume on which the object data was archived.

Code	Field	Description
RSLT	Result	<p>The completion status of the archive removal process:</p> <ul style="list-style-type: none"> • SUCS: successful • ARUN: failed (external archival storage system unavailable) • GERR: failed (general error)

ASCE: Archive Object Store End

This message indicates that writing a content block to an external archival storage system has ended.

Code	Field	Description
CBID	Content Block Identifier	The identifier of the content block stored on the external archival storage system.
VLID	Volume Identifier	The unique identifier of the archive volume to which the object data is written.
VREN	Verification Enabled	<p>Indicates if verification is performed for content blocks. Currently defined values are:</p> <ul style="list-style-type: none"> • VENA: verification is enabled • VDSA: verification is disabled
MCLS	Management Class	A string identifying the TSM Management Class to which the content block is assigned if applicable.

Code	Field	Description
RSLT	Result	<p>Indicates the result of the archive process. Currently defined values are:</p> <ul style="list-style-type: none"> • SUCS: successful (archiving process succeeded) • OFFL: failed (archiving is offline) • VRFL: failed (object verification failed) • ARUN: failed (external archival storage system unavailable) • GERR: failed (general error)

This audit message means that the specified content block has been written to the external archival storage system. If the write fails, the result provides basic troubleshooting information about where the failure occurred. More detailed information about archive failures can be found by examining Archive Node attributes in the StorageGRID system.

ASCT: Archive Store Cloud-Tier

This message is generated when archived object data is stored to an external archival storage system, which connects to StorageGRID through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was retrieved.
CSIZ	Content Size	The size of the object in bytes.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the cloud-tier the content was stored to.
TIME	Time	Total processing time for the request in microseconds.

ATCE: Archive Object Store Begin

This message indicates that writing a content block to an external archival storage has started.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the content block to be archived.
VLID	Volume Identifier	The unique identifier of the volume to which the content block is written. If the operation fails, a volume ID of 0 is returned.
RSLT	Result	Indicates the result of the transfer of the content block. Currently defined values are: <ul style="list-style-type: none"> • SUCS: success (content block stored successfully) • EXIS: ignored (content block was already stored) • ISFD: failed (insufficient disk space) • STER: failed (error storing the CBID) • OFFL: failed (archiving is offline) • GERR: failed (general error)

AVCC: Archive Validate Cloud-Tier Configuration

This message is generated when the configuration settings are validated for a Cloud Tiering - Simple Storage Service (S3) target type.

Code	Field	Description
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	UUID associated with the external archival storage system being validated.

CBRB: Object Receive Begin

During normal system operations, content blocks are continuously transferred between different nodes as data is accessed, replicated and retained. When transfer of a content block from one node to another is initiated, this message is issued by the destination entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the first sequence count requested. If successful, the transfer begins from this sequence count.
CTES	Expected End Sequence Count	Indicates the last sequence count requested. If successful, the transfer is considered complete when this sequence count has been received.
RSLT	Transfer Start Status	Status at the time the transfer was started: SUCS: Transfer started successfully.

This audit message means a node-to-node data transfer operation was initiated on a single piece of content, as identified by its Content Block Identifier. The operation requests data from "Start Sequence Count" to "Expected End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow, and when combined with storage audit messages, to verify replica counts.

CBRE: Object Receive End

When transfer of a content block from one node to another is completed, this message is issued by the destination entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.

Code	Field	Description
RSLT	Transfer Result	<p>The result of the transfer operation (from the perspective of the sending entity):</p> <p>SUCS: transfer successfully completed; all requested sequence counts were sent.</p> <p>CONL: connection lost during transfer</p> <p>CTMO: connection timed-out during establishment or transfer</p> <p>UNRE: destination node ID unreachable</p> <p>CRPT: transfer ended due to reception of corrupt or invalid data (might indicate tampering)</p>

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from "Start Sequence Count" to "Actual End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

CBSB: Object Send Begin

During normal system operations, content blocks are continuously transferred between different nodes as data is accessed, replicated and retained. When transfer of a content block from one node to another is initiated, this message is issued by the source entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	<p>Indicates if the CBID transfer was push-initiated or pull-initiated:</p> <p>PUSH: The transfer operation was requested by the sending entity.</p> <p>PULL: The transfer operation was requested by the receiving entity.</p>

Code	Field	Description
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the first sequence count requested. If successful, the transfer begins from this sequence count.
CTES	Expected End Sequence Count	Indicates the last sequence count requested. If successful, the transfer is considered complete when this sequence count has been received.
RSLT	Transfer Start Status	Status at the time the transfer was started: SUCC: transfer started successfully.

