



StorageGRID® 11.3

Monitoring and Troubleshooting Guide

January 2020 | 215-14213_2020-01_en-us
doccomments@netapp.com

Contents

Monitoring a StorageGRID system	6
Web browser requirements	6
Viewing the Dashboard	6
Using the Nodes page	8
Viewing the Overview tab	9
Viewing the Hardware tab	10
Viewing the Network tab	11
Viewing the Storage tab	12
Viewing the Events tab	13
Using the Task tab to reboot a grid node	15
Viewing the Objects tab	16
Viewing the ILM tab	17
Viewing the Load Balancer tab	18
Viewing the Platform Services tab	20
Viewing information about appliance Storage Nodes	21
Viewing information about appliance Admin Nodes and Gateway Nodes ...	29
Viewing the Grid Topology tree	35
Information you should monitor regularly	36
Viewing node icons	36
Viewing and acknowledging alarms	37
Monitoring storage capacity	39
Monitoring the recovery point objective through ILM	46
Monitoring object verification operations	48
Monitoring the SSM service	49
Monitoring the Total Events (SMTT) alarm	50
Monitoring archival capacity	51
Using reports	52
Types of charts	52
Chart legend	53
Displaying charts	54
Generating charts	55
Types of text reports	56
Generating text reports	57
Exporting text reports	58
Managing alarms	60
Alarm classes	60
Alarm triggering logic	61
Alarm triggering examples	61
Alarms of same severity	63
Alarm class overrides	64
Severity changes	64

Notifications	64
New services	65
Alarms and tables	65
Viewing Default alarms	65
Creating custom service or component alarms	66
Creating Global Custom alarms	68
Disabling alarms	71
Disabling a Default alarm system wide	71
Disabling Default alarms for services	72
Disabling Global Custom alarms system wide	73
Disabling Global Custom alarms for services	74
Clearing triggered alarms	75
Configuring email notifications for alarms	75
Types of alarm notifications	76
Configuring email server settings for alarms	76
Creating email templates	78
Creating mailing lists	79
Configuring global email notifications	80
Suppressing email notifications for a mailing list	81
Suppressing email notifications system wide	82
Managing alerts (preview mode for 11.3)	84
What alerts are	84
Viewing all alerts	87
Viewing a specific alert	90
Managing alert rules	92
Viewing alert rules	92
Creating custom alert rules	94
Editing alert rules	96
Removing custom alert rules	99
Managing alert notifications	99
Setting up alert notifications	100
Information included in alert notifications	102
How StorageGRID groups alerts in notifications	103
Silencing alert notifications	103
Troubleshooting alert notifications	105
Alerts reference	106
Commonly used Prometheus metrics	110
Troubleshooting a StorageGRID system	114
Overview of problem determination	114
Defining the problem	115
Assessing the risk and impact on the system	115
Collecting data	115
Analyzing data	135
Escalation information checklist	136
Troubleshooting Admin Nodes	137

Troubleshooting sign-on errors	137
Checking the status of an unavailable Admin Node	140
Troubleshooting Storage Nodes	141
Object store (storage volume) failures	141
Troubleshooting SAVP Total Usable Space (Percent) alarm	142
Troubleshooting Total Events (SMTT) alarms	143
Troubleshooting Storage Status (SSTS) alarms	144
Troubleshooting Low object data storage alerts	148
Lost and missing object data	149
Troubleshooting SVST (Services: Status - Cassandra) alarm	157
Other StorageGRID troubleshooting tips	160
User interface issues	160
Time synchronization	161
Network connectivity	161
Linux: Node status is “orphaned”	161
Linux: Enabling IPv6 support in the kernel	162
Linux: Changing trigger values for CPU Load Average	164
Alarms reference	166
Log files	191
StorageGRID software logs	191
Deployment and maintenance logs	194
Logs for third-party software	195
About the bycast.log	196
File rotation for bycast.log	196
Messages in bycast.log	196
Message severities in bycast.log	197
Error codes in bycast.log	197
Copyright	202
Trademark	203
How to send comments about documentation and receive update notifications	204

Monitoring a StorageGRID system

You can use the Grid Manager to monitor the daily activities of your StorageGRID system, including its health.

Web browser requirements

You must use a supported web browser.

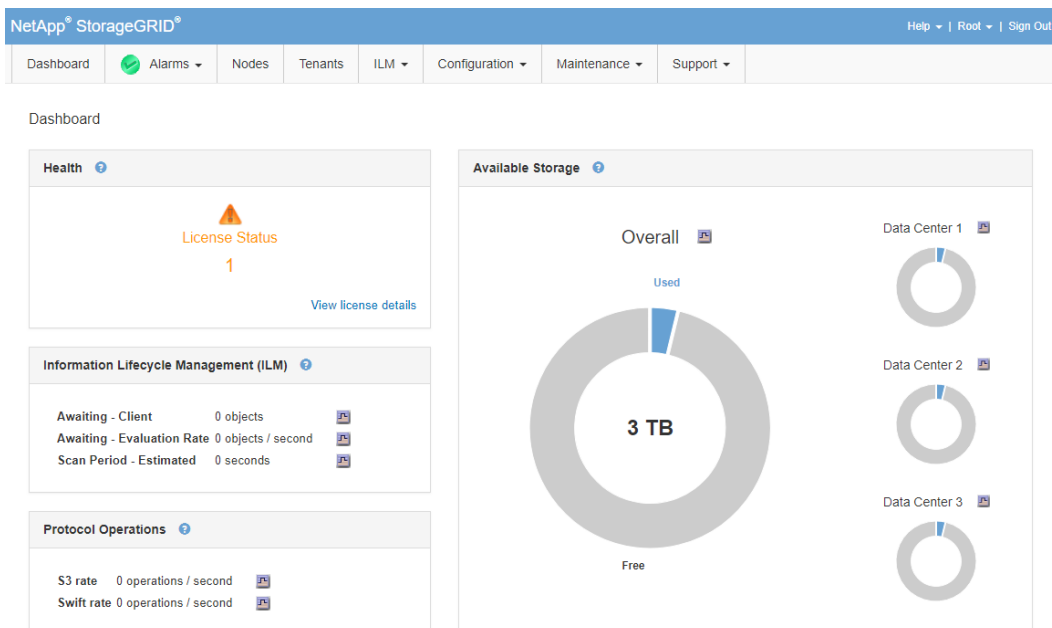
Web browser	Minimum supported version
Google Chrome	74
Microsoft Internet Explorer	11 (Native Mode)
Mozilla Firefox	67

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Viewing the Dashboard

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



For an explanation of the information on each panel, click the help icon (?) for the panel.

Panel	Description	View additional details	Learn more
Health	<p>Provides an indication of the system's health by showing:</p> <ul style="list-style-type: none"> The number of disconnected grid nodes, if any. The number of current alarms, if any. License status if there are license-related alarms. 	<p>When issues exist, links appear that allow you to view additional details:</p> <ul style="list-style-type: none"> To see details for any disconnected grid nodes, select View grid node details. To see details for any current alarms, select View current alarms. To see the current license or for more information about a license-related alarm, select View license details. 	<ul style="list-style-type: none"> Using the Nodes page Viewing and acknowledging alarms on page 37 Managing alarms on page 60 Alarms reference on page 166
Available Storage	<p>Displays the available and used storage capacity in the entire grid, not including archival media.</p> <p>The Overall chart presents grid-wide totals. If this is a multi-site grid, additional charts appear for each data center site.</p> <p>You can use this information to compare the used storage with the available storage. If you have a multi-site grid, you can determine which site is consuming more storage.</p>	<ul style="list-style-type: none"> To view the capacity, place your cursor over the chart's available and used capacity sections. To view capacity trends over a date range, click the chart icon for the overall grid, or for a data center site. To see details, select Nodes. Then, view the Storage tab for the entire grid, an entire site, or a single Storage Node. 	<ul style="list-style-type: none"> Viewing the Storage tab on page 12 Monitoring storage capacity on page 39 Administering StorageGRID
Information Lifecycle Management (ILM)	<p>Displays current ILM operations and ILM queues for your system.</p> <p>You can use this information to monitor your system's workload.</p>	<ul style="list-style-type: none"> To see details, select Nodes. Then, view the ILM tab for the entire grid, an entire site, or a single Storage Node. To see the existing ILM rules, select ILM > Rules. To see the existing ILM policies, select ILM > Policies. 	<ul style="list-style-type: none"> Viewing the ILM tab on page 17 Administering StorageGRID

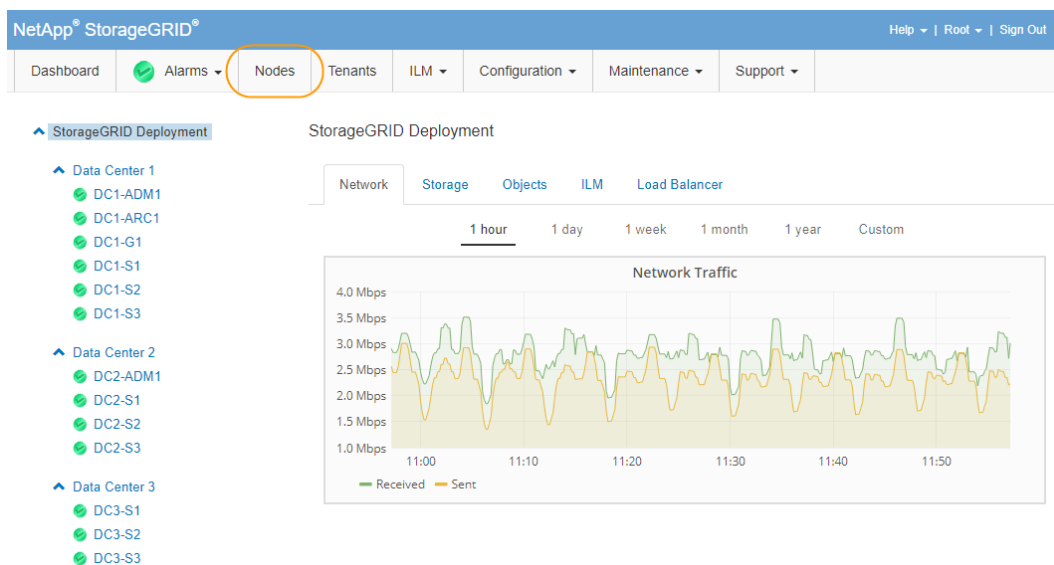
Panel	Description	View additional details	Learn more
Protocol Operations	Displays the number of protocol-specific operations (S3 and Swift) performed by your system. You can use this information to monitor your system's workloads and efficiencies. Protocol rates are averaged over the last two minutes.	<ul style="list-style-type: none"> To see details, select Nodes. Then, view the Objects tab for the entire grid, an entire site, or a single Storage Node. To view trends over a date range, click the chart icon to the right of the S3 or Swift protocol rate. 	<ul style="list-style-type: none"> Viewing the Objects tab on page 16 Administering StorageGRID Implementing S3 client applications Implementing Swift client applications

Using 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.

The Nodes page displays information about your entire StorageGRID system, each site in the grid, and each node within a site. 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 all sites. (The screenshot shows a grid 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.



The graphs on the Nodes page use the Grafana visualization tool and the Prometheus systems monitoring tool. Grafana displays time-series data in graph and chart formats, while Prometheus serves as the backend data source.

Choices

- [Viewing the Overview tab](#) on page 9
- [Viewing the Hardware tab](#) on page 10
- [Viewing the Network tab](#) on page 11
- [Viewing the Storage tab](#) on page 12
- [Viewing the Events tab](#) on page 13
- [Using the Task tab to reboot a grid node](#) on page 15
- [Viewing the Objects tab](#) on page 16
- [Viewing the ILM tab](#) on page 17
- [Viewing the Load Balancer tab](#) on page 18
- [Viewing the Platform Services tab](#) on page 20
- [Viewing information about appliance Storage Nodes](#) on page 21
- [Viewing information about appliance Admin Nodes and Gateway Nodes](#) on page 29

Viewing the Overview tab

The Overview tab provides basic information about each node. It also shows any current, unacknowledged alarms affecting the node.

The Overview tab is shown for all nodes.

Node Information

The Node Information section of the Overview tab provides general information about each node, including its ID (also referred to as the UUID), name, type, StorageGRID software version, and IP address for the Grid Network.


DC1-S1 (Storage Node)

[Overview](#)
[Hardware](#)
[Network](#)
[Storage](#)
[Objects](#)
[ILM](#)
[Events](#)
[Tasks](#)

Node Information ⓘ

ID	c452e1eb-7ba7-41cb-8439-e74672685f4f
Name	DC1-S1
Type	Storage Node
Software Version	11.3.0 (build 20190524.2116.9fa401d)
IP Addresses	10.96.98.113 Show more ▼

Alarms ⓘ



No unacknowledged alarms

In the **IP Addresses** field, click **Show more** to view the node's IPv4 and IPv6 addresses on the Grid Network (eth0), Admin Network (eth1), and Client Network (eth2).

Alarms

The Alarms section of the Overview tab includes a table listing any unacknowledged alarms for the node. You can view the details for an unacknowledged alarm or select the service name in the table to acknowledge the alarm.

Alarms ?							
Severity	Attribute Name	Code	Service	Description	Time Triggered	Trigger Value	Current Value
Major	Outbound Replication Status	ORSU	ARC	Storage Unavailable	2019-05-24 21:42:02 MDT	Storage Unavailable	Storage Unavailable

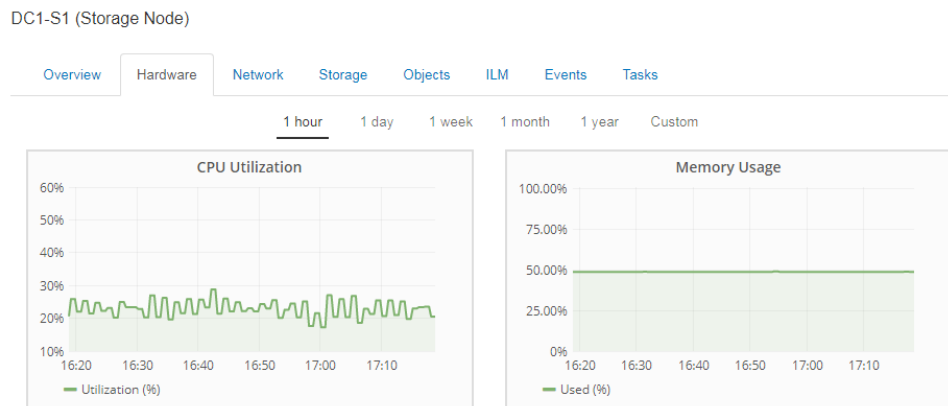
Related tasks

[Viewing and acknowledging alarms](#) on page 37

Viewing the Hardware tab

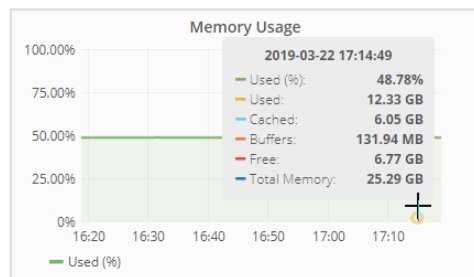
The Hardware tab displays CPU utilization and memory usage for each node, and additional hardware information about appliances.

The Hardware tab is shown for all nodes.



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

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



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

Related tasks

[Viewing information about appliance Storage Nodes](#) on page 21

[Viewing information about appliance Admin Nodes and Gateway Nodes](#) on page 29

Viewing the Network tab

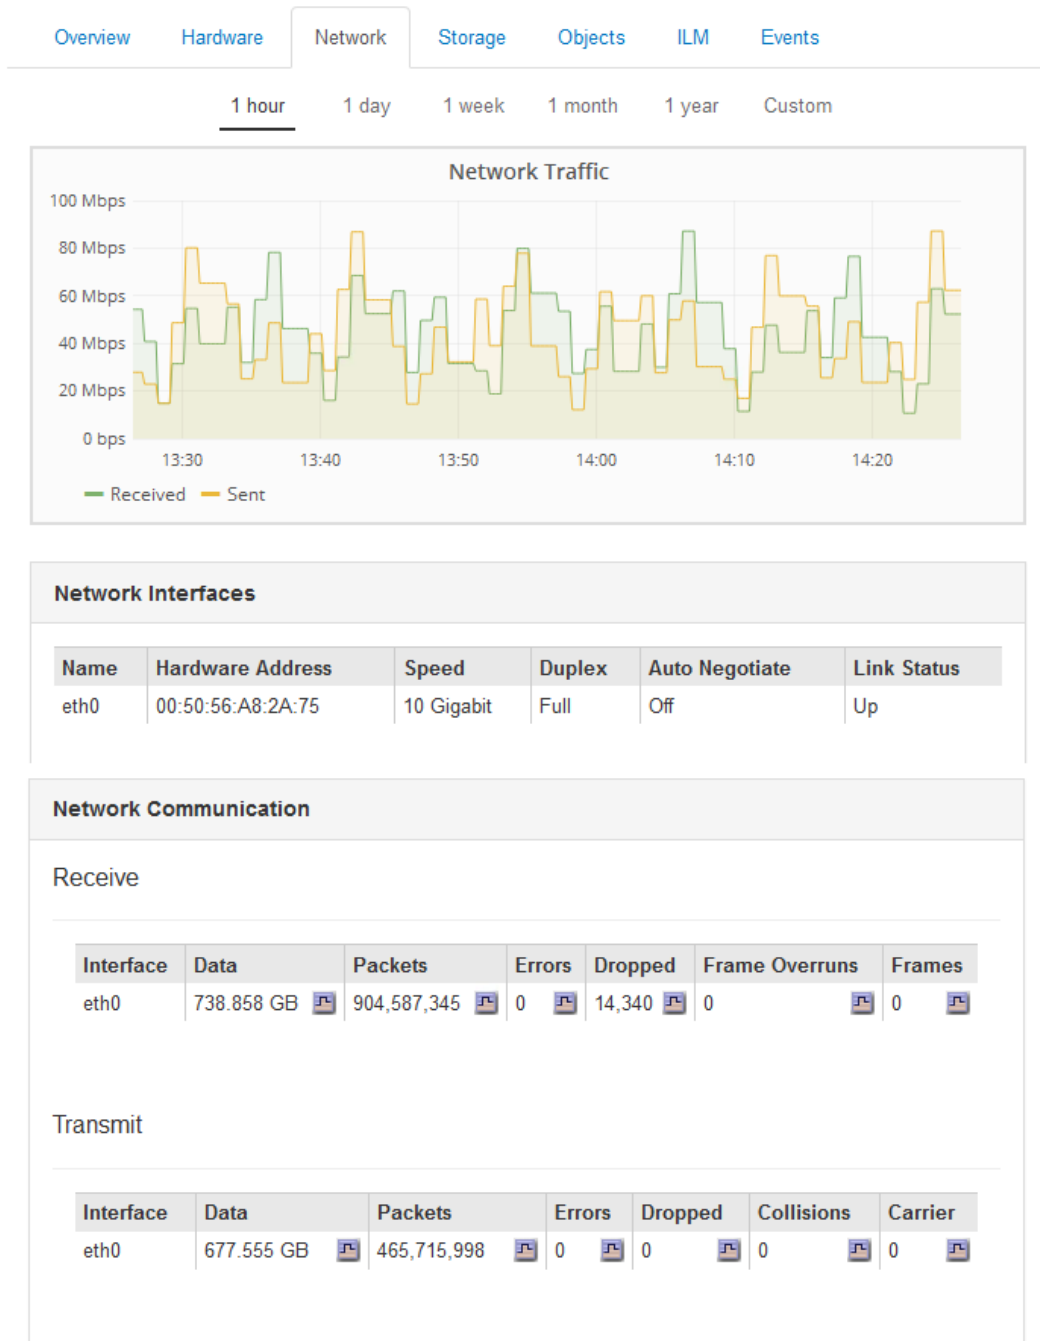
The Network tab displays a graph showing the network traffic received and sent across all of the network interfaces on the node, site, or grid.