This audit message means a node-to-node data transfer operation was initiated on a single piece of content, as identified by its Content Block Identifier. The operation requests data from "Start Sequence Count" to "Expected End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow, and when combined with storage audit messages, to verify replica counts.

CBSE: Object Send End

When transfer of a content block from one node to another is completed, this message is issued by the source entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.

Code	Field	Description
CTDR	Transfer Direction	<p>Indicates if the CBID transfer was push-initiated or pull-initiated:</p> <p>PUSH: The transfer operation was requested by the sending entity.</p> <p>PULL: The transfer operation was requested by the receiving entity.</p>
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.
RSLT	Transfer Result	<p>The result of the transfer operation (from the perspective of the sending entity):</p> <p>SUCS: Transfer successfully completed; all requested sequence counts were sent.</p> <p>CONL: connection lost during transfer</p> <p>CTMO: connection timed-out during establishment or transfer</p> <p>UNRE: destination node ID unreachable</p> <p>CRPT: transfer ended due to reception of corrupt or invalid data (might indicate tampering)</p>

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from "Start Sequence Count" to "Actual End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system

data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

ECOC: Corrupt Erasure Coded Data Fragment

This audit message indicates that the system has detected a corrupt erasure-coded data fragment.

Code	Field	Description
VCCO	VCS ID	The name of the VCS that contains the corrupt chunk.
VLID	Volume ID	The RangeDB Volume that contains the corrupt erasure-coded fragment.
CCID	Chunk ID	The identifier of the corrupt erasure-coded fragment.
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this particular message. 'NONE' is used rather than 'SUCCS' so that this message is not filtered.

ETAF: Security Authentication Failed

This message is generated when a connection attempt using Transport Layer Security (TLS) has failed.

Code	Field	Description
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection over which the authentication failed.
RUID	User Identity	A service dependent identifier representing the identity of the remote user.

Code	Field	Description
RSLT	Reason Code	The reason for the failure: SCNI: Secure connection establishment failed. CERM: Certificate was missing. CERT: Certificate was invalid. CERE: Certificate was expired. CERR: Certificate was revoked. CSGN: Certificate signature was invalid. CSGU: Certificate signer was unknown. UCRM: User credentials were missing. UCRI: User credentials were invalid. UCRU: User credentials were disallowed. TOUT: Authentication timed out.

When a connection is established to a secure service that uses TLS, the credentials of the remote entity are verified using the TLS profile and additional logic built into the service. If this authentication fails due to invalid, unexpected, or disallowed certificates or credentials, an audit message is logged. This enables queries for unauthorized access attempts and other security-related connection problems.

The message could result from a remote entity having an incorrect configuration, or from attempts to present invalid or disallowed credentials to the system. This audit message should be monitored to detect attempts to gain unauthorized access to the system.

GNRG: GNDS Registration

The CMN service generates this audit message when a service has updated or registered information about itself in the StorageGRID system.

Code	Field	Description
RSLT	Result	The result of the update request: <ul style="list-style-type: none"> • SUCS: Successful • SUNV: Service Unavailable • GERR: Other failure
GNID	Node ID	The node ID of the service that initiated the update request.
GNTP	Device Type	The grid node's device type (for example, BLDR for an LDR service).
GNDV	Device Model version	The string identifying the grid node's device model version in the DMDL bundle.
GNGP	Group	The group to which the grid node belongs (in the context of link costs and service-query ranking).
GNIA	IP Address	The grid node's IP address.

This message is generated whenever a grid node updates its entry in the Grid Nodes Bundle.

GNUR: GNDS Unregistration

The CMN service generates this audit message when a service has unregistered information about itself from the StorageGRID system.

Code	Field	Description
RSLT	Result	The result of the update request: <ul style="list-style-type: none"> • SUCS: Successful • SUNV: Service Unavailable • GERR: Other failure
GNID	Node ID	The node ID of the service that initiated the update request.

GTED: Grid Task Ended

This audit message indicates that the CMN service has finished processing the specified grid task and has moved the task to the Historical table. If the result is SUCS, ABRT, or ROLF, there will be a corresponding Grid Task Started audit message. The other results

indicate that processing of this grid task never started.

Code	Field	Description
TSID	Task ID	<p>This field uniquely identifies a generated grid task and allows the grid task to be managed over its lifecycle.</p> <p>Note: The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.</p>
RSLT	Result	<p>The final status result of the grid task:</p> <ul style="list-style-type: none">• SUCS: The grid task completed successfully.• ABRT: The grid task was aborted without a rollback error.• ROLF: The grid task was aborted and was unable to complete the rollback process.• CANC: The grid task was canceled by the user before it was started.• EXPR: The grid task expired before it was started.• IVLD: The grid task was invalid.• AUTH: The grid task was unauthorized.• DUPL: The grid task was rejected as a duplicate.

GTST: Grid Task Started

This audit message indicates that the CMN service has started to process the specified grid task. The audit message immediately follows the Grid Task Submitted message for grid tasks initiated by the internal Grid Task Submission service and selected for automatic activation. For grid tasks submitted into the Pending table, this message is generated when the user starts the grid task.

Code	Field	Description
TSID	Task ID	<p>This field uniquely identifies a generated grid task and allows the task to be managed over its lifecycle.</p> <p>Note: The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.</p>
RSLT	Result	<p>The result. This field has only one value:</p> <ul style="list-style-type: none"> • SUCS: The grid task was started successfully.

GTSU: Grid Task Submitted

This audit message indicates that a grid task has been submitted to the CMN service.

Code	Field	Description
TSID	Task ID	<p>Uniquely identifies a generated grid task and allows the task to be managed over its lifecycle.</p> <p>Note: The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.</p>
TTYP	Task Type	The type of grid task.
TVER	Task Version	A number indicating the version of the grid task.
TDSC	Task Description	A human-readable description of the grid task.

Code	Field	Description
VATS	Valid After Timestamp	The earliest time (UINT64 microseconds from January 1, 1970 - UNIX time) at which the grid task is valid.
VBTS	Valid Before Timestamp	The latest time (UINT64 microseconds from January 1, 1970 - UNIX time) at which the grid task is valid.
TSRC	Source	The source of the task: <ul style="list-style-type: none"> • TXTB: The grid task was submitted through the StorageGRID system as a signed text block. • GRID: The grid task was submitted through the internal Grid Task Submission Service.
ACTV	Activation Type	The type of activation: <ul style="list-style-type: none"> • AUTO: The grid task was submitted for automatic activation. • PEND: The grid task was submitted into the pending table. This is the only possibility for the TXTB source.
RSLT	Result	The result of the submission: <ul style="list-style-type: none"> • SUCS: The grid task was submitted successfully. • FAIL: The task has been moved directly to the historical table.

IDEI: ILM Initiated Delete

This message is generated when ILM starts the process of deleting an object.

The IDEI message is generated in either of these situations:

- **For objects in compliant S3 buckets:** This message is generated when ILM starts the process of auto-deleting an object because its retention period has expired (assuming the auto-delete setting is enabled and legal hold is off).
- **For objects in non-compliant S3 buckets or Swift containers.** This message is generated when ILM

starts the process of deleting an object because no placement instructions in the active ILM policy currently apply to the object.

Code	Field	Description
CBID	Content Block Identifier	The CBID of the object.
CMPA	Compliance: Auto delete	For objects in compliant S3 buckets only. 0 (false) or 1 (true), indicating whether a compliant object should be deleted automatically when its retention period ends, unless the bucket is under a legal hold.
CMPL	Compliance: Legal hold	For objects in compliant S3 buckets only. 0 (false) or 1 (true), indicating whether the bucket is currently under a legal hold.
CMPR	Compliance: Retention period	For objects in compliant S3 buckets only. The length of the object's retention period in minutes.
CTME	Compliance: Ingest time	For objects in compliant S3 buckets only. The object's ingest time. You can add the retention period in minutes to this value to determine when the object can be deleted from the bucket.
DMRK	Delete Marker Version ID	The version ID of the delete marker created when deleting an object from a versioned bucket. Operations on buckets do not include this field.
CSIZ	Content size	The size of the object in bytes.