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

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

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

DC1-S1-226 (Storage Node)



Viewing the Storage tab

The Storage tab summarizes storage availability and other storage metrics.






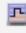



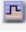

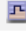
The Storage tab is shown for all nodes, each site, and the entire grid.

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




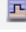




For Storage Nodes, each site, and the entire grid, the Storage tab also includes graphs showing object data storage and metadata storage used over time.

DC1-ADM1-225 (Admin Node)

[Overview](#)[Hardware](#)[Network](#)[Storage](#)[Events](#)**Disk Devices**

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

Volumes

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








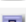





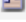

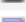
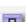







Viewing the Events tab


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 node types.


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.

Overview	Hardware	Network	Storage	Events
--------------------------	--------------------------	-------------------------	-------------------------	---------------

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

[Reset event counts](#) 

You can perform these tasks from the Events tab:

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

Related concepts

[Monitoring events](#) on page 119

Related tasks

[Displaying charts](#) on page 54

[Resetting event counts](#) on page 121

Using the Task tab to reboot a grid node

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

Before you begin

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

About this task

Rebooting a grid node from the Grid Manager issues the `reboot` command on the target node.

When you reboot a grid node, the node shuts down and restarts. All services are restarted automatically.

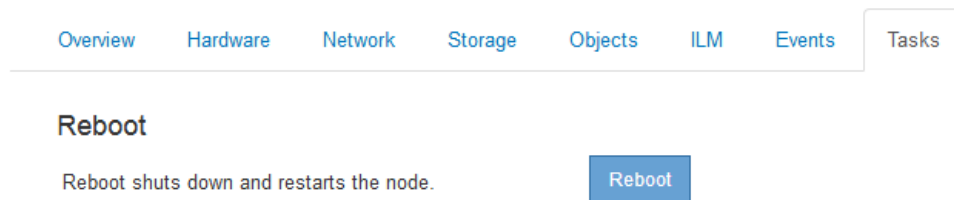
If you plan to reboot Storage Nodes, note the following:

- 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.

Steps

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

DC3-S3 (Storage Node)



4. Click **Reboot**.

A confirmation dialog box appears.

⚠ Reboot Node DC3-S3

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.

Provisioning Passphrase

Cancel

OK

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

5. Enter the provisioning passphrase, and click **OK**.

6. Wait for the node to reboot.

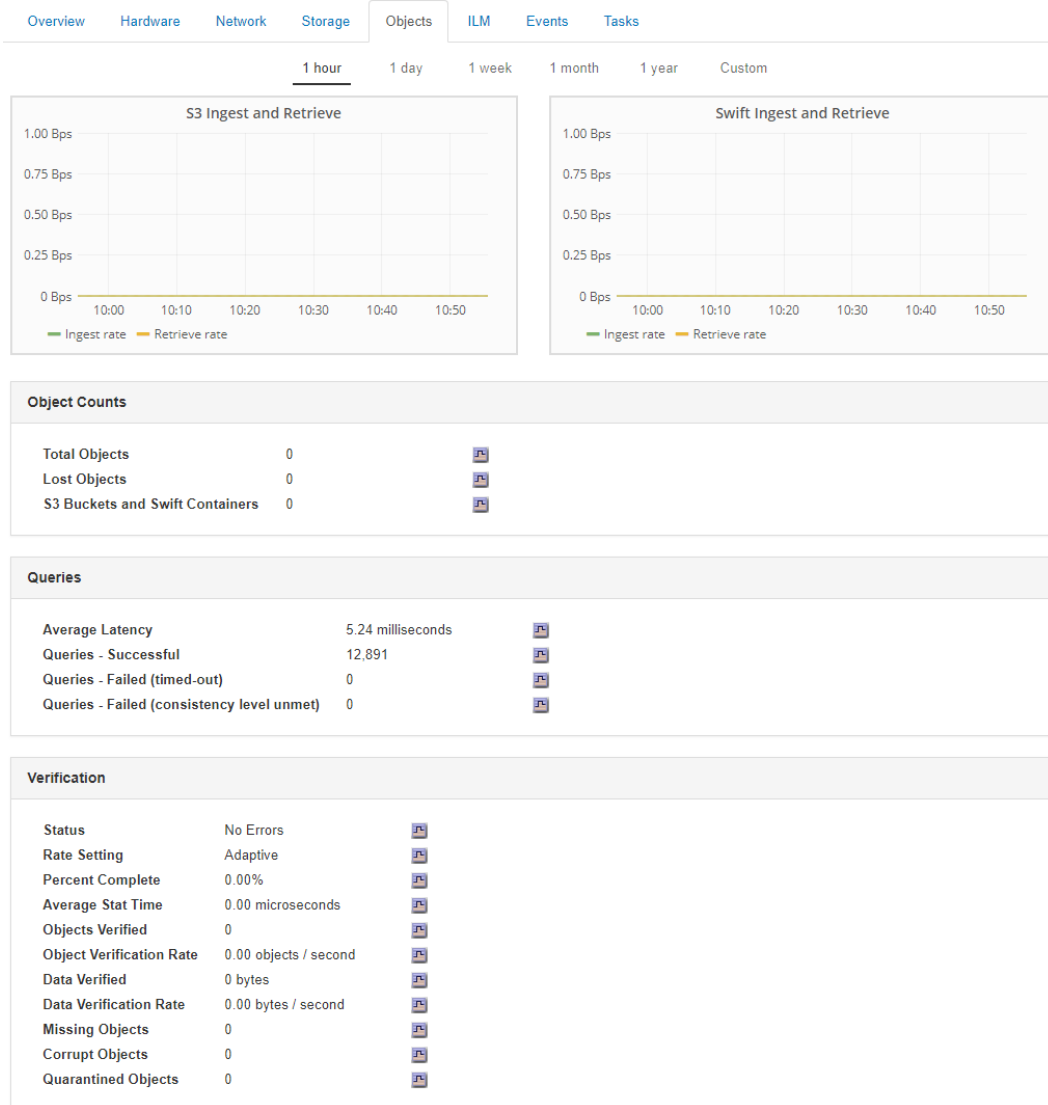
It might take some time for services to shut down.

When the node is rebooting, the gray icon (Administratively Down) appears on the left side of the Nodes page. When all services have started again, the icon changes back to its original color.

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 store queries and background verification.



Related information

[Implementing S3 client applications](#)

[Implementing Swift client applications](#)

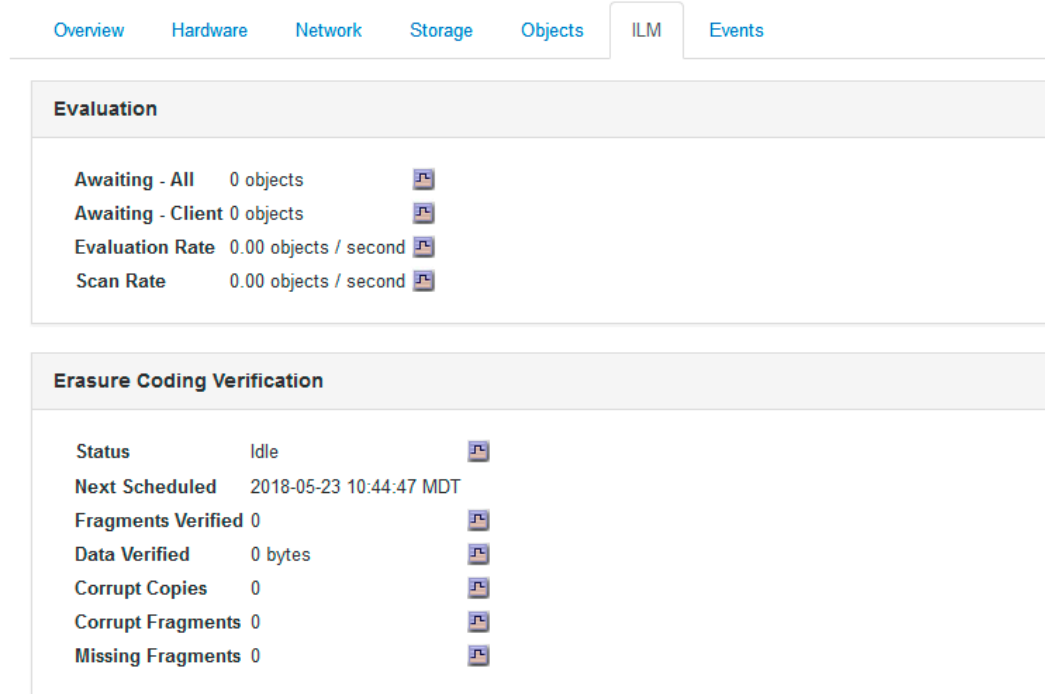
Viewing the ILM tab

The ILM tab provides information about Information Lifecycle Management (ILM) operations.

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

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

DC1-S1 (Storage Node)

**Related information**

[Administering StorageGRID](#)

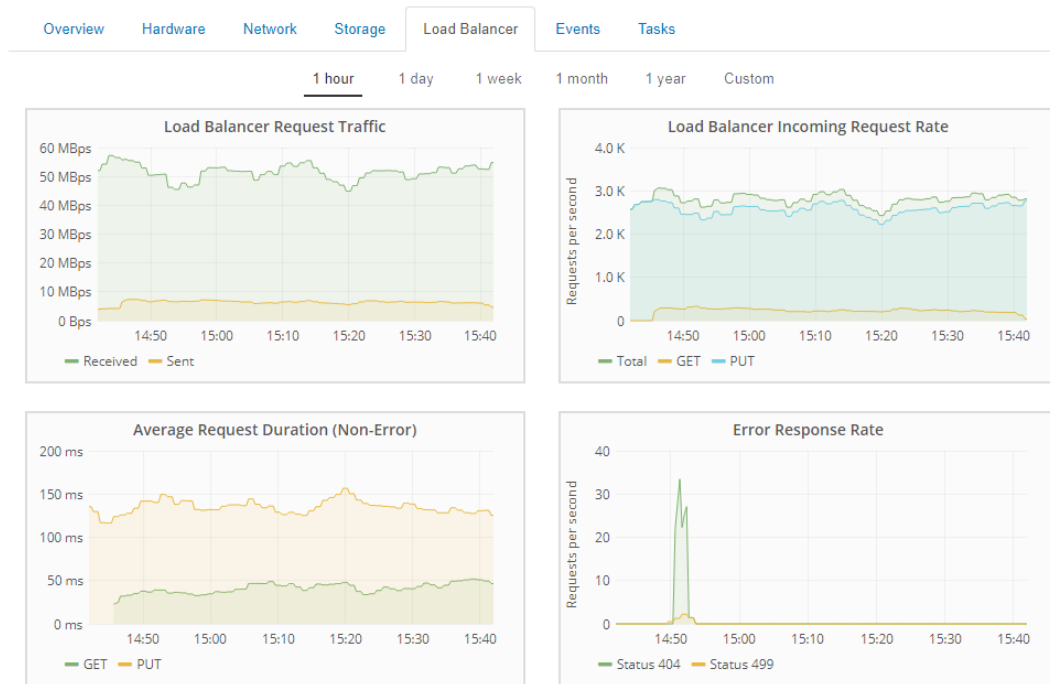
Viewing the Load Balancer tab

The Load Balancer tab includes performance and diagnostic graphs related to load balancer operations.

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 (or there is no load balancer configured), the graphs display “No data points.”

GW-SG1000-003-076 (Gateway 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.

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

Load Balancer Incoming Request Rate

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

Average Request Duration (Non-Error)

This graph provides a 3-minute moving average of request durations, broken down by request type (GET, PUT, HEAD, and DELETE). Each request duration starts when a request header is parsed by the load balancer 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

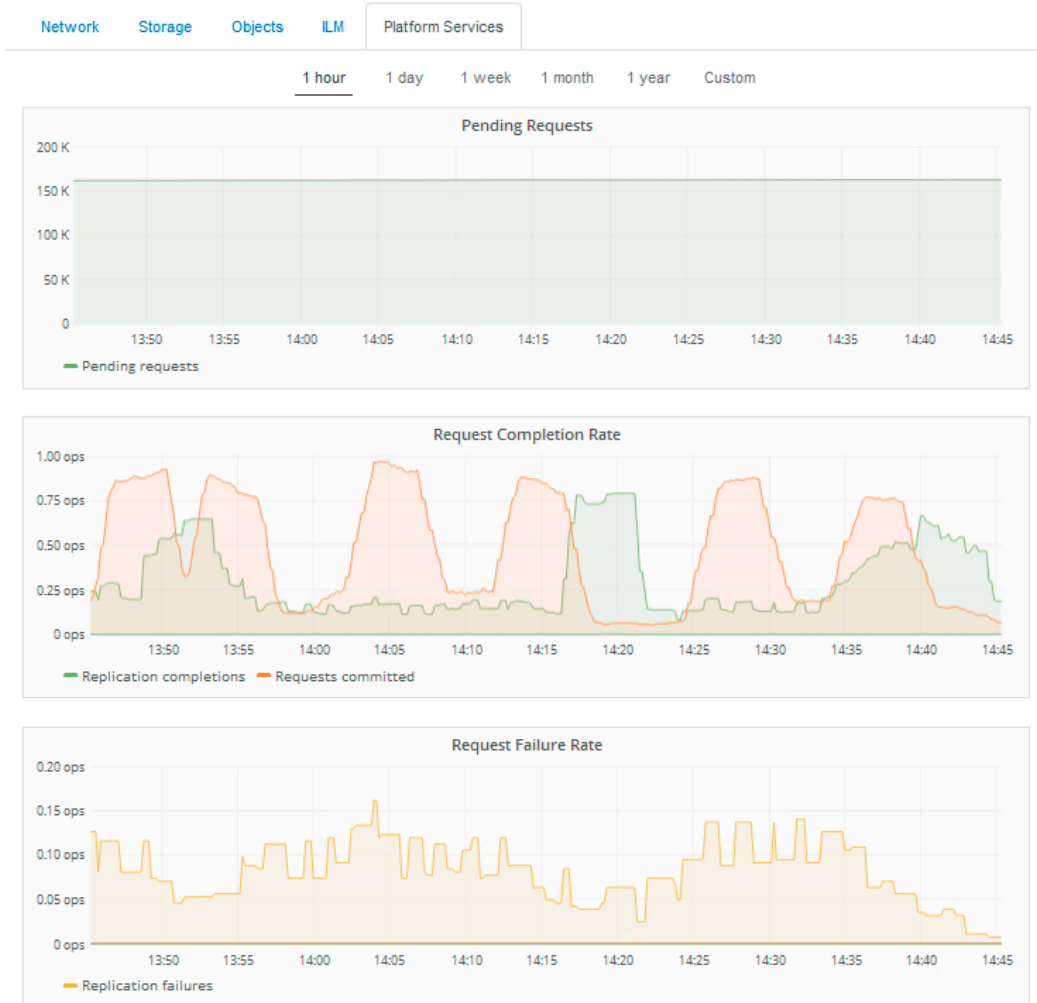
[Administering StorageGRID](#)

Viewing the Platform Services tab

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

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

Data Center 1



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

Related information

[Administering StorageGRID](#)

Viewing information about appliance Storage Nodes

The Nodes page lists information about service health and all computational, disk device, and network resources for each appliance Storage Node. You can also see memory, storage hardware, network resources, network interfaces, network addresses, and receipt and transmittal information.

Steps

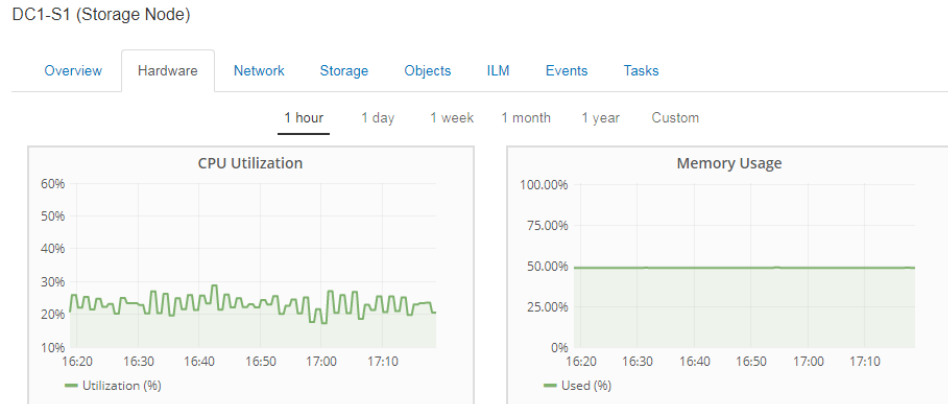
1. From the Nodes page, select an appliance Storage Node.
2. Select **Overview**.

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* is the Grid Network, Admin Network, or Client Network
- *hic* is a bonded port
- *mtc* are the management IP addresses on the appliance

Node Information ?	
ID	4e89eef5-fe86-4a95-949b-b4a29d9612de
Name	S1-104-50
Type	Storage Node
Software Version	11.3.0 (build 20190527.2302.d4b9c97)
IP Addresses	10.96.104.50, 10.96.97.40, 169.254.0.1 Show less ^
Interface	IP Address
eth0	10.96.104.50
eth0	fe80::2a0:98ff:fe6a:54b1
eth1	10.96.97.40
eth1	fe80::280:e5ff:fe43:19a4
hic2	10.96.104.50
hic4	10.96.104.50
mtc1	10.96.97.40
mtc2	169.254.0.1












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




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

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

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

StorageGRID Appliance		
Appliance Model	SG6060	
Storage Controller Name	StorageGRID-Q18-SGA-LAB51	
Storage Controller A Management IP	10.224.3.64	
Storage Controller B Management IP	10.224.3.65	
Storage Controller WWID	600a098000fc9ad2000000005cc9a9c4	
Storage Appliance Chassis Serial Number	721911500152	
Storage Hardware	Nominal	
Storage Controller Failed Drive Count	0	
Storage Controller A	Nominal	
Storage Controller B	Nominal	
Storage Controller Power Supply A	Nominal	
Storage Controller Power Supply B	Nominal	
Storage Data Drive Type	NL-SAS HDD	
Storage Data Drive Size	4.00 TB	
Storage RAID Mode	DDP	
Storage Multipath Connectivity	Nominal	
Overall Power Supply	Nominal	
Compute Controller BMC IP	10.224.3.31	
Compute Controller Serial Number	721911500171	
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
721911500152	99	Nominal 	N/A	Nominal	Nominal	Nominal

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.