Code	Field	Description
LOCS	Locations	<p>The storage location of object data within the StorageGRID system. The value for LOCS is "" if the object has no locations (for example, it has been deleted).</p> <p>CLEC: for erasure-coded objects, the erasure coding profile ID and the erasure coding group ID that is applied to the object's data.</p> <p>CLDI: for replicated objects, the LDR node ID and the volume ID of the object's location.</p> <p>CLNL: ARC node ID of the object's location if the object data is archived.</p>
PATH	S3 Bucket/Key or Swift Container/Object ID	The S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
RSLT	Result	<p>The result of the ILM operation.</p> <p>SUCS: The ILM operation was successful.</p>
RULE	Rules Label	<ul style="list-style-type: none"> If an object in a compliant S3 bucket is being deleted automatically because its retention period has expired, this field is blank. If the object is being deleted because there are no more placement instructions that currently apply to the object, this field shows the human-readable label of the last ILM rule that applied to the object.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was deleted. Operations on buckets and objects in unversioned buckets do not include this field.

LKCU: Overwritten Object Cleanup

This message is generated when StorageGRID removes an overwritten object that previously required cleanup to free up storage space. An object is overwritten when an S3 or Swift client writes an object to a path already containing a object. The removal process occurs automatically and in the background.

Code	Field	Description
CSIZ	Content size	The size of the object in bytes.
LTYP	Type of cleanup	<i>Internal use only.</i>
LUID	Removed Object UUID	The identifier of the object that was removed.
PATH	S3 Bucket/Key or Swift Container/Object ID	The S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
SEGC	Container UUID	UUID of the container for the segmented object. This value is available only if the object is segmented.
UUID	Universally Unique Identifier	The identifier of the object that still exists. This value is available only if the object has not been deleted.

LLST: Location Lost

This message is generated whenever a location for an object copy (replicated or erasure coded) cannot be found.

Code	Field	Description
CBIL	CBID	The affected CBID.
NOID	Source Node ID	The node ID on which the locations were lost.
UUID	Universally Unique ID	The identifier of the affected object in the StorageGRID system.
ECPR	Erasure Coding Profile	For erasure-coded object data. The ID of the Erasure Coding profile used.

Code	Field	Description
LTYP	Location Type	CLDI (Online): For replicated object data CLEC (Online): For erasure-coded object data CLNL (Nearline): For archived replicated object data
PCLD	Path to replicated object	The complete path to the disk location of the lost object data. Only returned when LTYP has a value of CLDI (that is, for replicated objects). Takes the form /var/local/rangedb/2/p/13/13/00oJs6X%{h{U)SeUFxE@
RSLT	Result	Always NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.
TSRC	Triggering Source	USER: User triggered SYST: System triggered

MGAU: Management audit message

The Management category logs user requests to the Management API. Every request that is not a GET or HEAD request to the API logs a response with the username, IP, and type of request to the API.

Code	Field	Description
MDIP	Destination IP Address	The server (destination) IP address.
MDNA	Domain name	The host domain name.
MPAT	Request PATH	The request path.
MPQP	Request query parameters	The query parameters for the request.

Code	Field	Description
MRBD	Request body	<p>The content of the request body. While the response body is logged by default, the request body is logged in certain cases when the response body is empty. Because the following information is not available in the response body, it is taken from the request body for the following POST methods:</p> <ul style="list-style-type: none"> • Username and account ID in POST authorize • New subnets configuration in POST /grid/grid-networks/update • New NTP servers in POST /grid/ntp-servers/update • Decommissioned server IDs in POST /grid/servers/decommission <p>Note: Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).</p>
MRMD	Request method	<p>The HTTP request method:</p> <ul style="list-style-type: none"> • POST • PUT • DELETE • PATCH
MRSC	Response code	The response code.
MRSP	Response body	<p>The content of the response (the response body) is logged by default.</p> <p>Note: Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).</p>
MSIP	Source IP address	The client (source) IP address.

Code	Field	Description
MUUN	User URN	The URN (uniform resource name) of the user who sent the request.
RSLT	Result	Returns successful (SUCS) or the error reported by the backend.

OLST: System Detected Lost Object

This message is generated when the DDS service cannot locate any copies of an object within the StorageGRID system.

Code	Field	Description
CBID	Content Block Identifier	The CBID of the lost object.
NOID	Node ID	If available, the last known direct or nearline location of the lost object. It is possible to have just the Node ID without a Volume ID if the volume information is not available.
PATH	S3 Bucket/Key or Swift Container/Object ID	If available, the S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.
UUID	Universally Unique ID	The identifier of the lost object within the StorageGRID system.
VOLI	Volume ID	If available, the Volume ID of the Storage Node or Archive Node for the last known location of the lost object.

ORLM: Object Rules Met

This message is generated when the object is successfully stored and copied as specified by the ILM rules.



The ORLM message is not generated when an object is successfully stored by the default Make 2 Copies rule if another rule in the policy uses the Object Size advanced filter.

Code	Field	Description
CBID	Content Block Identifier	The CBID of the object.
CSIZ	Content size	The size of the object in bytes.
LOCS	Locations	<p>The storage location of object data within the StorageGRID system. The value for LOCS is "" if the object has no locations (for example, it has been deleted).</p> <p>CLEC: for erasure-coded objects, the erasure coding profile ID and the erasure coding group ID that is applied to the object's data.</p> <p>CLDI: for replicated objects, the LDR node ID and the volume ID of the object's location.</p> <p>CLNL: ARC node ID of the object's location if the object data is archived.</p>
PATH	S3 Bucket/Key or Swift Container/Object ID	The S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
RSLT	Result	<p>The result of the ILM operation.</p> <p>SUCS: The ILM operation was successful.</p>
RULE	Rules Label	The human-readable label given to the ILM rule applied to this object.
SEGC	Container UUID	UUID of the container for the segmented object. This value is available only if the object is segmented.
SGCB	Container CBID	CBID of the container for the segmented object. This value is available only if the object is segmented.

Code	Field	Description
STAT	Status	<p>The status of ILM operation.</p> <p>DONE: ILM operations against the object have completed.</p> <p>DFER: The object has been marked for future ILM re-evaluation.</p> <p>PRGD: The object has been deleted from the StorageGRID system.</p> <p>NLOC: The object data can no longer be found in the StorageGRID system. This status might indicate that all copies of object data are missing or damaged.</p>
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.

The ORLM audit message can be issued a number of times for a single object. For instance, it is issued whenever one of the following events take place:

- ILM rules for the object are satisfied forever.
- ILM rules for the object are satisfied for this epoch.
- ILM rules have deleted the object.
- The background verification process detects that a copy of replicated object data is corrupt. The StorageGRID system performs an ILM evaluation to replace the corrupt object.

Related information

[Object ingest transactions](#)

[Object delete transactions](#)

OVWR: Object Overwrite

This message is generated when an external (client-requested) operation causes one object to be overwritten by another object.

Code	Field	Description
CBID	Content Block Identifier (new)	The CBID for the new object.
CSIZ	Previous Object Size	The size, in bytes, of the object being overwritten.

Code	Field	Description
OCBD	Content Block Identifier (previous)	The CBID for the previous object.
UUID	Universally Unique ID (new)	The identifier of the new object within the StorageGRID system.
OUID	Universally Unique ID (previous)	The identifier for the previous object within the StorageGRID system.
PATH	S3 or Swift Object Path	The S3 or Swift object path used for both the previous and new object
RSLT	Result Code	Result of the Object Overwrite transaction. Result is always: SUCS: Successful

SADD: Security Audit Disable

This message indicates that the originating service (node ID) has turned off audit message logging; audit messages are no longer being collected or delivered.

Code	Field	Description
AETM	Enable Method	The method used to disable the audit.
AEUN	User Name	The user name that executed the command to disable audit logging.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.

The message implies that logging was previously enabled, but has now been disabled. This is typically used only during bulk ingest to improve system performance. Following the bulk activity, auditing is restored (SADE) and the capability to disable auditing is then permanently blocked.

SADE: Security Audit Enable

This message indicates that the originating service (node ID) has restored audit message logging; audit messages are again being collected and delivered.