Field in the Appliance table	Description
Storage Controller B Management IP	IP address for management port 1 on storage controller B. You use this IP to access SANtricity software to troubleshoot storage issues. Some appliance models do not have a storage controller B.
Storage Controller WWID	The worldwide identifier of the storage controller shown in SANtricity software.
Storage Appliance Chassis Serial Number	The chassis serial number of the appliance.
Storage Hardware	The overall status of the storage controller hardware. If the Storage Node is a StorageGRID appliance and it needs attention, then both the StorageGRID and SANtricity systems indicate that the storage hardware needs attention. If the status is “needs attention,” first check the storage controller using SANtricity software. Then, ensure that no other alarms exist that apply to the compute controller.
Storage Controller Failed Drive Count	The number of drives that are not optimal.
Storage Controller A	The status of storage controller A.
Storage Controller B	The status of storage controller B. Some appliance models do not have a storage controller B.
Storage Controller Power Supply A	The status of power supply A for the storage controller.
Storage Controller Power Supply B	The status of power supply B for the storage controller.
Storage Data Drive Type	The type of drives in the appliance, such as HDD (hard disk drive) or SSD (solid state drive).
Storage Data Drive Size	The total capacity including all data drives in the appliance.
Storage RAID Mode	The RAID mode configured for the appliance.
Storage Multipath Connectivity	The multipath connectivity state. For details about resolving performance or fault tolerance issues, refer to the E-Series documents.
Overall Power Supply	The status of all power supplies for the appliance.
Compute Controller BMC IP	The IP address of the baseboard management controller (BMC) port in the compute controller. You use this IP to connect to the BMC interface to monitor and diagnose the appliance hardware. This field is not displayed for appliance models that do not contain a BMC.
Compute Controller Serial Number	The serial number of the compute controller.

Field in the Appliance table	Description
Compute Hardware	The status of the compute controller hardware. This field is not displayed for appliance models that do not have separate compute hardware and storage hardware.
Compute Controller CPU Temperature	The temperature status of the compute controller's CPU.
Compute Controller Chassis Temperature	The temperature status of the compute controller.

Column in the Storage Shelves table	Description
Shelf Chassis Serial Number	The serial number for the storage shelf chassis.
Shelf ID	The numeric identifier for the storage shelf.
Shelf Status	The overall status of the storage shelf.
IOM Status	The status of the input/output modules (IOMs) for the storage shelf. If no expansion shelves are present, the status is N/A.
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. If the appliance does not contain drawers, the status is N/A.
Fan Status	The overall status of the cooling fans in the storage shelf.

- c. Confirm that all statuses are “Nominal.”

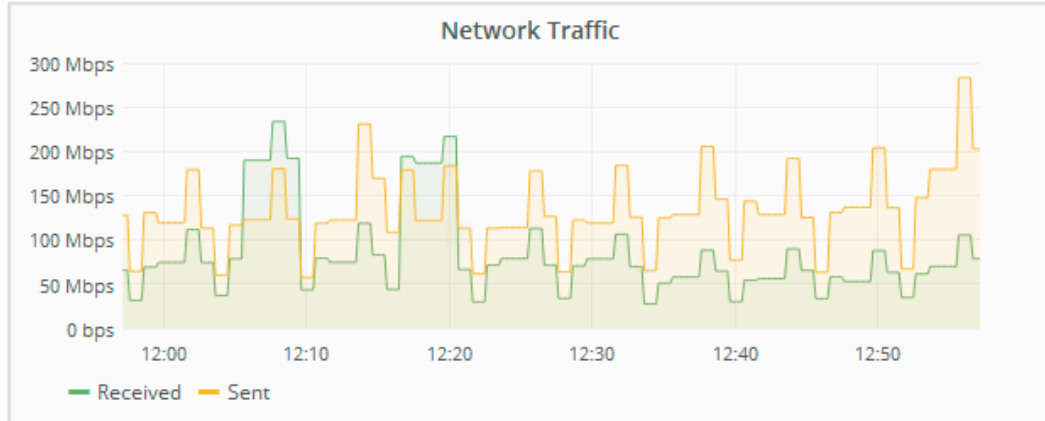
The statuses in this section correspond to the following alarm codes.

Field	Alarm code
Storage Controller Failed Drive Count	BADD
Compute Controller Chassis Temperature	BRDT
Compute Controller Hardware	CCNA
Compute Controller Power Supply A	CPSA
Compute Controller Power Supply B	CPSB
Compute Controller CPU Temperature	CPUT
Overall Power Supply	OPST
Storage Controller Power Supply A	PSAS
Storage Controller Power Supply B	PSBS
Storage Controller A	SCSA
Storage Controller B	SCSB
Storage Hardware	SOSS

For details about alarms in StorageGRID, see the information about troubleshooting.

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

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



- a. Review the **Network Interfaces** section.

Network Interfaces					
Name	Hardware Address	Speed	Duplex	Auto Negotiate	Link Status
eth0	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
eth1	D8:C4:97:2A:E4:9E	Gigabit	Full	Off	Up
eth2	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
hic1	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic2	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic3	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic4	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
mtc1	D8:C4:97:2A:E4:9E	Gigabit	Full	On	Up
mtc2	D8:C4:97:2A:E4:9F	Gigabit	Full	On	Up

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

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

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

Link mode	Bond mode	Individual HIC link speed (hic1, hic2, hic3, hic4)	Expected Grid/Client Network speed (eth0,eth2)
Fixed	Active/Backup	10	10

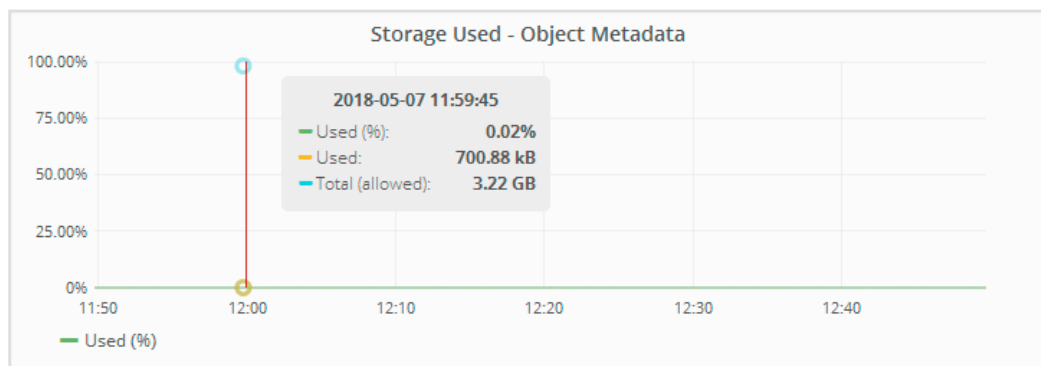
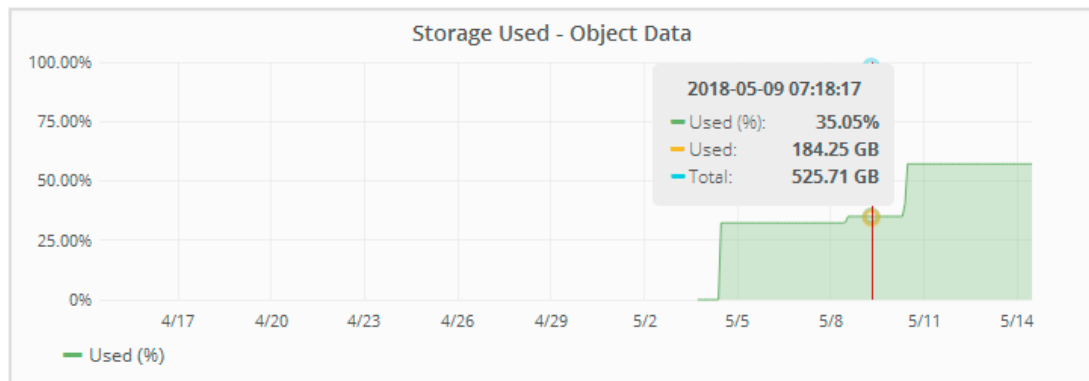
See the installation and maintenance instructions for your appliance for more information about configuring the 10/25-GbE ports.

- b. Review the **Network Communication** section.

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




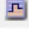
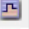










5. Select **Storage** to view graphs that show the percentages of storage used over time for object data and object metadata, as well as information about disk devices, volumes, and object stores.



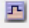
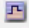

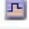












- a. Scroll down to view the amounts of available storage for each volume and object store.

The Worldwide Name for each disk matches the volume world-wide identifier (WWID) that appears when you view standard volume properties in SANtricity software (the management software connected to the appliance's storage controller).

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

Disk Devices						
Name	World Wide Name	I/O Load	Read Rate	Write Rate		
croot(8:1,sda1)	N/A	0.03%	 0 bytes/s	 4 KB/s		
cvloc(8:2,sda2)	N/A	0.37%	 0 bytes/s	 29 KB/s		
sdc(8:16,sdb)	N/A	0.00%	 0 bytes/s	 0 bytes/s		
sdd(8:32,sdc)	N/A	0.00%	 0 bytes/s	 183 bytes/s		
sde(8:48,sdd)	N/A	0.00%	 0 bytes/s	 12 bytes/s		

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

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

Related information

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)

Viewing information about appliance Admin Nodes and Gateway Nodes

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

Steps

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

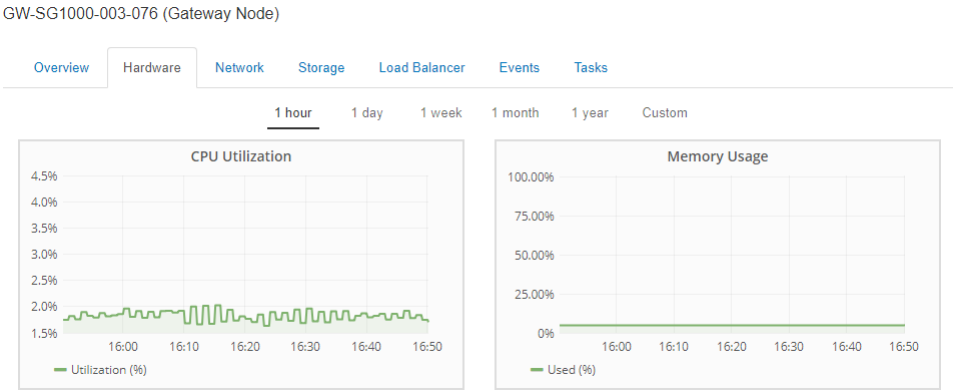
The Node Information table on the Overview tab displays the node's ID and name, the node type, the software version installed, and the IP addresses associated with the node. The Interface column contains the name of the interface, as follows:

- *adllb* and *adlli* are shown if active/backup bonding is used for the Admin Network interface
- *eth* is the Grid Network, Admin Network, or Client Network







- *hic* is a bonded port
- *mtc* are the management IP addresses on the appliance

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

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



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

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

Field in the Appliance table	Description
Appliance Model	The model number for this StorageGRID appliance.
Storage Controller Failed Drive Count	The number of drives that are not optimal.
Storage Data Drive Type	The type of drives in the appliance, such as HDD (hard disk drive) or SSD (solid state drive).
Storage Data Drive Size	The total capacity including all data drives in the appliance.
Storage RAID Mode	The RAID mode for the appliance.
Overall Power Supply	The status of all power supplies in the appliance.
Compute Controller BMC IP	The IP address of the baseboard management controller (BMC) port in the compute controller. You can use this IP to connect to the BMC interface to monitor and diagnose the appliance hardware.
Compute Controller Serial Number	The serial number of the compute controller.

Field in the Appliance table	Description
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.

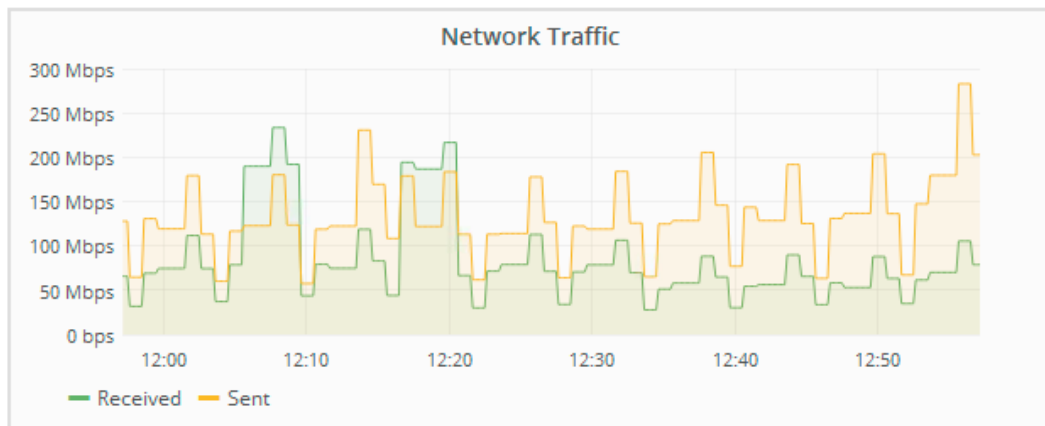
- c. Confirm that all statuses are “Nominal.”

The statuses in this section correspond to the following alarm codes.

Field	Alarm code
Storage Multipath Connectivity	APMS
Storage Controller Failed Drive Count	BADD
Compute Controller Chassis Temperature	BRDT
Compute Controller Hardware	CCNA
Compute Controller Power Supply A	CPSA
Compute Controller Power Supply B	CPSB
Compute Controller CPU Temperature	CPUT
Overall Power Supply	OPST

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
adllb	C2:20:17:59:1C:F3	10 Gigabit	Full	Off	Up
adlli	42:8F:84:80:00:A9	10 Gigabit	Full	Off	Up
eth0	98:03:9B:98:A2:72	400 Gigabit	Full	Off	Up
eth1	B4:A9:FC:08:4E:49	10 Gigabit	Full	Off	Up
eth2	98:03:9B:98:A2:72	400 Gigabit	Full	Off	Up
hic1	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
hic2	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
hic3	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
hic4	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
mtc1	B4:A9:FC:08:4E:49	Gigabit	Full	On	Up
mtc2	B4:A9:FC:08:4E:49	Gigabit	Full	On	Up








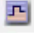


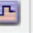
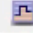




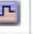





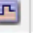
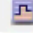

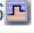




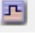

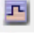
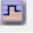

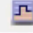






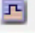
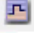
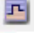

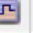
















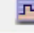











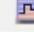









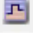

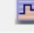
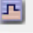
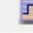







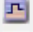

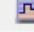

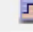


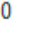



Use the following table with the values in the **Speed** column in the Network Interfaces table to determine whether the four 40/100-GbE network ports on the appliance were configured to use active/backup mode or LACP mode.

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

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

- b. Review the **Network Communication** section.

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

Network Communication									
Receive									
Interface	Data	Packets	Errors	Dropped	Frame Overruns	Frames			
eth0	3.250 TB 	5,610,578,144 	0 	8,327 	0 	0 			
eth1	1.205 GB 	9,828,095 	0 	32,049 	0 	0 			
eth2	849.829 GB 	186,349,407 	0 	10,269 	0 	0 			
hic1	114.864 GB 	303,443,393 	0 	0 	0 	0 			
hic2	2.315 TB 	5,351,180,956 	0 	305 	0 	0 			
hic3	1.690 TB 	1,793,580,230 	0 	0 	0 	0 			
hic4	194.283 GB 	331,640,075 	0 	0 	0 	0 			
mtc1	1.205 GB 	9,828,096 	0 	0 	0 	0 			
mtc2	1.168 GB 	9,564,173 	0 	32,050 	0 	0 			
Transmit									
Interface	Data	Packets	Errors	Dropped	Collisions	Carrier			
eth0	5.759 TB 	5,789,638,626 	0 	0 	0 	0 			
eth1	4.563 MB 	41,520 	0 	0 	0 	0 			
eth2	855.404 GB 	139,975,194 	0 	0 	0 	0 			
hic1	289.248 GB 	326,321,151 	5 	0 	0 	5 			
hic2	1.636 TB 	2,640,416,419 	18 	0 	0 	18 			
hic3	3.219 TB 	4,571,516,003 	33 	0 	0 	33 			
hic4	1.687 TB 	1,658,180,262 	22 	0 	0 	22 			
mtc1	4.563 MB 	41,520 	0 	0 	0 	0 			
mtc2	49.678 KB 	609 	0 	0 	0 	0 			

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

GW-SG1000-003-076 (Gateway Node)

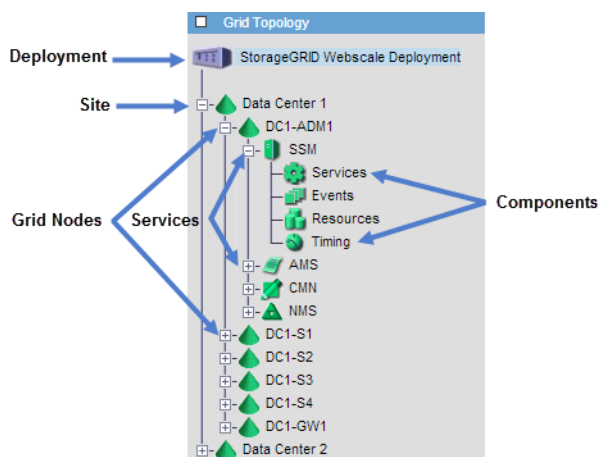
Overview	Hardware	Network	Storage	Load Balancer	Events	Tasks
Disk Devices						
Name	World Wide Name	I/O Load	Read Rate	Write Rate		
croot(253:2,dm-2)	N/A	0.00%	0 bytes/s	8 KB/s		
cvloc(253:3,dm-3)	N/A	0.01%	0 bytes/s	405 KB/s		
Volumes						
Mount Point	Device	Status	Size	Available	Write Cache Status	
/	croot	Online	21.00 GB	13.09 GB	Unknown	
/var/local	cvloc	Online	903.78 GB	894.55 GB	Unknown	

Related information[*SG1000 appliance installation and maintenance*](#)

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 > 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.

Information you should monitor regularly

You should monitor the StorageGRID system's key attributes and metrics regularly. Becoming familiar with system operations can help you spot trends before they turn into problems.

The following table lists the tasks you should perform on a regular basis.

Task	Frequency
Monitor system status. Note what has changed from the previous day.	Daily
Monitor the rate at which Storage Node capacity is being consumed.	Weekly
Check the capacity of the external archival storage system.	Weekly

The important items to monitor relate to:

- Object storage capacity
- Metadata storage capacity
- Object management related activities

When monitoring capacity, look at the absolute value and at the rate at which capacity is being consumed over time. Consumption rates can help you estimate when additional capacity might be required.



Note: The attributes and metrics used for storage capacity and metadata storage capacity do not include the capacity used for archived content.






Steps

1. [Viewing node icons](#) on page 36
2. [Viewing and acknowledging alarms](#) on page 37
3. [Monitoring storage capacity](#) on page 39
4. [Monitoring the recovery point objective through ILM](#) on page 46
5. [Monitoring object verification operations](#) on page 48
6. [Monitoring the SSM service](#) on page 49
7. [Monitoring the Total Events \(SMTT\) alarm](#) on page 50
8. [Monitoring archival capacity](#) on page 51

Viewing node icons

The Grid Manager uses the following icons to show the health of each grid node. The icons indicate whether nodes are connected to the grid and show if any alarms are active on the node.