Code	Field	Description
AETM	Enable Method	The method used to enable the audit.
AEUN	User Name	The user name that executed the command to enable audit logging.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.

The message implies that logging was previously disabled (SADD), but has now been restored. This is typically only used during bulk ingest to improve system performance. Following the bulk activity, auditing is restored and the capability to disable auditing is then permanently blocked.

SCMT: Object Store Commit

Grid content is not made available or recognized as stored until it has been committed (meaning it has been stored persistently). Persistently stored content has been completely written to disk, and has passed related integrity checks. This message is issued when a content block is committed to storage.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block committed to permanent storage.
RSLT	Result Code	Status at the time the object was stored to disk: SUCS: Object successfully stored.

This message means a given content block has been completely stored and verified, and can now be requested. It can be used to track data flow within the system.

SDEL: S3 DELETE

When an S3 client issues a DELETE transaction, a request is made to remove the specified object or bucket. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the deleted object in bytes. Operations on buckets do not include this field.
DMRK	Delete Marker Version ID	The version ID of the delete marker created when deleting an object from a versioned bucket. Operations on buckets do not include this field.
HTRH	HTTP Request Header	<p>List of logged HTTP request header names and values as selected during configuration.</p> <p>Note: X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field).</p>
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	<p>Result of the DELETE transaction. Result is always:</p> <p>SUCS: Successful</p>
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.

Code	Field	Description
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: urn:sgws:identity::03393893651506583485:root Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.

Code	Field	Description
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was deleted. Operations on buckets and objects in unversioned buckets do not include this field.

SGET: S3 GET

When an S3 client issues a GET transaction, a request is made to retrieve an object or list the objects in a bucket. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets do not include this field.
HTRH	HTTP Request Header	<p>List of logged HTTP request header names and values as selected during configuration.</p> <p>Note: X-Forwarded-For is automatically included if it is present in the request and if the x-Forwarded-For value is different from the request sender IP address (SAIP audit field).</p>

Code	Field	Description
RANG	Range Read	For range read operations only. Indicates the range of bytes that was read by this request. The value after the slash (/) shows the size of the entire object.
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.

Code	Field	Description
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: urn:sgws:identity::03393893651506583485:root Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets do not include this field.

SHEA: S3 HEAD

When an S3 client issues a HEAD transaction, a request is made to check for the existence of an object or bucket and retrieve the metadata about an object. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the checked object in bytes. Operations on buckets do not include this field.

Code	Field	Description
HTRH	HTTP Request Header	<p>List of logged HTTP request header names and values as selected during configuration.</p> <p>Note: X-Forwarded-For is automatically included if it is present in the request and if the x-Forwarded-For value is different from the request sender IP address (SAIP audit field).</p>
RSLT	Result Code	<p>Result of the GET transaction.</p> <p>Result is always:</p> <p>SUCS: Successful</p>
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.

Code	Field	Description
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: urn:sgws:identity::03393893651506583485:root Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets do not include this field.

SPOS: S3 POST

When an S3 client issues a POST Object restore request, a request is made to restore an object from AWS Glacier storage to a Cloud Storage Pool. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes.

Code	Field	Description
HTRH	HTTP Request Header	<p>List of logged HTTP request header names and values as selected during configuration.</p> <p>Note: X-Forwarded-For is automatically included if it is present in the request and if the x-Forwarded-For value is different from the request sender IP address (SAIP audit field).</p>
RSLT	Result Code	<p>Result of the POST Object restore request. Result is always:</p> <p>SUCS: Successful</p>
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.

Code	Field	Description
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SRCF	Subresource Configuration	Restore information.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: urn:sgws:identity::03393893651506583485:root Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets do not include this field.

SPUT: S3 PUT

When an S3 client issues a PUT transaction, a request is made to create a new object or bucket. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.

Code	Field	Description
CMPS	Compliance Settings	The compliance settings used when creating the bucket, if present in the PUT Bucket request (truncated to the first 1024 characters)
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets do not include this field.
HTRH	HTTP Request Header	<p>List of logged HTTP request header names and values as selected during configuration.</p> <p>Note: X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field).</p>
LKEN	Object Lock Enabled	Value of the request header x-amz-bucket-object-lock-enabled, if present in the PUT Bucket request.
LKLH	Object Lock Legal Hold	Value of the request header x-amz-object-lock-legal-hold, if present in the PUT Object request.
LKMD	Object Lock Retention Mode	Value of the request header x-amz-object-lock-mode, if present in the PUT Object request.
LKRU	Object Lock Retain Until Date	Value of the request header x-amz-object-lock-retain-until-date, if present in the PUT Object request.

Code	Field	Description
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	Result of the PUT transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3KY	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SRCF	Subresource Configuration	The new subresource configuration (truncated to the first 1024 characters).

Code	Field	Description
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: urn:sgws:identity::03393893651506583485:root Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
ULID	Upload ID	Included only in SPUT messages for Complete Multipart Upload operations. Indicates that all parts have been uploaded and assembled.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of a new object created in a versioned bucket. Operations on buckets and objects in unversioned buckets do not include this field.
VSST	Versioning State	The new versioning state of a bucket. Two states are used: "enabled" or "suspended." Operations on objects do not include this field.

SREM: Object Store Remove

This message is issued when content is removed from persistent storage and is no longer accessible through regular APIs.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block deleted from permanent storage.

Code	Field	Description
RSLT	Result Code	Indicates the result of the content removal operations. The only defined value is: SUCS: Content removed from persistent storage

This audit message means a given content block has been deleted from a node and can no longer be requested directly. The message can be used to track the flow of deleted content within the system.

SUPD: S3 Metadata Updated

This message is generated by the S3 API when an S3 client updates the metadata for an ingested object. The message is issued by the server if the metadata update is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets do not include this field.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request, when updating a bucket's compliance settings.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. Note: X-Forwarded-For is automatically included if it is present in the request and if the x-Forwarded-For value is different from the request sender IP address (SAIP audit field).

Code	Field	Description
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets do not include this field.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: urn:sgws:identity::0339389 3651506583485:root Empty for anonymous requests.

Code	Field	Description
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object whose metadata was updated. Operations on buckets and objects in unversioned buckets do not include this field.

SVRF: Object Store Verify Fail

This message is issued whenever a content block fails the verification process. Each time replicated object data is read from or written to disk, several verification and integrity checks are performed to ensure the data sent to the requesting user is identical to the data originally ingested into the system. If any of these checks fail, the system automatically quarantines the corrupt replicated object data to prevent it from being retrieved again.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block which failed verification.

Code	Field	Description
RSLT	Result Code	Verification failure type: CRCF: Cyclic redundancy check (CRC) failed. HMAC: Hash-based message authentication code (HMAC) check failed. EHSH: Unexpected encrypted content hash. PHS: Unexpected original content hash. SEQC: Incorrect data sequence on disk. PERR: Invalid structure of disk file. DERR: Disk error. FNAM: Bad file name.

Note: This message should be monitored closely. Content verification failures can indicate attempts to tamper with content or impending hardware failures.

To determine what operation triggered the message, see the value of the AMID (Module ID) field. For example, an SVFY value indicates that the message was generated by the Storage Verifier module, that is, background verification, and STOR indicates that the message was triggered by content retrieval.

SVRU: Object Store Verify Unknown

The LDR service's Storage component continuously scans all copies of replicated object data in the object store. This message is issued when an unknown or unexpected copy of replicated object data is detected in the object store and moved to the quarantine directory.

Code	Field	Description
FPTH	File Path	The file path of the unexpected object copy.
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

Note: The SVRU: Object Store Verify Unknown audit message should be monitored closely. It means unexpected copies of object data were detected in the object store. This situation should be investigated immediately to determine how these copies were created, because it can indicate attempts to tamper with content or impending hardware failures.

SYSD: Node Stop

When a service is stopped gracefully, this message is generated to indicate the shutdown was requested. Typically this message is sent only after a subsequent restart, because the audit message queue is not cleared prior to shutdown. Look for the SYST message, sent at the beginning of the shutdown sequence, if the service has not restarted.

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shutdown.

The message does not indicate if the host server is being stopped, only the reporting service. The RSLT of a SYSD cannot indicate a "dirty" shutdown, because the message is generated only by "clean" shutdowns.

SYST: Node Stopping

When a service is gracefully stopped, this message is generated to indicate the shutdown was requested and that the service has initiated its shutdown sequence. SYST can be used to determine if the shutdown was requested, before the service is restarted (unlike SYSD, which is typically sent after the service restarts.)