Icon	Color	Node state	Alarm severity	Meaning
	Green	Connected	Normal	The node is functioning normally. It is connected to the grid and there are no alarms.
	Yellow	Connected	Notice	The node is connected to the grid, but an unusual condition exists that does not affect normal operations.

Icon	Color	Node state	Alarm severity	Meaning
	Light Orange	Connected	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	Connected	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	Connected	Critical	The node is connected to the grid, but an abnormal condition exists that has stopped normal operations. You should address the issue immediately.
	Gray	Disconnected	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.
	Blue	Disconnected	Unknown	<p>The node is not connected to the grid. This situation requires immediate attention. For example, the network connection between nodes has been lost or the power is down. This is the most severe condition.</p> <p>Note: You might see transient blue nodes during managed shutdown operations. You can ignore these alarms.</p>

Viewing and acknowledging alarms

When an alarm is triggered, an icon is shown on the Alarms menu, on the Nodes page for the affected node and site, and in the Grid Topology tree. You can view any currently active alarms from the Nodes page and acknowledge them as required.

Before you begin

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


About this task

When you acknowledge an alarm, the alarm no longer appears on the Nodes page, and the corresponding node icons no longer appear. The alarm is shown again only if it is triggered at the next severity level, or if it is resolved and occurs again.

Steps

1. Select **Nodes**.
2. Select a grid node that has an alarm icon.
3. On the Overview tab, review the information for the alarm in the Alarms table.

This table lists any unacknowledged alarms for the node.

Alarms ?							
Severity	Attribute Name	Code	Service	Description	Time Triggered	Trigger Value	Current Value
 Major	Outbound Replication Status ?	ORSU	ARC	Storage Unavailable	2019-05-24 21:42:02 MDT	Storage Unavailable	Storage Unavailable

4. To learn more about the alarm and its trigger values, click the help icon ? next to the attribute name in the table.

Help

Outbound Replication Status (ORSU)

The current status of the Replication component for outbound replication.

Attribute values are:

- 0 = No Errors
- 10 = Storage Unavailable
- 20 = Disabled

Outbound Replication Status ?

ORSU

ARC

5. To view additional details about an alarm or to acknowledge it, click the service name in the table.

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

Overview


Alarms


Reports


Configuration

Main

History


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

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

Apply Changes 

6. To acknowledge the alarm, select the **Acknowledge** check box for the alarm, and click **Apply Changes**.

The alarm no longer appears on the Nodes page and the corresponding node icons are not shown.

Note: When you acknowledge an alarm, the acknowledgment is not copied to other Admin Nodes. For this reason, if you view the Nodes page from another Admin Node, you might continue to see alarm icons.

7. As required, view acknowledged alarms on the **Alarms > Current Alarms** page.
 - a. Select **Alarms**. Then, in the Alarms section of the menu, select **Current Alarms**.
 - b. Select **Show Acknowledged Alarms**.

Related concepts

[Viewing node icons](#) on page 36

[Managing alarms](#) on page 60

Related references

[Alarms reference](#) on page 166

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.

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

Steps

1. [Monitoring storage capacity for the entire grid](#) on page 39
2. [Monitoring storage capacity for each Storage Node](#) on page 41
3. [Monitoring object metadata capacity for each Storage Node](#) on page 44

Monitoring storage capacity for the entire grid

The Dashboard shows the available storage capacity for the entire grid and for individual data centers. You can use the Available Storage charts to quickly assess storage use for an entire StorageGRID system. In a multi-site grid, you can compare storage usage between sites (data centers).

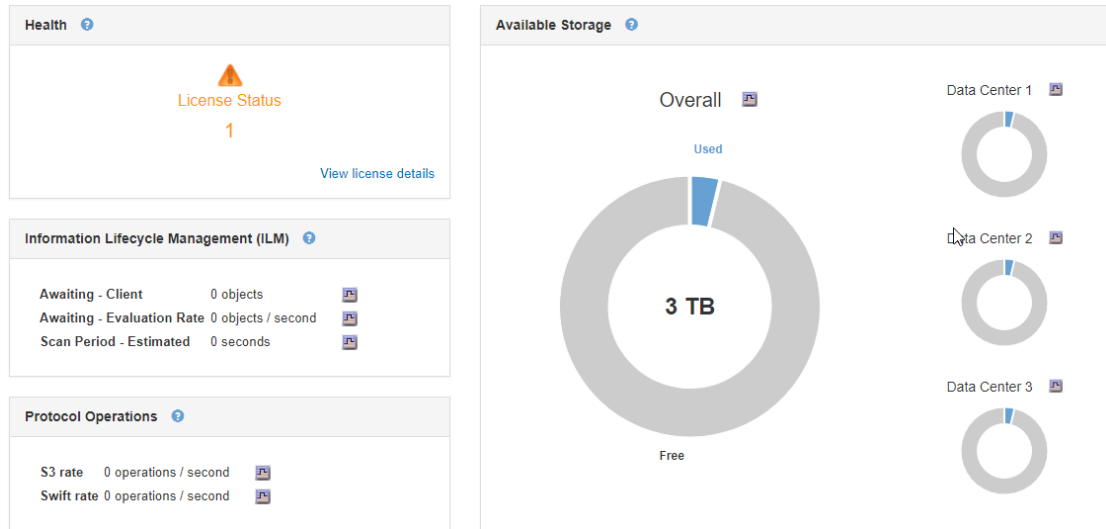
Before you begin

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

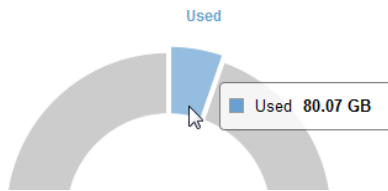
Steps


1. Select **Dashboard**.
2. In the **Available Storage** panel, note the overall summary of free and used storage capacity. For multi-site grids, review the chart for each data center.

Note: The summary does not include archival media.



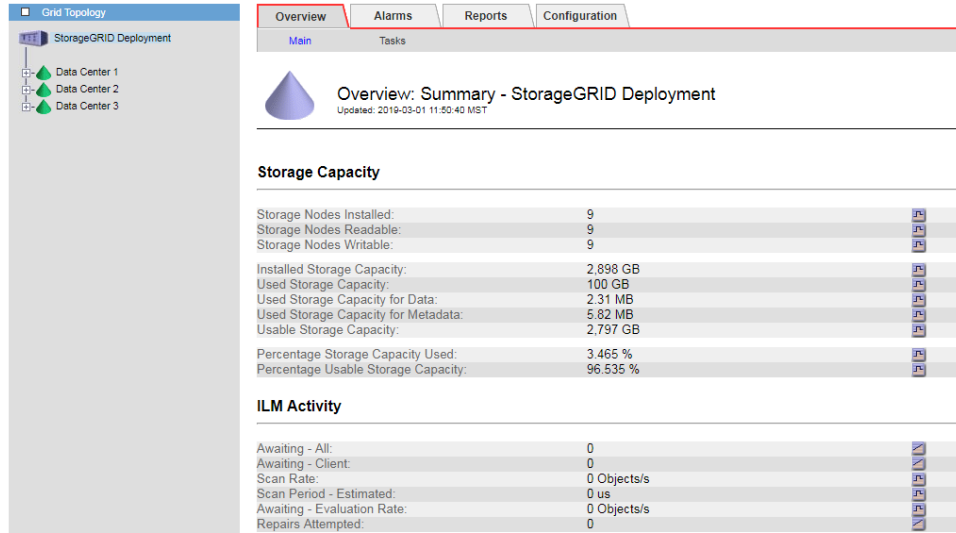
3. Place your cursor over the chart's Free or Used capacity sections to see exactly how much space is free or used.



4. Click the Chart (Reports) 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.

5. To monitor additional storage attributes:
 - Go to **Nodes > Storage Node > Storage**, and view the graphs and tables.
 - As directed by technical support, select **Support > Grid Topology**. Then select **StorageGRID deployment > Overview > Main**.



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

Related information

[Expanding a StorageGRID system](#)

Monitoring storage capacity for each Storage Node

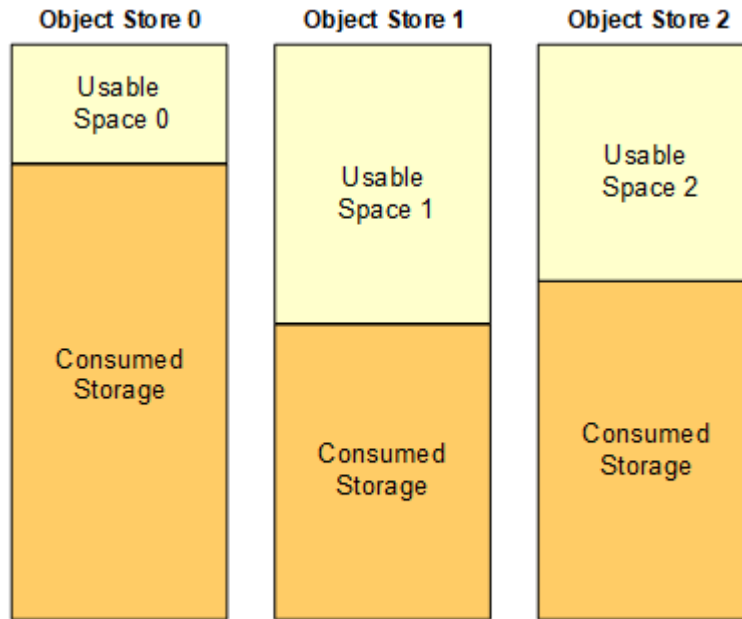
You must monitor the value of the Total Usable Space (STAS) attribute for each Storage Node to ensure that the node has enough space for new object data.

Before you begin

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

About this task

STAS is calculated by adding together the available space on all object stores within the Storage Node.



$$\text{Total Usable Space} = \text{Usable Space 0} + \text{Usable Space 1} + \text{Usable Space 2}$$

Note: If the value of STAS for a Storage Node falls below the value of the Storage Volume Hard Read-Only Watermark, the Storage Node becomes read-only and can no longer store new objects. You should add storage volumes or Storage Nodes before this happens.

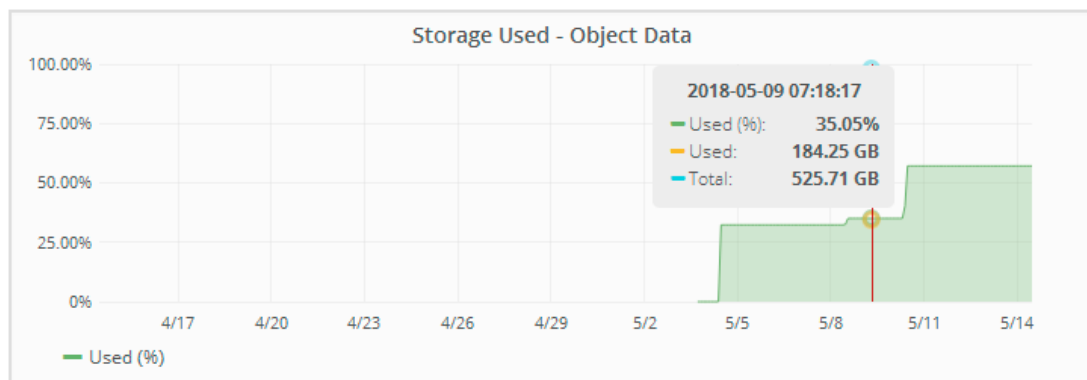
Steps


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
















The graphs and tables for the node appear.











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







A pop-up displays Used (%), Used, and Total capacities. The Total value is the Total Usable Space (STAT) attribute.



3. Review the values in the tables below the graphs. To view graphs of the values, click the Reports (Charts) icon  in the **Available** columns in the **Volumes** and **Object Stores** tables.

Disk Devices							
Name	World Wide Name	I/O Load		Read Rate		Write Rate	
croot(8:1,sda1)	N/A	0.03%		0 bytes/s		4 KB/s	
cvloc(8:2,sda2)	N/A	0.37%		0 bytes/s		29 KB/s	
sdc(8:16,sdb)	N/A	0.00%		0 bytes/s		0 bytes/s	
sdd(8:32,sdc)	N/A	0.00%		0 bytes/s		183 bytes/s	
sde(8:48,sdd)	N/A	0.00%		0 bytes/s		12 bytes/s	

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

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

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

Usable space is the actual amount of storage space available to store objects.

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

The STSS alarm and the **Low object storage** alert are triggered when insufficient space remains for storing object data on a Storage Node.

Related tasks

[Troubleshooting Storage Status \(STS\) alarms](#) on page 144

[Troubleshooting Low object data storage alerts](#) on page 148

Related information

[Administering StorageGRID](#)

[Expanding a StorageGRID system](#)

Monitoring object metadata capacity for each Storage Node

You must monitor the metadata usage for each Storage Node to ensure that adequate space remains available for essential database operations. If the CDLP alarm and the **Low metadata storage** alert are triggered, you must add new Storage Nodes.

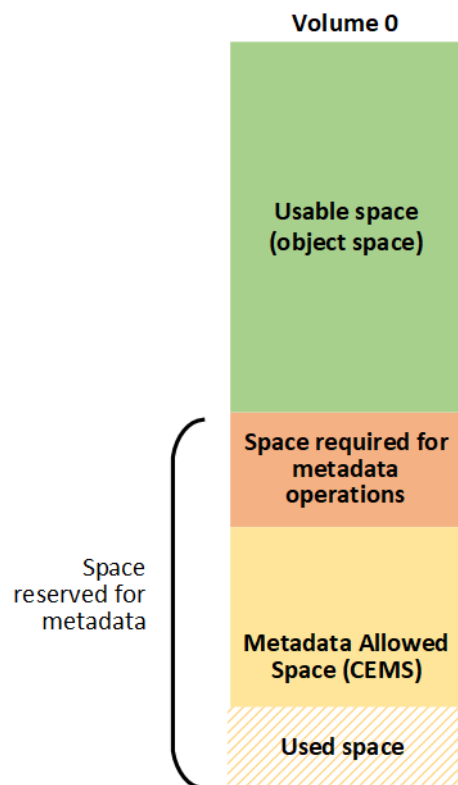
Before you begin

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

About this task

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

The total space reserved on each Storage Node for metadata is known as the Metadata Reserved Space (CAWM). The Metadata Reserved Space is subdivided into the space available for object metadata (the Metadata Allowed Space, or CEMS) and the space required for essential database operations, such as compaction and repair.



If the space used by object metadata is more than 100% of the Metadata Allowed Space, database operations cannot run efficiently and errors will occur.

The Metadata Used Space (Percent) attribute, or CDLP, measures how full the Metadata Allowed Space is. When the Metadata Used Space (Percent) reaches the following thresholds, the CDLP alarm and the **Low metadata storage** alert are triggered:

- Minor:** Object metadata is using 70% or more of the Metadata Allowed Space. You should add new Storage Nodes as soon as possible.

- **Major:** Object metadata is using 90% or more of the Metadata Allowed Space. You must add new Storage Nodes immediately.
Attention: When the CDLP alarm is triggered at the major level, a warning appears on the Dashboard. If this warning appears, you must add new Storage Nodes immediately. You must never allow object metadata to use more than 100% of the allowed space.
- **Critical:** Object metadata is using 100% or more of the Metadata Allowed Space and is starting to consume the space required for essential database operations. You must stop the ingest of new objects, and you must add new Storage Nodes immediately.

When you add the new nodes, the system automatically rebalances object metadata across all Storage Nodes, and the CDLP alarm and the **Low metadata storage** alert are cleared.

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

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

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

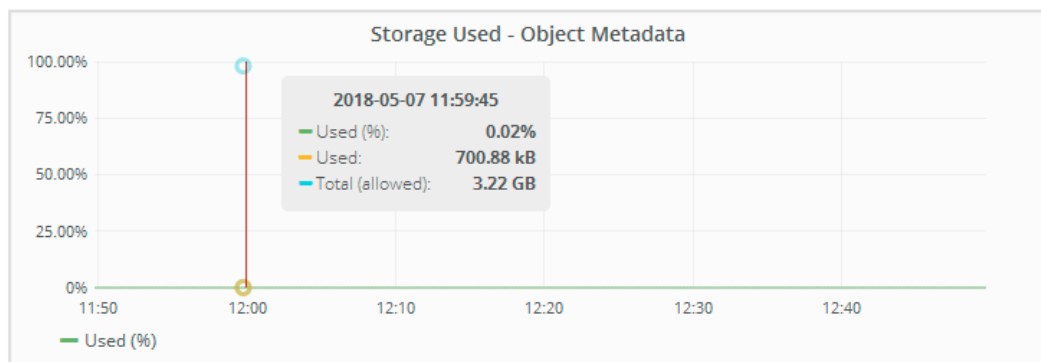
Undesirable results can occur if object metadata uses more than 100% of the allowed space. You must add new Storage Nodes immediately or contact support.

Attention: If the first storage volume is smaller than the Metadata Reserved Space (for example, in a non-production environment), the calculation for the CDLP alarm and the **Low metadata storage** alert might be inaccurate.

Steps

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

The value for Used % is the Metadata Used Space (Percent) (CDLP) attribute.



3. If the **Used %** value is 70% or higher, expand the StorageGRID system by adding Storage Nodes.
4. To view details about the CDLP alarm:
 - a. Select **Alarms**. Then, in the Alarms section of the menu, select **Current Alarms**.

- b. Review the details for the alarm in the table.


CDLP (Metadata Used Space (Percent)) alarms of all severities (minor, major, and critical) are displayed on this page.

Example

Current Alarms

Last Refreshed: 2018-05-04 10:56:02 MDT

☐ Show Acknowledged Alarms (1 - 18 of 18)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value
 Minor	CDLP (Metadata Used Space (Percent))	Data Center 1/SGA-Lab11/DDS	The metadata store is more than 70% full. You should add new Storage Nodes as soon as possible.	2018-05-04 10:53:51 MDT	89.319 %	89.319 %

Note: You can also see the CDLP alarm for a Storage Node by selecting **Support > Grid Topology > Storage Node > DDS > Data Store > Alarms**.

5. To view details about the **Low metadata storage** alert:
 - a. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **Alerts**.
 - b. From the table of alerts, expand the **Low metadata storage** alert group, if required, and select the specific alert you want to view.
 - c. Review the details in the alert dialog box.
6. If a major or critical CDLP alarm or **Low metadata storage** alert has been triggered, perform an expansion to add Storage Nodes immediately.

When you add the new nodes, the system automatically rebalances object metadata across all Storage Nodes, and the alarms clear.

Related information

[Administering StorageGRID](#)

[Expanding a StorageGRID system](#)

Monitoring the recovery point objective through ILM


You can track ILM evaluation attributes to determine the recovery point objective (RPO) of the StorageGRID system as defined by the ILM policy. The RPO defines the maximum tolerable period in which data might be lost because of a site failure, a Storage Node failure, or both.