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shutdown.

The message does not indicate if the host server is being stopped, only the reporting service. The RSLT code of a SYST message cannot indicate a "dirty" shutdown, because the message is generated only by "clean" shutdowns.

SYSU: Node Start

When a service is restarted, this message is generated to indicate if the previous shutdown was clean (commanded) or disorderly (unexpected).

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shut down. DSDN: System was not cleanly shut down. VRGN: System was started for the first time after server installation (or re-installation).

The message does not indicate if the host server was started, only the reporting service. This message can be used to:

- Detect discontinuity in the audit trail.
- Determine if a service is failing during operation (as the distributed nature of the StorageGRID system can mask these failures). Server Manager restarts a failed service automatically.

VLST: User Initiated Volume Lost

This message is issued whenever the `/proc/CMSI/Volume_Lost` command is run.

Code	Field	Description
VOLL	Volume Identifier Lower	The lower end of the affected volume range or a single volume.
VOLU	Volume Identifier Upper	The upper end of the affected volume range. Equal to VOLL if a single volume.
NOID	Source Node ID	The node ID on which the locations were lost.
LTYP	Location Type	'CLDI' (Online) or 'CLNL' (Nearline). If not specified, defaults to 'CLDI'.
RSLT	Result	Always 'NONE'. RSLT is a mandatory message field, but is not relevant for this message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

WDEL: Swift DELETE

When a Swift client issues a DELETE transaction, a request is made to remove the

specified object or container. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on containers do not include this field.
CSIZ	Content Size	The size of the deleted object in bytes. Operations on containers do not include this field.
HTRH	HTTP Request Header	<p>List of logged HTTP request header names and values as selected during configuration.</p> <p>Note: X-Forwarded-For is automatically included if it is present in the request and if the x-Forwarded-For value is different from the request sender IP address (SAIP audit field).</p>
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	<p>Result of the DELETE transaction. Result is always:</p> <p>SUCS: Successful</p>
SAIP	IP address of requesting client	The IP address of the client application that made the request.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.

Code	Field	Description
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers do not include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

WGET: Swift GET

When a Swift client issues a GET transaction, a request is made to retrieve an object, list the objects in a container, or list the containers in an account. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on accounts and containers do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers do not include this field.
HTRH	HTTP Request Header	<p>List of logged HTTP request header names and values as selected during configuration.</p> <p>Note: X-Forwarded-For is automatically included if it is present in the request and if the x-Forwarded-For value is different from the request sender IP address (SAIP audit field).</p>
RSLT	Result Code	<p>Result of the GET transaction. Result is always SUCS: successful</p>
SAIP	IP address of requesting client	The IP address of the client application that made the request.

Code	Field	Description
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name. Operations on accounts do not include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers do not include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

WHEA: Swift HEAD

When a Swift client issues a HEAD transaction, a request is made to check for the existence of an account, container, or object, and retrieve any relevant metadata. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on accounts and containers do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers do not include this field.

Code	Field	Description
HTRH	HTTP Request Header	<p>List of logged HTTP request header names and values as selected during configuration.</p> <p>Note: X-Forwarded-For is automatically included if it is present in the request and if the x-Forwarded-For value is different from the request sender IP address (SAIP audit field).</p>
RSLT	Result Code	<p>Result of the HEAD transaction.</p> <p>Result is always:</p> <p>SUCS: successful</p>
SAIP	IP address of requesting client	The IP address of the client application that made the request.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name. Operations on accounts do not include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers do not include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

WPUT: Swift PUT

When a Swift client issues a PUT transaction, a request is made to create a new object or

container. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on containers do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on containers do not include this field.
HTRH	HTTP Request Header	<p>List of logged HTTP request header names and values as selected during configuration.</p> <p>Note: X-Forwarded-For is automatically included if it is present in the request and if the x-Forwarded-For value is different from the request sender IP address (SAIP audit field).</p>
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	<p>Result of the PUT transaction. Result is always:</p> <p>SUCS: successful</p>
SAIP	IP address of requesting client	The IP address of the client application that made the request.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.

Code	Field	Description
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers do not include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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