Before you begin

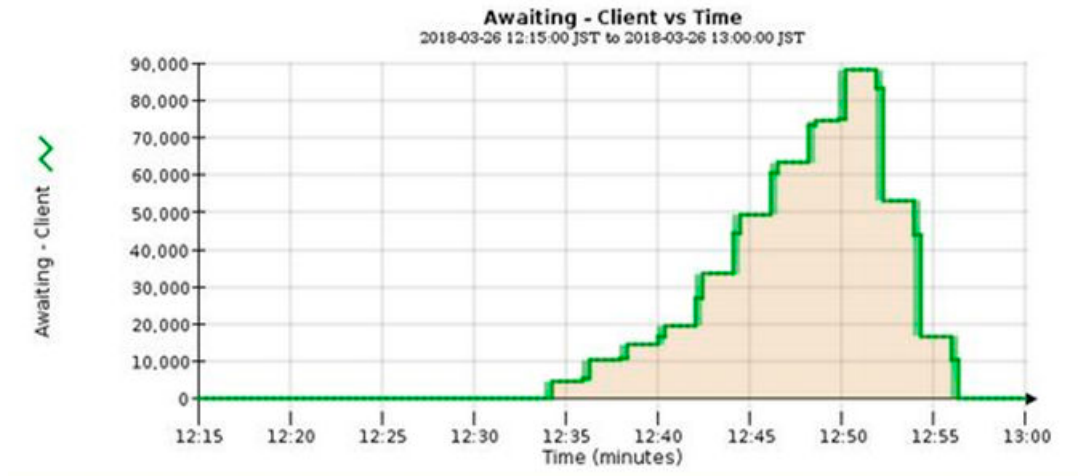
You must be signed in to the Grid Manager using a supported browser.

About this task

The StorageGRID system manages objects by applying the active ILM policy. The ILM policy and associated ILM rules determine how many copies are made, how those copies are made, the appropriate placement, and the length of time each copy is retained.

Ingest or other activity can exceed the rate at which the system can process ILM. When this scenario occurs, the system might begin to queue objects whose ILM can no longer be fulfilled in near real time. The chart of Awaiting - Client can be useful in determining whether this situation has occurred. You can find the chart in the Grid Manager by going to **Dashboard > Information Lifecycle Management (ILM) > Awaiting - Client** and clicking the  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.



You can further investigate ILM queues using the **Nodes** tab.

Steps

1. Select **Nodes**.
2. Select **deployment > 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:
 - **Scan Period - Estimated:** The estimated time to complete a full ILM scan of all objects.
Note: A full scan does not guarantee that ILM has been applied to all objects.
 - **Repairs Attempted:** The total number of object repair operations for replicated data that have been attempted. This count increments each time a Storage Node tries to repair a high-risk object. High-risk ILM repairs are prioritized if the grid becomes busy.
Note: The same object repair might increment again if replication failed after the repair.

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

Monitoring object verification operations

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

Before you begin

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

About this task

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

The foreground verification process 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.

You can look at the Nodes page for a Storage Node to review results from background and foreground verifications, such as corrupt or missing objects detected. 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	
Quarantined Objects	0	

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

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

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

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

Related tasks

[Running foreground verification](#) on page 130

Related information

[Administering StorageGRID](#)

Monitoring the SSM service

The Server Status Monitor (SSM) service is present on all grid nodes and monitors the node's status, services, and resources. You can look at the SSM entry for each grid node to see the status of services on that node.

About this task

The SSM service monitors the condition of the server and related hardware, polls the server and hardware drivers for information, and displays the processed data. The information monitored includes:

- Service status
- Computation resources (restarts, runtime, uptime, load)
- Memory information (installed, available)
- CPU information (type, mode, speed)
- Volumes (status, available space)
- Network (addresses, interfaces, resources)
- NTP synchronization

Steps

1. Select **Support > Grid Topology**.

2. Select the grid node, and select **SSM**.

The state and status of the node's SSM service is shown.

3. Select **SSM > Services**.

The status of each service on the node is shown, as in this example for a primary Admin Node.

Service	Version	Status	Threads	Load	Memory
ADE Exporter Service	11.3.0-20190628.1946.557d4e0	Running	12	0.006 %	15 MB
Account Service Forwarder	11.3.0-20190715.1915.432cb81	Running	3	0.078 %	37.4 MB
Audit Management System (AMS)	11.3.0-20190715.2307.4de3acc	Running	43	0.041 %	49.4 MB
CIFS Filesharing (nmbd)	2.4.5.16+dfsg-1+deb9u2	Not Running	0	0 %	0 B
CIFS Filesharing (smbd)	2.4.5.16+dfsg-1+deb9u2	Not Running	0	0 %	0 B
CIFS Filesharing (winbindd)	2.4.5.16+dfsg-1+deb9u2	Not Running	0	0 %	0 B
Configuration Management Node (CMN)	11.3.0-20190715.2307.4de3acc	Running	43	0.076 %	56.8 MB
Database Engine	10.1.38-0+deb9u1	Running	64	0.784 %	879 MB
Dynamic IP Service	11.3.0-20190713.0309.9efe958	Running	7	0.234 %	45.7 MB
Grid Deployment Utility Server	11.3.0-20190704.2314.4caae0d	Running	3	0 %	42.8 MB
High Availability Service	1:1.3.2-1	Not Running	0	0 %	0 B
Load Balancer Service	11.3.0-	Running	21	0.002 %	56.7 MB

- To monitor the events for the node, select **SSM > Events**.

Note: This information is the same as displayed on the **Nodes > Events** tab.

- To monitor the node's resources, select **SSM > Resources**.

As required, you can use the Configuration tab to reset network error counters.

- To monitor the clock settings for the node, select **SSM > Timing**.

The timing attributes report on the state of the grid node's local clock and the state of neighboring grid node clocks. In addition, these attributes report on NTP synchronization.


Monitoring the Total Events (SMTT) alarm



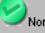
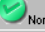

When the Total Events (SMTT) alarm is raised, monitor the situation and take appropriate action.

Category	Code	Service	Notes
Total events	SMTT	SSM	<p>The total number of logged error or fault events includes errors such as network errors. Unless these errors have been cleared (that is, the count has been reset to 0), total events alarms can be triggered.</p> <p>Note: This alarm is safe to ignore only if the events that triggered the alarm have been investigated.</p>

Overview Alarms Reports Configuration

Main History

 **Alarms: SSM (CSN1-A-1) - Events**
Updated: 2010-03-27 12:49:00 PDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
 Normal	SMST (Log Monitor State)						<input type="checkbox"/>
 Notice	SMTT (Total Events)	At least 1	2010-03-26 22:11:01 PDT	1	1		<input type="checkbox"/>
 Normal	AMQS (Audit Messages Queued)						<input type="checkbox"/>
 Normal	NRLY (Available Audit Relays)						<input type="checkbox"/>
 Normal	ABRL (Available Attribute Relays)						<input type="checkbox"/>

Related concepts

[Troubleshooting Total Events \(SMTT\) alarms](#) on page 143

Monitoring archival capacity

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

Before you begin

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

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.


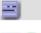






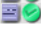

Steps

1. Select **Support > 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.

Overview Alarms Reports Configuration

Main

 **Overview: ARC (DC1-ARC1-98-165) - ARC**
Updated: 2015-09-15 15:59:21 PDT

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

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

Related information

[Administering StorageGRID](#)

Using reports

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

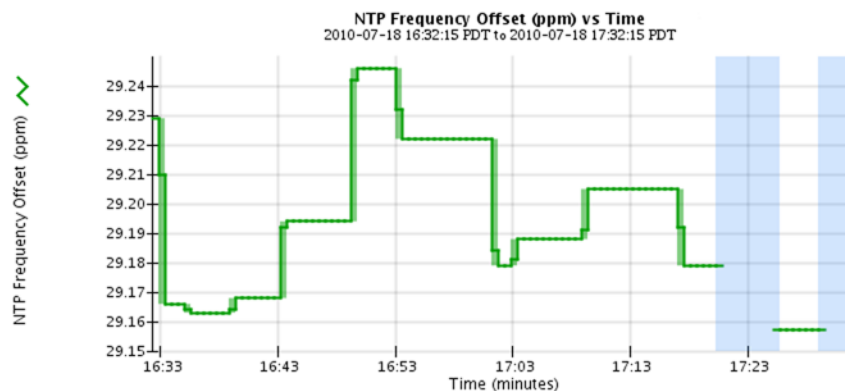
Types of charts

In addition to the summary pie charts shown on the Dashboard, you can access more detailed graphs that present the data with the attribute value (vertical axis) over a specified time span (horizontal axis).

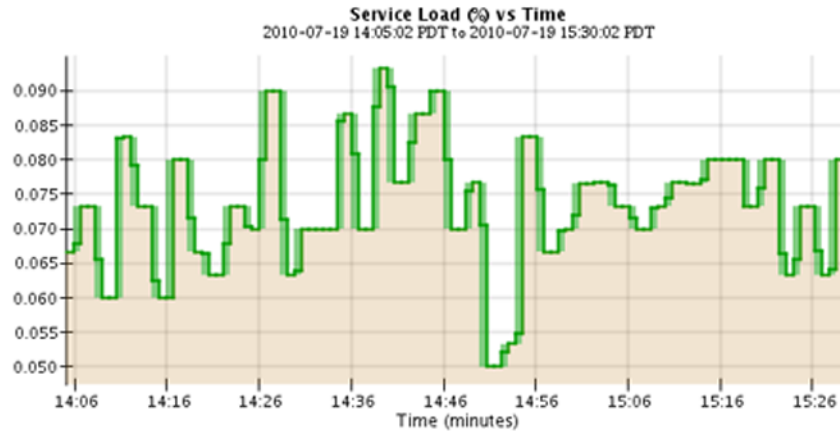
The Dashboard provides access to pie charts summarizing available storage as well as graphs for ILM and protocol operations.

In addition, graphs are available from the Nodes page and the Grid Topology tree. There are three types of graphs:

- Line graph: Used to plot the values of an attribute that has 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 graph: Used to plot volumetric quantities, such as object counts or service load values. Area graphs are similar to line graphs, but include a light brown shading below the line. The changes in the value are plotted in regular data intervals (bins) over time.



- **State graph:** State graphs are used to plot 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.

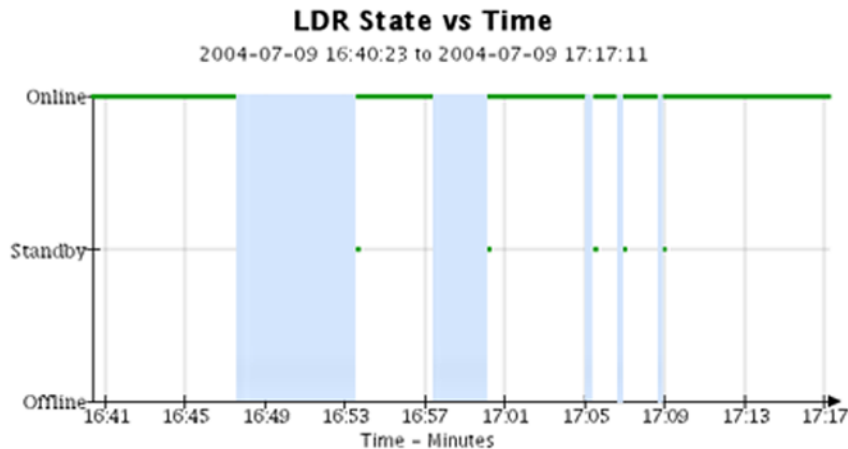
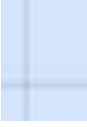




Chart legend

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

Sample	Meaning
	Reported attribute values are plotted using dark green lines.
	Light green shading around dark green lines indicates that the actual values in that time range vary and have been “binned” for faster plotting. The dark line represents the weighted average. The range in light green indicates the maximum and minimum values within the bin. Light brown shading is used for area graphs to indicate volumetric data.
	Blank areas (no data plotted) indicate that the attribute values were unavailable. The background can be blue, gray, or a mixture of gray and blue, depending on the state of the service reporting the attribute.

Sample	Meaning
	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

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

Before you begin

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

About this task

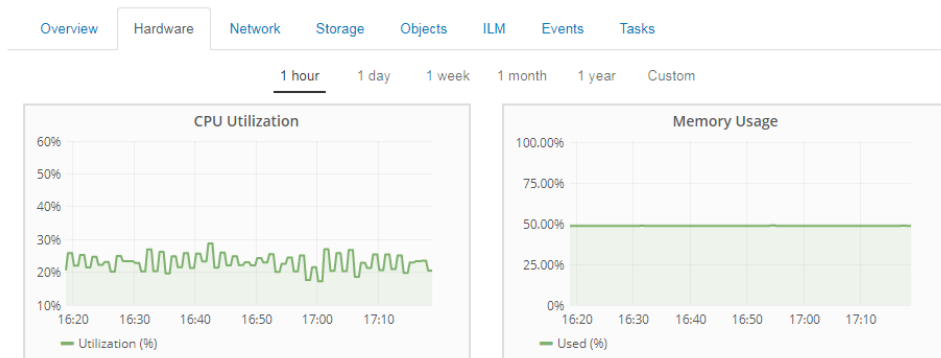
Note: You cannot create charts for all attributes; for example, text attributes such as Node ID, version number, and build number.

Steps

1. Select **Nodes** > *grid* or *grid node*.
2. Select the appropriate tab, or select **Chart**  to the right of an attribute to display a chart.

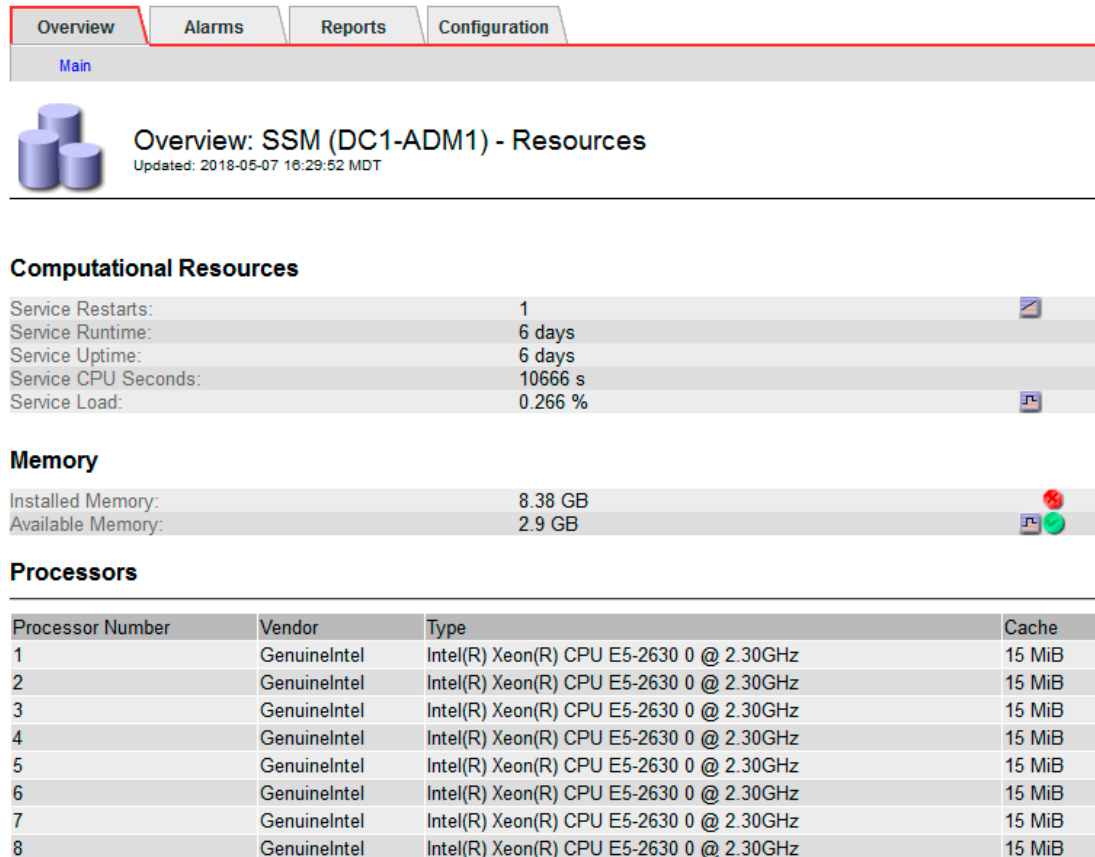
For some attributes, the chart appears when you select the tab. For other attributes, you select the **Chart** icon to display the chart.

DC1-S1 (Storage Node)



3. To display additional attributes and charts, select **Support** > **Grid Topology**.

4. Select *grid node > component or service > Overview > 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

5. Click **Chart** next to the attribute.

The display automatically changes to the **Reports > Charts** page. The chart displays the attribute's data over the past day.

Generating charts

Charts display a graphical representation of attribute data values. You can report on a data center site, grid node, component, or service.

Before you begin

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

Steps

- Select **Support > 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/DD HH:MM:SS* in local time. Leading zeros are required to match the format. For example, 2017/4/6 7:30:00 fails validation. The correct format is: 2017/04/06 07:30:00.

8. Click **Update**.

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

9. If you want to print the chart, right-click and select **Print**, and modify any necessary printer settings and click **Print**.

Types of text reports

Text reports display a textual representation of attribute data values that have been processed by the NMS service. There are two types of reports generated depending on the time period you are reporting on: raw text reports for periods less than a week, and aggregate text reports for time periods greater than a week.

Raw text reports

A raw text report displays details about the selected attribute:

- Time Received: Local date and time that a sample value of an attribute's data was processed by the NMS service.
- Sample Time: Local date and time that an attribute value was sampled or changed at the source.
- Value: Attribute value at sample time.

Text Results for Services: Load - System Logging

2010-07-18 15:58:39 PDT To 2010-07-19 15:58:39 PDT

Time Received	Sample Time	Value
2010-07-19 15:58:09	2010-07-19 15:58:09	0.016 %
2010-07-19 15:56:06	2010-07-19 15:56:06	0.024 %
2010-07-19 15:54:02	2010-07-19 15:54:02	0.033 %
2010-07-19 15:52:00	2010-07-19 15:52:00	0.016 %
2010-07-19 15:49:57	2010-07-19 15:49:57	0.008 %
2010-07-19 15:47:54	2010-07-19 15:47:54	0.024 %
2010-07-19 15:45:50	2010-07-19 15:45:50	0.016 %
2010-07-19 15:43:47	2010-07-19 15:43:47	0.024 %
2010-07-19 15:41:43	2010-07-19 15:41:43	0.032 %
2010-07-19 15:39:40	2010-07-19 15:39:40	0.024 %
2010-07-19 15:37:37	2010-07-19 15:37:37	0.008 %
2010-07-19 15:35:34	2010-07-19 15:35:34	0.016 %
2010-07-19 15:33:31	2010-07-19 15:33:31	0.024 %
2010-07-19 15:31:27	2010-07-19 15:31:27	0.032 %
2010-07-19 15:29:24	2010-07-19 15:29:24	0.032 %
2010-07-19 15:27:21	2010-07-19 15:27:21	0.049 %
2010-07-19 15:25:18	2010-07-19 15:25:18	0.024 %
2010-07-19 15:21:12	2010-07-19 15:21:12	0.016 %
2010-07-19 15:19:09	2010-07-19 15:19:09	0.008 %
2010-07-19 15:17:07	2010-07-19 15:17:07	0.016 %

Aggregate text reports

An aggregate text report displays data over a longer period of time (usually a week) than a raw text report. Each entry is the result of summarizing multiple attribute values (an aggregate of attribute values) by the NMS service over time into a single entry with average, maximum, and minimum values that are derived from the aggregation.

Each entry displays the following information:

- **Aggregate Time:** Last local date and time that the NMS service aggregated (collected) a set of changed attribute values.
- **Average Value:** The average of the attribute's value over the aggregated time period.
- **Minimum Value:** The minimum value over the aggregated time period.
- **Maximum Value:** The maximum value over the aggregated time period.

Text Results for Attribute Send to Relay Rate

2010-07-11 16:02:46 PDT To 2010-07-19 16:02:46 PDT

Aggregate Time	Average Value	Minimum Value	Maximum Value
2010-07-19 15:59:52	0.271072196 Messages/s	0.266649743 Messages/s	0.274983464 Messages/s
2010-07-19 15:53:52	0.275585378 Messages/s	0.266562352 Messages/s	0.283302736 Messages/s
2010-07-19 15:49:52	0.279315709 Messages/s	0.233318712 Messages/s	0.333313579 Messages/s
2010-07-19 15:43:52	0.28181323 Messages/s	0.241651024 Messages/s	0.374976601 Messages/s
2010-07-19 15:39:52	0.284233141 Messages/s	0.249982001 Messages/s	0.324971987 Messages/s
2010-07-19 15:33:52	0.325752083 Messages/s	0.266641993 Messages/s	0.358306197 Messages/s
2010-07-19 15:29:52	0.278531507 Messages/s	0.274984766 Messages/s	0.283320999 Messages/s
2010-07-19 15:23:52	0.281437642 Messages/s	0.274981961 Messages/s	0.291577735 Messages/s
2010-07-19 15:17:52	0.261563307 Messages/s	0.258318006 Messages/s	0.266655787 Messages/s
2010-07-19 15:13:52	0.265159147 Messages/s	0.258318557 Messages/s	0.26663986 Messages/s

Generating text reports

Text reports display a textual representation of attribute data values that have been processed by the NMS service. You can report on a data center site, grid node, component, or service.

Before you begin

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

About this task

For attribute data that is expected to be continuously changing, this attribute data is sampled by the NMS service (at the source) at regular intervals. For attribute data that changes infrequently (for example, data based on events such as state or status changes), an attribute value is sent to the NMS service when the value changes.

The type of report displayed depends on the configured time period. By default, aggregate text reports are generated for time periods longer than one week.

Grey 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 > 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), 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 report appears after a few moments. Allow several minutes for tabulation of long time ranges.
7. If you selected Custom Query, you need to customize the time period to report on by entering the **Start Date** and **End Date**.
Use the format *YYYY/MM/DD HH:MM:SS* in local time. Leading zeros are required to match the format. For example, 2017/4/6 7:30:00 fails validation. The correct format is: 2017/04/06 07:30:00.
8. Click **Update**.
A text report is generated after a few moments. Allow several minutes for tabulation of long time ranges. Depending on the length of time set for the query, either a raw text report or aggregate text report is displayed.
9. If you want to print the report, right-click and select **Print**, and modify any necessary printer settings and click **Print**.

Exporting text reports

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

About this task


The copied data can then be saved into a new document (for example, a spreadsheet) and used to analyze the performance of the StorageGRID system.

Steps

1. Select **Support > Grid Topology**.
2. Create a text report.
3. Click **Export** .

Overview Alarms **Reports** Configuration


Charts **Text**

 **Reports (Text): SSM (170-176) - Events**

Attribute: Results Per Page: Start Date: End Date:

Quick Query: Raw Data: ☒

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 

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

Previous « 1 2 3 4 5 » Next

The Export Text Report window opens displaying the report.

```
Grid ID: 000000
OID: 2.16.124.113590.2.1.400019.1.1.1.1.16996732.200
Node Path: Site/170-176/SSM/Events
Attribute: Attribute Send to Relay Rate (ABSR)
Query Start Date: 2010-07-19 08:42:09 PDT
Query End Date: 2010-07-20 08:42:09 PDT
Time Received,Time Received (Epoch),Sample Time,Sample Time (Epoch),Value,Type
2010-07-20 08:40:46,1279640446559000,2010-07-20 08:40:46,1279640446537209,0.274981485 Messages/s,U
2010-07-20 08:38:46,1279640326561000,2010-07-20 08:38:46,1279640326529124,0.274989 Messages/s,U
2010-07-20 08:36:46,1279640206556000,2010-07-20 08:36:46,1279640206524330,0.283317543 Messages/s,U
2010-07-20 08:34:46,1279640086540000,2010-07-20 08:34:46,1279640086517645,0.274982493 Messages/s,U
2010-07-20 08:32:46,1279639966543000,2010-07-20 08:32:46,1279639966510022,0.291646426 Messages/s,U
2010-07-20 08:30:46,1279639846561000,2010-07-20 08:30:46,1279639846501672,0.308315369 Messages/s,U
2010-07-20 08:28:46,1279639726527000,2010-07-20 08:28:46,1279639726494673,0.291657509 Messages/s,U
2010-07-20 08:26:46,1279639606526000,2010-07-20 08:26:46,1279639606490890,0.266627739 Messages/s,U
2010-07-20 08:24:46,1279639486495000,2010-07-20 08:24:46,1279639486473368,0.258318523 Messages/s,U
2010-07-20 08:22:46,1279639366480000,2010-07-20 08:22:46,1279639366466497,0.274985902 Messages/s,U
2010-07-20 08:20:46,1279639246469000,2010-07-20 08:20:46,1279639246460346,0.283253871 Messages/s,U
2010-07-20 08:18:46,1279639126469000,2010-07-20 08:18:46,1279639126426669,0.274982804 Messages/s,U
2010-07-20 08:16:46,1279639006437000,2010-07-20 08:16:46,1279639006419168,0.283315503 Messages/s,U
```

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

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

Managing alarms

Alarms help you evaluate and quickly resolve trouble spots that sometimes occur during normal operation. You can customize StorageGRID alarms to address your specific monitoring requirements.

You can configure custom alarms either globally (Global Custom alarms) or for individual services (Custom alarms). You can create custom alarms with alarm levels that override Default alarms, and you can create alarms for attributes that do not have a Default alarm. Alarm customization is restricted to accounts with the Grid Topology Page Configuration and Other Grid Configuration permissions.

Attention: Using the Default alarm settings is recommended. 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 until it prevents a critical operation from completing. If you do need to change an alarm setting, you should discuss your proposed changes with technical support.

Steps

1. [Alarm classes](#) on page 60
2. [Alarm triggering logic](#) on page 61
3. [Viewing Default alarms](#) on page 65
4. [Creating custom service or component alarms](#) on page 66
5. [Creating Global Custom alarms](#) on page 68
6. [Disabling alarms](#) on page 71
7. [Configuring email notifications for alarms](#) on page 75

Related references

[Alarms reference](#) on page 166

Related information

[Administering StorageGRID](#)

Alarm classes

Alarms are separated into three 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 or 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 system-wide. You can also create a new Global Custom alarm that will monitor status system-wide. This can be useful for monitoring any customized conditions of your StorageGRID system.

Custom alarms

Custom alarms monitor the status of a single service or component. You can create a Custom alarm to override a Default alarm or Global Custom alarm at the service or component level. You can also create new Custom alarms based on the service's unique requirements.

Related tasks

- [Viewing Default alarms](#) on page 65
- [Disabling Default alarms for services](#) on page 72
- [Disabling a Default alarm system wide](#) on page 71
- [Creating Global Custom alarms](#) on page 68
- [Disabling Global Custom alarms for services](#) on page 74
- [Disabling Global Custom alarms system wide](#) on page 73
- [Creating custom service or component alarms](#) on page 66

Alarm triggering logic

Each alarm class is organized into a hierarchy of five severity levels from Normal to Critical. An alarm is triggered when a StorageGRID attribute reaches a threshold value that evaluates to true against a combination of alarm class and alarm severity level.

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 in the Alarms table on the Alarms page for a specific service or component (for example, **LDR > Storage > Alarms > Main**). 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. Custom alarms with alarm severities from Critical down to Notice.
2. Global Custom alarms with alarm severities from Critical down to Notice.
3. 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 Custom alarm for an attribute, the NMS service only evaluates the attribute value against Custom alarms. Global Custom alarms and Default alarms are not evaluated. Thus, an enabled Global Custom alarm for an attribute can meet the criteria needed to trigger an alarm, but it will not be triggered because a 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 examples

You can use these examples to understand how Custom alarms, Global Custom alarms, and Default alarms are triggered.

Example 1

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

	Threshold Values	
	Global Custom alarm (enabled)	Default alarm (enabled)
Notice	≥ 1500	≥ 1000
Minor	$\geq 15,000$	≥ 1000

	Threshold Values	
	Global Custom alarm (enabled)	Default alarm (enabled)
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.

Note, however, that if there are Custom alarms for an attribute, these alarms are still evaluated as Custom alarms have a higher priority than Global Custom alarms.

Example 2

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

	Threshold Values		
	Custom alarm (enabled)	Global Custom alarm (enabled)	Default alarm (enabled)
Notice	≥ 500	≥ 1500	≥ 1000
Minor	≥ 750	$\geq 15,000$	$\geq 10,000$
Major	$\geq 1,000$	$\geq 150,000$	$\geq 250,000$

If the attribute is evaluated when its value is 1000, a Major alarm is triggered and an email notification is sent to the designated personnel. The Custom alarm takes precedence over both the Global Custom alarm and Default alarm. A value of 1000 reaches the threshold value of the Major severity level for the Custom alarm. As a result, the attribute value triggers a Major level alarm.

Within the same scenario, if the Custom alarm is then disabled and the attribute value reevaluated at 1000, the alarm level is changed to Normal. The attribute value is evaluated against the threshold values of the Global Custom alarm, the next alarm class that is defined and enabled. A value of 1000 does not reach any threshold level for this alarm class. As a result, the attribute value is evaluated to be Normal and no notification is sent. The Notice level alarm from the previous evaluation is cleared.

Example 3

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

	Threshold Values		
	Custom alarm (disabled)	Global Custom alarm (enabled)	Default alarm (enabled)
Notice	≥ 500	≥ 1500	≥ 1000
Minor	≥ 750	$\geq 15,000$	$\geq 10,000$

	Threshold Values		
	Custom alarm (disabled)	Global Custom alarm (enabled)	Default alarm (enabled)
Major	>=1,000	>= 150,000	>= 250,000

If the attribute is evaluated when its value is 10,000, a Notice alarm is triggered and an email notification is sent to the designated personnel.

The Custom alarm is defined, but disabled; therefore, the attribute value is evaluated against the next alarm class. The Global Custom alarm is defined, enabled, and it takes precedence over the Default alarm. The attribute value is evaluated against the threshold values set for the Global Custom alarm class. A value of 10,000 reaches the Notice severity level for this alarm class. As a result, the attribute value triggers a Notice level alarm.

If the Global Custom alarm is then disabled and the attribute value reevaluated at 10,000, a Minor level alarm is triggered. The attribute value is evaluated against the threshold values for the Default alarm class, the only alarm class in that is both defined and enabled.

A value of 10,000 reaches the threshold value for a Minor level alarm. As a result, the Notice level alarm from the previous evaluation is cleared and the alarm level changes to Minor. An email notification is sent to the designated personnel.

Alarms of same severity

If two Global Custom or 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 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%	<=	100C		
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	Under 5C	=	500C		

Default Alarms

Filter by: Disabled Defaults

0 Result(s)

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

Apply Changes

Alarm class overrides

To override a class of alarms, disable all alarms within that class. If all alarms within a class for an attribute are disabled, the NMS service interprets the class as having no alarms configured for the attribute and evaluates the next lower class for enabled alarms.

For example, if an alarm is triggered at the Global Custom alarm class level, it means that there are no enabled alarms at the Custom alarms class level for that attribute.

To override a Default alarm, add a Global Custom alarm or Custom alarm for that attribute. This override is achieved because the NMS service does not evaluate lower priority alarm classes once an alarm setting is detected within a class. If this override is performed after an alarm has already been triggered, the override will not take effect until the alarm is triggered again.

Severity changes

If an alarm's severity changes, the severity is propagated up the network hierarchy as needed. If there is a notification configured, a notification is sent. The notification is sent only at the time the alarm enters or leaves the new severity level.

Notifications

A notification reports the occurrence of an alarm or the change of state for a service. It is an email communication to designated personnel that the system requires attention.

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



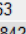

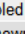


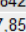

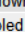


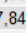

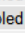
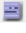

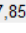

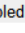





As new services are added through the addition of new grid nodes or sites, they inherit Default alarms and Global Custom alarms.

Alarms and tables

Alarm attributes displayed in tables can be disabled at the service, component, or system level. Alarms cannot be disabled for individual rows in a table.

For example, in the following figure, there are two critical Entries Available (VMFI) alarms. You can disable the VMFI alarm so that the Critical level VMFI alarm is not triggered (both currently Critical alarms would appear in the table as green); however, you cannot disable a single alarm in a table row so that one VMFI alarm displays as a Critical level alarm while the other remains green.

Volumes

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

Viewing Default alarms


You can view the Default alarms for a particular service or component, or you can view the Default alarms for the entire StorageGRID system.

Before you begin

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

Steps

1. To view the Default alarms for a particular service or component:
 - a. Select **Support > Grid Topology**.
 - b. Select **service or component > Configuration > Alarms**.

The Default alarms for the selected service or component are displayed.
2. To view the Default alarms for the entire StorageGRID system:
 - a. Select **Alarms**. Then, in the Alarms section of the menu, select **Global Alarms**.
 - b. For **Filter by** select **Attribute Code** or **Attribute Name**.
 - c. For **equals**, enter an asterisk: *
 - d. Click the arrow  or press **Enter**.

All Default alarms are listed.



Global Alarms

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

Global Custom Alarms (0 Result(s))

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

Default Alarms

Filter by Attribute Code equals *

221 Result(s)

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

Creating custom service or component alarms

Customizing alarm settings enables you to create a customized methodology for monitoring the StorageGRID system. You can create alarms on individual services or components in addition to creating global alarms.

Before you begin

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

About this task

You 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

Steps

1. Select **Support > Grid Topology**.
2. Select a service or component in the Grid Topology tree.
3. Select **Configuration > Alarms**.

Overview


Alarms

Reports





Configuration

Main

Alarms


Configuration (Alarms): LDR (DC2-S2-105-62) - LDR
Updated: 2017-03-28 10:48:58 PDT






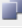


















Custom Alarms (0 Result(s))

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




Global Custom Alarms (0 Result(s))

Enabled	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions


Default Alarms (14 Result(s))






Enabled	Attribute	Severity	Message	Operator	Value	Actions
<input checked="" type="checkbox"/>	HTAS (Auto-Start HTTP)	 Notice	No	=	0	 
<input checked="" type="checkbox"/>	LDRE (LDR State)	 Notice	Standby	=	10	 
<input checked="" type="checkbox"/>	IRSU (Inbound Replication Status)	 Notice	Disabled	=	20	 
<input checked="" type="checkbox"/>	ORSU (Outbound Replication Status)	 Major	Storage Unavailable	=	10	 
<input checked="" type="checkbox"/>	ORSU (Outbound Replication Status)	 Notice	Disabled	=	20	 
<input checked="" type="checkbox"/>	HSTE (HTTP/CDMI State)	 Major	Offline	=	0,10,40	 
<input checked="" type="checkbox"/>	HSTU (HTTP/CDMI Status)	 Notice	Not Started	=	2	 
<input checked="" type="checkbox"/>	SSTS (Storage Status)	 Major	Volume(s) Unavailable	=	30	 

4. Add a new row to the Custom alarms table:

- Click **Edit**  (if this is the first entry) or **Insert**  to add a new alarm.
- Copy an alarm from the Default alarms or Global Custom alarms tables. Click **Copy**  next to the alarm you want to customize.

5. Make any necessary changes to the Custom alarm settings:

Heading	Description
Enabled	Select or unselect the check box to enable or disable the alarm.
Attribute	Select the name and code of the attribute being monitored from the list of all attributes applicable to the selected service or component. To display information about the attribute, click Info  next to the attribute's name.
Severity	The icon and text indicating the level of the alarm.
Message	The reason for the alarm (connection lost, storage space below 10%, and so on).

Heading	Description
Operator	<p>Operators for testing the current attribute value against the Value threshold:</p> <ul style="list-style-type: none"> • = equal to • > greater than • < less than • >= greater than or equal to • <= less than or equal to • ≠ not equal to
Value	<p>The alarm's threshold value used to test against the attribute's actual value using the operator.</p> <p>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/or ranges.</p>
Additional Recipients	<p>A supplementary list of email addresses to be notified when the alarm is triggered, in addition to the mailing list's configuration on the Configuration > Notifications 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.</p> <p>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

6. Click **Apply Changes**.

Creating Global Custom alarms

You can configure Global Custom alarms when you require a unique alarm that is the same for every service of the same type. Customizing alarm settings enables you to create a customized methodology for monitoring the StorageGRID system.



Before you begin


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

About this task





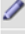







Global alarms override Default alarms. You should not change Default alarm values unless absolutely necessary. By changing Default alarms, you run the risk of concealing problems that might otherwise trigger an alarm.

Steps


1. Select **Alarms**. Then, in the Alarms section of the menu, select **Global Alarms**.
2. Add a new row to the Global Custom alarms table:
 - To add a new alarm, click **Edit**  (if this is the first entry) or **Insert** .


Global Alarms
 Updated: 2016-03-18 14:00:28 PDT



















Global Custom Alarms (0 Result(s))


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



Default Alarms





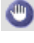

Filter by equals 

9 Result(s)

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

Apply Changes 

- To modify a Default alarm, search for the Default alarm.
 - a. Under Filter by, select either **Attribute Code** or **Attribute Name**.
 - b. Type a search string.
Specify four characters or use wildcards (for example, A??? or AB*). Asterisks (*) represent multiple characters, and question marks (?) represent a single character.
 - c. Click the arrow , or press **Enter**.
 - d. In the list of results, click **Copy**  next to the alarm you want to modify.
The Default alarm is copied to the Global Custom alarms table.
3. Make any necessary changes to the Global Custom alarms settings:

Heading	Description
Enabled	Select or unselect the check box to enable or disable the alarm.
Attribute	<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	<p>The alarm's threshold value used to test against the attribute's actual value using the operator.</p> <p>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.</p>
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.</p> <p>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 tasks

[Configuring email server settings for alarms](#) on page 76

Disabling alarms

Alarms are enabled by default, but you can disable alarms that are not required.

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.

Attention: There are consequences to disabling alarms and extreme care should be taken. Disabling an alarm can result in no alarm being triggered. Because alarms are evaluated by alarm class and then severity level within the class, disabling an alarm at a higher class does not necessarily result in a lower class alarm being evaluated. All alarms for a specific attribute must be disabled before a lower alarm class will be evaluated.

Steps

1. [Disabling a Default alarm system wide](#) on page 71
2. [Disabling Default alarms for services](#) on page 72
3. [Disabling Global Custom alarms system wide](#) on page 73
4. [Disabling Global Custom alarms for services](#) on page 74
5. [Clearing triggered alarms](#) on page 75


Disabling a Default alarm system wide


You can temporarily disable a Default alarm system wide.

Before you begin

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

Steps

1. Select **Alarms**. Then, in the Alarms section of the menu, select **Global Alarms**.
2. Search for the Default alarm to disable.
 - a. In the Default Alarms section, select **Filter by > Attribute Code** or **Attribute Name**.
 - b. Type a search string.
Specify four characters or use wildcards (for example, A??? or AB*). Asterisks (*) represent multiple characters, and question marks (?) represent a single character.
 - c. Click the arrow , or press **Enter**.

Note: Selecting **Disabled Defaults** displays a list of all currently disabled Default alarms.
3. In the Default Alarms table, click the Edit icon  next to the alarm you want to disable.
4. Clear the **Enabled** check box.



Global Alarms

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

Global Custom Alarms (0 Result(s))

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

Default Alarms

Filter by: Attribute Code equals U*

3 Result(s)

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

Apply Changes

5. Click **Apply Changes**.

The Default alarm is disabled system wide.

Note: When you disable a Default alarm system wide, an asterisk (*) is shown next to the **Enabled** check box on the **Support > Grid Topology > grid node > service or component > Configuration > Alarms** page. The asterisk indicates that the Default alarm has already been disabled system wide, even through the **Enabled** check box is selected.

Disabling Default alarms for services

To temporarily stop alarms for a specific service, you can disable Default alarms for that service.

Before you begin

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

Steps

- Select **Support > Grid Topology**.
- Select a service or component in the Grid Topology tree.
- Select **Configuration > Alarms**.
- In the **Default Alarms** table, click the Edit icon next to the alarm you want to disable.
- Clear the **Enabled** check box for the alarm.


Overview

Alarms

Reports

Configuration

Main Alarms


Configuration (Alarms): LDR (DC2-S2-105-62) - LDR
Updated: 2017-03-28 10:48:58 PDT

Custom Alarms (0 Result(s))

Enabled	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input type="checkbox"/>	LDRA (LDR Status)	Critical	Error	=	5		
<input checked="" type="checkbox"/>	HSTE (HTTP/CDMI State)	Critical		=			

Global Custom Alarms (0 Result(s))

Enabled	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
---------	-----------	----------	---------	----------	-------	-----------------------	---------

Default Alarms (14 Result(s))

Enabled	Attribute	Severity	Message	Operator	Value	Actions
<input checked="" type="checkbox"/>	HTAS (Auto-Start HTTP)	Notice	No	=	0	
<input type="checkbox"/>	LDRE (LDR State)	Notice	Standby	=	10	
<input checked="" type="checkbox"/>	IRSU (Inbound Replication Status)	Notice	Disabled	=	20	
<input checked="" type="checkbox"/>	ORSU (Outbound Replication Status)	Major	Storage Unavailable	=	10	

Note: If an asterisk (*) is shown next to the **Enabled** check box, the Default alarm has already been disabled system wide, even though the **Enabled** check box is selected.

6. Click **Apply Changes**.

The Default alarm is disabled for the service or component.

Disabling Global Custom alarms system wide

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


Before you begin

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

About this task

Note: Alarms cannot be disabled for individual rows in a table.

Steps

- Select **Alarms**. Then, in the Alarms section of the menu, select **Global Alarms**.
- In the Global Custom Alarms table, click **Edit**  next to the alarm you want to disable.
- Clear the **Enabled** check box.

Global Alarms
Updated: 2018-03-21 11:21:08 PDT

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

Filter by: Disabled Defaults

0 Result(s)

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

Apply Changes

4. Click **Apply Changes**.

The Global Custom alarm is disabled system wide.

Disabling Global Custom alarms for services

To disable a global alarm for a service, create another enabled global alarm for the attribute. You must create another enabled global alarm, because if all alarms within a class for an attribute are disabled, the NMS service interprets the class as having no alarms configured for the attribute and evaluates the next lower class for the enabled alarm.

Before you begin

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

About this task

Instead of creating a Global Custom alarm and disabling it for selected services, reconfigure the alarms such that you create individual local Custom alarms for the services that require the alarm. If you want to ensure that all these Custom alarms have the same configuration, you can create a Global Custom alarm, disable it, and then enable it for selected services as a Custom alarm.

If you want to create a Global Custom alarm and disable it for selected services, you must create a local Custom alarm for that service that will never be triggered. A local Custom alarm that is never triggered overrides all Global Custom alarms for that service.

Note: Alarms cannot be disabled for individual rows in a table.

Steps

- Select **Alarms**. Then, in the Alarms section of the menu, select **Global Alarms**.
- In the Global Custom alarm table, click next to the alarm you want to disable.
The alarm is copied to the Custom Alarms table.
- Clear the **Enabled** check box for the alarm.
- Click **Apply Changes**.

Related tasks

[Creating custom service or component alarms](#) on page 66

Clearing triggered alarms

You can clear a triggered alarm instead of acknowledging it.

Before you begin

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

About this task

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

Steps

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

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

3. Restart the NMS service:


```
service nms restart
```
4. Log out of the Admin Node:


```
exit
```

The alarm is cleared.

Related concepts

[Disabling alarms](#) on page 71

Configuring email notifications for alarms

StorageGRID system can automatically send email notifications when an alarm is triggered or a service state changes. By default, alarm email notifications are not sent. You must configure the email server and specify the email recipients.

Steps

1. [Types of alarm notifications](#) on page 76
2. [Configuring email server settings for alarms](#) on page 76
3. [Creating email templates](#) on page 78

4. [Creating mailing lists](#) on page 79
5. [Configuring global email notifications](#) on page 80
6. [Suppressing email notifications for a mailing list](#) on page 81
7. [Suppressing email notifications system wide](#) on page 82

Types of alarm notifications

The StorageGRID systems sends out two types of alarm notifications: severity level and service state.

Severity level notifications

Severity level notifications are sent when an event triggers an alarm at the selected alarm 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

Service state notifications are 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.

- Unknown
- Administratively Down

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

Related tasks

[Configuring global email notifications](#) on page 80

Configuring email server settings for alarms

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

Before you begin

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

About this task

Use these settings to define the SMTP server used for alarm email notifications and AutoSupport email messages. These settings are not used for alert notifications.

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

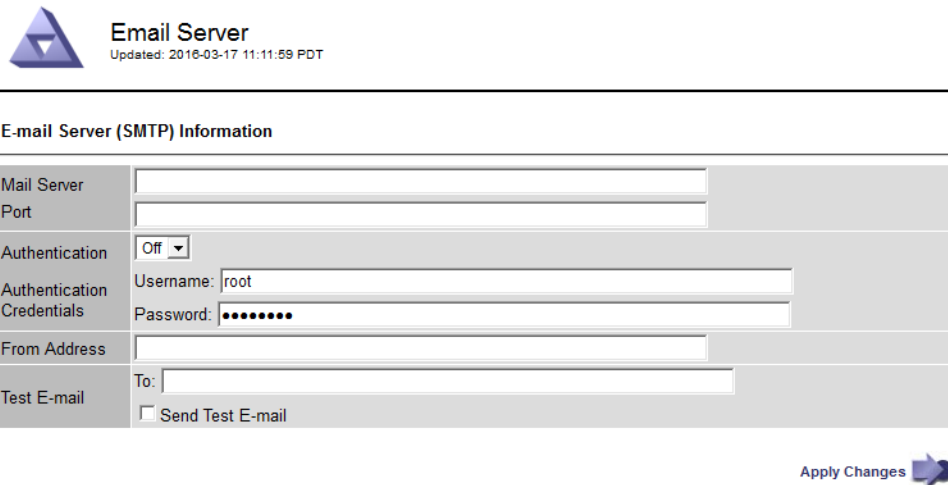
SMTP is the only protocol supported for sending email.

Steps

1. Select **Alarms**. Then, in the Alarms section of the menu, select **Email Setup**.

- From the Email menu, select **Server**.

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



- Add the following SMTP mail server settings:

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

- 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-delineated list of email addresses. Because the NMS service does not confirm success or failure when a test email is sent, you must be able to check the test recipient's inbox.
 - 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 MINS (NMS Notification Status) Minor alarm.

Related information*Administering StorageGRID***Creating email templates**

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

Before you begin



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

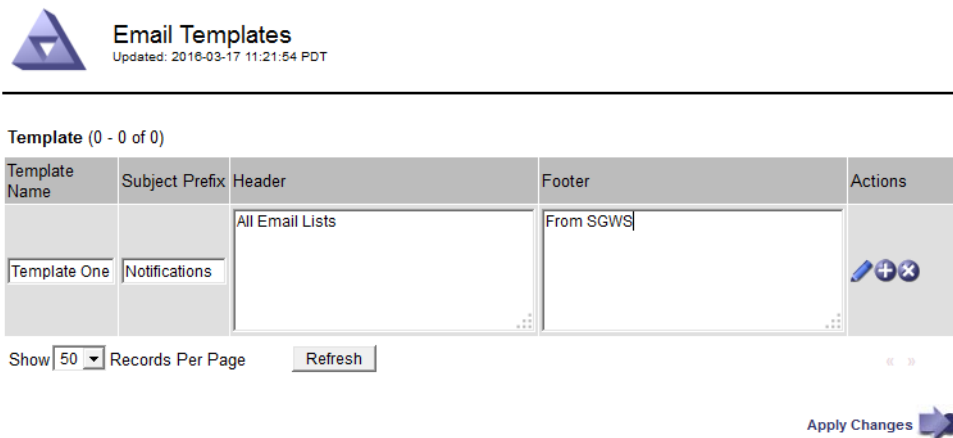
About this task

Use these settings to define the email templates used for 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 **Alarms**. Then, in the Alarms section of the menu, select **Email Setup**.
2. From the Email menu, select **Templates**.
3. Click **Edit**  (or **Insert**  if this is not the first template).



4. In the new row add the following:

Item	Description
Template Name	Unique name used to identify the template. Template names cannot be duplicated.
Subject Prefix	Optional. Prefix that will appear at the beginning of an email's subject line. Prefixes can be used to easily configure email filters and organize notifications.
Header	Optional. Header text that appears at the beginning of the email message body. Header text can be used to preface the content of the email message with information such as company name and address.

Item	Description
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

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



Before you begin

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

About this task

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

Steps

1. Select **Alarms**. Then, in the Alarms section of the menu, select **Email Setup**.
2. From the Email menu, select **Lists**.
3. Click **Edit**  (or **Insert**  if this is not the first mailing list).




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="text"/>	  

Show Records Per Page

Apply Changes 

4. In the new row, add the following:

Item	Description
Group Name	<p>Unique name used to identify the mailing list. Mailing list names cannot be duplicated.</p> <p>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.</p>
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 tasks

[Creating email templates](#) on page 78

Configuring global email notifications

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

Before you begin



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

About this task

Use these settings to configure global notifications for 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 **Alarms**. Then, in the Alarms section of the menu, select **Email Setup**.
2. From the Email menu, select **Notifications**.
3. Click **Edit**  (or **Insert**  if this is not the first notification).
4. Under **E-mail List**, select the mailing list.

5. Select one or more alarm severity levels and service states:

Category	Notification Type	Description
Severity Level	Notice	An unusual condition exists that does not affect normal operation.
	Minor	An abnormal condition exists that could affect operation in the future.
	Major	An abnormal condition exists that is currently affecting operation.
	Critical	An abnormal condition exists that has stopped normal operation.
Service State	Unknown	An unknown condition exists that has stopped normal service operation.
	Administratively Down	A condition whereby a service has been purposefully stopped.

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 tasks

[Creating mailing lists](#) on page 79

Suppressing email notifications for a mailing list

You can suppress notifications for a mailing list system-wide when you do not want a mailing list to receive notifications, for example while performing maintenance procedures.


Before you begin

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

About this task

Use these settings to suppress email notifications for alarms. These settings are not used for alert notifications.

Steps

1. Select **Alarms**. Then, in the Alarms section of the menu, select **Email Setup**.
2. From the Email menu, select **Notifications**.
3. Click **Edit**  next to the mailing list for which you want to suppress notifications.
4. Under **Suppress**, select the check box next to the mailing list you want to suppress, or select **Suppress** at the top of the column to suppress all mailing lists.
5. Click **Apply Changes**.

Notifications are suppressed for the selected mailing lists.

Suppressing email notifications system wide

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

Before you begin

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

About this task

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

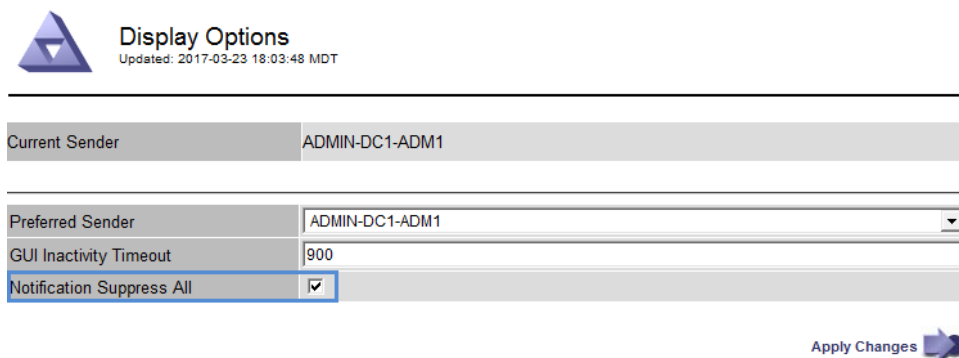
Note: This option does not suppress alert email notifications.

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

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

Steps

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



Display Options Updated: 2017-03-23 18:03:48 MDT	
Current Sender	ADMIN-DC1-ADM1
Preferred Sender	ADMIN-DC1-ADM1
GUI Inactivity Timeout	900
Notification Suppress All	<input checked="" type="checkbox"/>

Apply Changes

4. Click **Apply Changes**.

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






Notifications

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


All e-mail notifications are now suppressed.

Notifications (0 - 0 of 0)

	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 Records Per Page

(0 / 0)

Apply Changes 

Related information

[Administering StorageGRID](#)

[Upgrading StorageGRID](#)

Managing alerts (preview mode for 11.3)

Starting with StorageGRID 11.3, you can use alerts, in addition to alarms, to monitor various events and conditions within your StorageGRID system.

Choices

- [What alerts are](#) on page 84
- [Viewing all alerts](#) on page 87
- [Viewing a specific alert](#) on page 90
- [Managing alert rules](#) on page 92
- [Managing alert notifications](#) on page 99
- [Alerts reference](#) on page 106
- [Commonly used Prometheus metrics](#) on page 110

What alerts are

The new alerts system is available to preview in StorageGRID 11.3. The alerts system is designed to be easier to use and more powerful than the legacy alarms system. However, the alarms system continues to be the primary system for this release.

Benefits of the alerts system

The alerts system provides an easy-to-use interface for detecting, evaluating, and resolving the issues that can occur during StorageGRID operation. Available to preview in the StorageGRID 11.3 release, the alerts system offers significant benefits when compared to the alarms system:

- The new alerts system focuses on real problems in the system. Unlike some alarms in the legacy system, all of the new alerts are triggered for events that require your immediate attention, not for events that can safely be ignored.
- 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 are reporting the **Unable to communicate with node** alert, only one email is sent and the alert is shown as a group on the Alerts page.
- The Alerts page provides a more 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.
- Alerts use intuitive names and descriptions to help you understand more quickly what the problem is. Alert notifications include details about the node and site affected, the alert severity, the time when the alert rule was triggered, and the current value of metrics related to the alert.
- Both alert notifications and the alert listings on the Alerts page 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. You can silence an alert rule

for the entire grid, a single site, or a single node. The new silences functionality is more powerful than the acknowledge functionality in the alarms system.

- Creating custom alert rules is significantly easier and allows for greater functionality than creating custom alarms using the StorageGRID attributes system. 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.

Comparing alarms and alerts in StorageGRID 11.3

The new alerts system is available to preview in the StorageGRID 11.3 release; however, this new system does not currently offer comprehensive coverage or complete functionality. For example, alerts are not yet shown on the Dashboard or on the Nodes page, and alerts are not linked to the events logged in StorageGRID logs.

Attention: For StorageGRID 11.3, consider the alerts system to be a supplement to the alarms system, not a replacement for it. You must use the alarms system as your primary tool for detecting and resolving any issues with your system.

The following table shows the high-level similarities and differences between the alarms system and the preview alerts system in StorageGRID 11.3. Additional alerts and more complete alert functionality will be added in future StorageGRID releases.

	Alarms	Alerts (preview)
Where are they displayed	<p>When an alarm is triggered, you can see it in the following places:</p> <ul style="list-style-type: none"> • On the Dashboard • On the Nodes page • In the Grid Topology tree <p>You can also access details about current and historical alarms from the Alarms menu.</p>	<p>When an alert is triggered, you can only see it on the Alerts page (Alarms > Alerts (preview) > Alerts).</p>
Where are they managed	<p>Select Alarms. Then, use the options in the Alarms section of the menu.</p> <p>Managing alarms on page 60</p>	<p>Select Alarms. Then, use the options in the Alerts (Preview) section of the menu.</p> <p>Managing alerts (preview mode for 11.3) on page 84</p>

	Alarms	Alerts (preview)
Where are email notifications managed	<p>Select Alarms. Then, in the Alarms section of the menu, select Email Setup.</p> <p>Note: You can access the same Email Server page by selecting Support > AutoSupport and selecting Email Server from the AutoSupport menu.</p> <p>Configuring email notifications for alarms on page 75</p>	<p>Select Alarms. Then, in the Alerts (Preview) section of the menu, select Notifications.</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. You can use the same mail server for all notifications; however, the alerts system does not yet support SMTP server authentication (username, password).</p> <p>Managing alert notifications on page 99</p>
What user group permissions required	<ul style="list-style-type: none"> • Anyone who can sign in to the Grid Manager can monitor 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. 	You must have Root Access permission to view and manage alerts and notifications for the StorageGRID 11.3 preview.
Which Admin Nodes send notifications	A single Admin Node (the “preferred sender”).	All Admin Nodes, across all sites.
How are notifications suppressed	You can acknowledge an alarm after it has been triggered, or you can disable an alarm globally or for a particular service or component.	<p>You configure silences to suppress alert notifications for a specified amount of time.</p> <p>Each silence suppresses the notifications for an alert rule at one or more severities. You can silence an alert rule on the entire grid, a single site, or a single node.</p>
Where are recommended actions	<p>You must refer to the StorageGRID documentation.</p> <p>Alarms reference on page 166</p>	<p>Each alert rule includes a set of recommended actions, which are included in email notifications and available directly from the Alerts pages in the Grid Manager.</p> <p>As required, additional information is provided in the StorageGRID documentation.</p> <p>Alerts reference on page 106</p>

Viewing all alerts

You can view a list of all alerts currently affecting your StorageGRID system.

Before you begin

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

Steps

1. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **Alerts**.

The Alerts page appears. It lists all alerts currently affecting your StorageGRID system. 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 Alerts page is refreshed every two minutes.




Alerts

View the current alerts affecting your StorageGRID system.



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

2. Review the information in the table.

Column header	Description
Name	The name of the alert and its description.

Column header	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"> • 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. • 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. • 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.
Time triggered	How long ago the alert was triggered. If multiple alerts are grouped, the title row shows times for the most recent instance of the alert (<i>newest</i>) and the oldest instance of the alert (<i>oldest</i>).
Site/Node	The name of the site and node where the alert is occurring. If multiple alerts are grouped, the site and node names are not shown in the title row.
Status	Whether the alert is active or has been silenced. If multiple alerts are grouped and All alerts is selected in the drop-down, the title row shows how many instances of that alert are active and how many instances have been silenced.
Current values	<p>The current value of the metric that caused the alert to be triggered. For some alerts, additional values are shown to help you understand and investigate the alert. For example, the values shown for a Low object data storage alert include the percent 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:

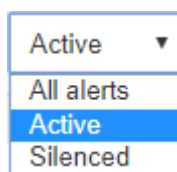
- To show the individual alerts in a group, click the down caret  in the heading, or click the group's name.
- To hide the individual alerts in a group, click the up caret  in the heading, or click the group's name.

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

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



5. To sort alerts or alert groups, click the up/down arrows  in each column header.
- When **Group alerts** is selected, both the alert groups and the individual alerts within each group are sorted. For example, you might want to sort the alerts in a group by **Time triggered** to find the most recent instance of a specific alert.
 - When **Group alerts** is unselected, the entire list of alarms 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 active and silenced alerts.
 - Select **Active** to view only the active alerts.
 - Select **Silenced** to view only the 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 “Viewing a specific alert.”

Related tasks

[Viewing a specific alert](#) on page 90

[Silencing alert notifications](#) on page 103

Viewing a specific alert

You can view more detailed information about a specific alert affecting your StorageGRID system. 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 an active alert or update the alert rule.

Before you begin

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

Steps

1. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **Alerts**.

The Alerts page appears.

2. From the table of alerts, expand the alert group, if required, and select the alert you want to view.

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

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

The Details dialog box for the alert appears.

Low installed node memory

The amount of installed memory on a node is low.

Recommended actions

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

See the instructions for your platform:

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

Time triggered

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

Status

Active ([silence this alert](#))

Site / Node

Data Center 2 / DC2-S1-99-56

Severity

8 Critical

Total RAM size

8.38 GB




Condition

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

Close

3. Review the details in the dialog box.

Information	Description
Title	The name of the alert.
First paragraph	The description of the alert.
Recommended actions	The recommended actions for resolving this alert.

Information	Description
Time triggered	The date and time the alert was triggered in your local time and in UTC.
Status	Whether the alert is active or has been silenced.
Site/Node	The name of the site and node where the alert is occurring.
Severity	<p>The severity of the alert.</p> <ul style="list-style-type: none"> • 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. • 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. • 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.
Current value	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.

4. To silence this alert:

- a. Click **silence this alert**.

The Create Silence page appears in a new browser tab. The **Alert Rule**, **Severity**, and **Nodes** fields are completed.

- b. Optionally, enter a description for this silence.
- c. Specify how long this alert should be silenced.
- d. Click **Save**.

The active alert is immediately silenced.

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

- a. Click **View conditions**.

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



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

6. To edit the alert rule that caused this alert to be triggered:

a. Click **Edit this rule**.

The Edit Rule page appears in a new browser tab.

b. Update the setting for the alert rule:

- If the alert was triggered by a default alert rule, you can disable the rule, update the conditions, or change the duration.
- If the alert was triggered by a custom alert rule, you can disable the rule, update the description and recommended actions, update the conditions, or change the duration.

c. Click **Save**.

Note: If you disabled the alert rule for an active alert, you must wait a few minutes for the alert to no longer display on the Alerts page.

Note: If you updated the conditions for an active 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 reflects the updated values.

7. To close the Details dialog box, click **Close**.

Related tasks

[Silencing alert notifications](#) on page 103

[Editing alert rules](#) on page 96

Managing alert rules

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

Steps

1. [Viewing alert rules](#) on page 92
2. [Creating custom alert rules](#) on page 94
3. [Editing alert rules](#) on page 96
4. [Removing custom alert rules](#) on page 99

Viewing alert rules

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

Before you begin

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

About this task

To view the list of all alert rules, go to “Alerts reference.”

Steps

1. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **Rules**.

The Alert Rules page appears.




Alert Rules

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.

<div> + Create custom rule Edit rule Remove custom rule </div>				
Name	Conditions	Type	Status	
<input type="radio"/> Cloud Storage Pool connectivity error The health check for Cloud Storage Pools detected one or more new errors.	sum(increase(storagegrid_data_mover_private_total_cloud_storage_pool_connection_errors[5m])) <i>Major</i> > 0	Default	Enabled	
<input type="radio"/> Expiration of server certificate for Management Interface The server certificate used for the management interface is about to expire.	storagegrid_servercertificate_management_interface_cert_expiry_days <i>Minor</i> <= 30 <i>Major</i> <= 14 <i>Critical</i> <= 0	Default	Enabled	
<input type="radio"/> Expiration of server certificate for Storage API Endpoints The server certificate used for the storage API endpoints is about to expire.	storagegrid_servercertificate_storage_api_endpoints_cert_expiry_days <i>Minor</i> <= 30 <i>Major</i> <= 14 <i>Critical</i> <= 0	Default	Enabled	
<input type="radio"/> Large audit queue The disk queue for audit messages is full.	sum(storagegrid_private_audit_messages_queued) without (job) <i>Major</i> >= 20000000 <i>Critical</i> >= 50000000	Default	Enabled	
<input type="radio"/> Low audit log disk capacity The space available for audit logs is low.	node_filesystem_avail(mountpoint="/var/local/audit/export") <i>Minor</i> <= 2000000000 <i>Major</i> <= 500000000 <i>Critical</i> <= 100000000	Default	Enabled	
<input type="radio"/> Low available node memory The amount of RAM available on a node is low.	node_memory_MemAvailable <i>Minor</i> <= 100000000 <i>Major</i> <= 50000000 <i>Critical</i> <= 10000000	Default	Enabled	
<input type="radio"/> Low installed node memory The amount of installed memory on a node is low.	node_memory_MemTotal <i>Major</i> < 24000000000 <i>Critical</i> <= 12000000000	Default	Enabled	
<input type="radio"/> Low metadata query performance The average time for Cassandra metadata queries is too long.	storagegrid_metadata_queries_average_latency_milliseconds <i>Minor</i> >= 1000	Default	Enabled	
<input type="radio"/> Low metadata storage The space available for storing object metadata is low.	storagegrid_storage_utilization_metadata_bytes / storagegrid_storage_utilization_metadata_allowed_bytes <i>Minor</i> >= 0.70 <i>Major</i> >= 0.90 <i>Critical</i> >= 1.00	Default	Enabled	

2. Review the information in the alert rules table:

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

Column header	Description
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	Whether this alert rule is currently enabled or disabled. The conditions for disabled alert rules are not evaluated, so no alerts are triggered.

Related references

[Alerts reference](#) on page 106

Creating custom alert rules

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

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have 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 < 24000000000`, first execute `node_memory_MemTotal >= 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 **Alarms**. Then, in the Alerts (Preview) section of the menu, select **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

5

minutes

Cancel

Save

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

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

4. Enter the following information:

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


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

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

A basic expression is usually of the form:

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

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

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 2.3.

Example

This expression causes an alert to be triggered if the amount of installed RAM for a node is less than 24,000,000,000 bytes (24 GB).

```
node_memory_MemTotal < 24000000000
```

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 references

[Commonly used Prometheus metrics](#) on page 110

Related information

[Administering StorageGRID](#)

[Prometheus: Query basics \(Version 2.3\)](#)

Editing alert rules

For default alert rules, you can change the rule's enabled/disabled state; the conditions for minor, major, and critical alerts; and the duration. For custom alert rules, you can also edit the rule's name, description, and recommended actions.

Before you begin

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

- You must have Root Access permission.

Steps

1. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **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 name, description, and recommended actions cannot be edited.

Edit Rule - Low installed node memory

Enabled ☒

Unique Name

Low installed node memory

Description

The amount of installed memory on a node is low.

Recommended Actions (optional)

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

Conditions ⓘ

Minor

Major

node_memory_MemTotal < 2400000000

Critical

node_memory_MemTotal <= 1200000000

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

Duration

2

minutes

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.

Note: If you disable the alert rule for an active alert, you must wait a few minutes for the alert to no longer display on the Alerts page.

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

Note: You cannot edit this information for default alert rules.


Field	Description
Unique Name	<p>A unique name for this rule.</p> <p>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.</p>

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.

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

Conditions ?

Minor	<input type="text"/>	
Major	<input type="text" value="node_memory_MemTotal < 24000000000"/>	
Critical	<input type="text" value="node_memory_MemTotal <= 14000000000"/>	

Note: If you update the conditions for an active 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 2.3.

Example

This expression causes an alert to be triggered if the amount of installed RAM for a node is less than 24,000,000,000 bytes (24 GB).

```
node_memory_MemTotal < 24000000000
```

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**.

Related references

[Commonly used Prometheus metrics](#) on page 110

Related information

[Administering StorageGRID](#)

[Prometheus: Query basics \(Version 2.3\)](#)

Removing custom alert rules

You can remove a custom alert rule if you no longer want to use it. You cannot remove a default alert rule.

Before you begin

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

Steps

1. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **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

You can manage the notifications that occur when an alert is triggered. You can configure the email server, specify the recipients of email notifications, and suppress (silence) the alert notifications for specific alerts.

Steps

1. [Setting up alert notifications](#) on page 100
2. [Information included in alert notifications](#) on page 102
3. [How StorageGRID groups alerts in notifications](#) on page 103
4. [Silencing alert notifications](#) on page 103
5. [Troubleshooting alert notifications](#) on page 105

Setting up alert notifications

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

Before you begin

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

About this task

If you want email notifications to be sent for alerts, you must configure an alerts email server and specify alert email recipients. Because alarms and alerts are independent systems, the email setup used for alarm notifications and AutoSupport messages is not used for alert notifications. You can use the same mail server for all notifications; however, the alerts system does not yet support SMTP server authentication (username, password).

In contrast to alarm notifications and AutoSupport messages, alert notifications are not currently sent by a single Admin Node (that is, they are not sent by the “preferred sender”). Instead, alert emails are sent by all Admin Nodes across all sites. For this reason, you must confirm that alert email can be sent from each Admin Node.

Steps

1. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **Notifications**.

The Notifications page appears.

Notifications

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

Enable Alert Notifications ☐

Save

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

The Email (SMTP) Server, Filters, and Recipients sections appear.

Notifications

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

Enable Alert Notifications ☒

Email (SMTP) Server

Mail Server


Port

Sender Email Address

Filters

Severity ☒ Minor, major, critical ☐ Major, critical ☐ Critical only

Recipients

Recipient 1 


3. In the **Email (SMTP) Server** section, enter the following information:

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.
Sender Email Address	A valid email address to use as the From address for alert notifications. For example, storagegrid-alerts@example.com.

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

Option	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.

5. 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.

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

- a. Click **Send Test Email**.

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

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

Note: If you do not receive the email within a few minutes, 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.

7. Click **Save**.

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

The email settings are saved.

Related tasks

[Troubleshooting alert notifications](#) on page 105

Information included in alert 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.

Alert email notifications are sent by all Admin Nodes across all sites in your StorageGRID system.

Email notifications include the following information:

NetApp StorageGRID

Low object data storage (6 alerts) 1

The space available for storing object data is low. 2

Recommended actions 3

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

DC1-S1-226

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

DC1-S2-227

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

Sent from: DC1-ADM1-225 5

Callout	Description
1	The name of the alert, followed by the number of active instances of this alert.
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 host name of the Admin Node that sent the notification.

How StorageGRID groups alerts in 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.

- 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 alert notifications are sent.
- 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. That is, if alert thresholds are reached for the minor, major, and critical severities, a notification is sent only 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.
- If an additional alert with the same name is triggered, StorageGRID waits 10 minutes before sending a new notification. The new notification reports all active alerts, even if they were reported previously.
- StorageGRID continues to send notifications once every 7 days until all instances of the alert are resolved or the alert rule is silenced.

Silencing alert notifications

Optionally, 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 silence an alert rule on the entire grid, a single site, or a single node.

Before you begin

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

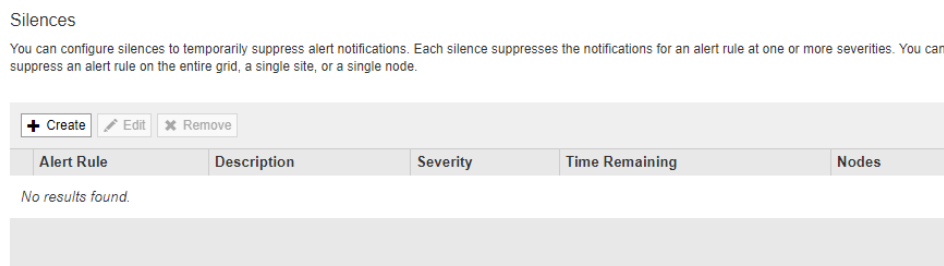
About this task

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

Steps

1. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **Silences**.

The Silences page appears.



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

Cancel
Save

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. Select All rules if you want to silence all alert rules using the criteria specified in this dialog box.
Description	Optionally, a description of the silence. For example, describe the purpose of this silence.
Duration	How long you want this silence to remain in effect, in minutes, hours, or days. A silence can be in effect from 5 minutes to 365 days.
Severity	Which alert severity or severities should be silenced. If the alert is triggered at one of the selected severities, no notifications are sent.
Nodes	<p>Which node or nodes you want this silence to apply to. You can suppress an alert rule 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 Alarms. Then, in the Alerts (Preview) section of the menu, select 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 Alarms. Then, in the Alerts (Preview) section of the menu, select 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 notifications to be sent and for the Alerts page to update.</p>

Troubleshooting alert notifications

If you are unable to receive the test alert email notification, follow these steps to resolve the issue.

Steps

1. Verify your settings.
 - a. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **Notifications**.
 - b. Verify that the Email (SMTP) Server settings are correct.
 - c. Verify that you have specified valid email addresses for the recipients.
2. Check your spam filter, and make sure that the email was not sent to a junk folder.
3. Ask your email administrator to confirm that emails from the sender address are not being blocked.
4. Collect a log file for the Admin Node, and then contact technical support.

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

Related tasks

[Collecting log files and system data](#) on page 124

Alerts reference

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

For information about the Prometheus metrics used in some of these alerts, see “Commonly used Prometheus metrics.”

Alert name	Related alarm	Description and recommended actions
Cloud Storage Pool connectivity error	<i>none</i>	<p>The health check for Cloud Storage Pools detected one or more new errors.</p> <ol style="list-style-type: none"> 1. Go to the Cloud Storage Pools section of the Storage Pools page. 2. Look at the Last Error column to determine which Cloud Storage Pool has an error. 3. See the instructions for administering StorageGRID. <p>Administering StorageGRID</p>
Expiration of server certificate for Management Interface	MCEP	<p>The server certificate used for the management interface is about to expire.</p> <ol style="list-style-type: none"> 1. Go to Configuration > Server Certificates. 2. In the Management Interface Server Certificate section, upload a new certificate. <p>Administering StorageGRID</p>
Expiration of server certificate for Storage API Endpoints	SCEP	<p>The server certificate used for accessing storage API endpoints is about to expire.</p> <ol style="list-style-type: none"> 1. Go to Configuration > Server Certificates. 2. In the Object Storage API Service Endpoints Server Certificate section, upload a new certificate. <p>Administering StorageGRID</p>
Large audit queue	AMQS	<p>The disk queue for audit messages is full.</p> <ol style="list-style-type: none"> 1. Check the load on the system—if there have been a significant number of transactions, the alert should resolve itself over time, and you can ignore the alert. 2. If the alert persists and increases in severity, view a chart of the queue size. If the number is steadily increasing over hours or days, the audit load has likely exceeded the audit capacity of the system. 3. Reduce the client operation rate or decrease the number of audit messages logged by changing the audit level for Client Writes and Client Reads to Error or Off (Configuration > Audit). <p>Understanding audit messages</p>

Alert name	Related alarm	Description and recommended actions
Low audit log disk capacity	VMFR	<p>The space available for audit logs is low.</p> <ol style="list-style-type: none"> 1. Monitor this alert to see if the issue resolves on its own and the disk space becomes available again. 2. Contact technical support if the available space continues to decrease.
Low available node memory	TMEM	<p>The amount of RAM available on a node is low.</p> <p>Low available RAM could indicate a change in the workload or a memory leak with one or more nodes.</p> <ol style="list-style-type: none"> 1. Monitor this alert to see if the issue resolves on its own. 2. If the available memory falls below the major alert threshold, contact technical support.
Low installed node memory	UMEM	<p>The amount of installed memory on a node is low.</p> <p>Increase the amount of RAM available to the virtual machine or Linux host. Check the threshold value for the major alert to determine the default minimum requirement for a StorageGRID node. See the installation instructions for your platform:</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux or CentOS installation • Ubuntu or Debian installation • VMware installation
Low metadata query performance	CQST	<p>The average time for Cassandra metadata queries is too long.</p> <p>An increase in query latency can be caused by a hardware change, such as replacing a disk, or a workload change, such as a sudden increase in ingests.</p> <ol style="list-style-type: none"> 1. Determine if there were any hardware or workload changes around the time the query latency increased. 2. If you are unable to resolve the problem, contact technical support.

Alert name	Related alarm	Description and recommended actions
Low metadata storage	CDLP	<p>The space available for storing object metadata is low.</p> <p>Critical alert</p> <ol style="list-style-type: none"> 1. Stop ingesting objects. 2. Immediately add Storage Nodes in an expansion procedure. <p>Major alert Immediately add Storage Nodes in an expansion procedure.</p> <p>Minor alert</p> <ol style="list-style-type: none"> 1. Monitor the rate at which object metadata space is being used. Select Nodes > Storage Nodes > Storage, and view the Storage Used - Object Metadata graph. 2. Add Storage Nodes in an expansion procedure as soon as possible. <p>Once new Storage Nodes are added, the system automatically rebalances object metadata across all Storage Nodes, and the alarm clears.</p> <p>Monitoring object metadata capacity for each Storage Node on page 44</p> <p>Expanding a StorageGRID system</p>
Low metrics disk capacity	VMFR	<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	SSTS	<p>The space available for storing object data is low.</p> <p>Perform an expansion procedure. You can add storage volumes (LUNs) to existing Storage Nodes, or you can add new Storage Nodes.</p> <p>Troubleshooting Low object data storage alerts on page 148</p> <p>Expanding a StorageGRID system</p>
Low root disk capacity	VMFR	<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 volume disk capacity	VMFR	<p>The space available for the /var/local mount point 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	Related alarm	Description and recommended actions
Node network connectivity error	NRER NTER	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.
Node not in sync with time source	NTSO	<p>The node's time is not in sync with the NTP time source. Monitor the alert for 10 minutes to see if the issue resolves on its own. If the alert persists:</p> <ol style="list-style-type: none"> 1. Verify that you have specified at least four external NTP sources, each providing a Stratum 3 or better reference. 2. Check that all NTP sources are operating normally. 3. Verify the connection to the NTP sources. Make sure they are not blocked by a firewall.
Objects lost	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. Lost and missing object data on page 149 2. When the underlying problem is resolved, reset the counter: <ol style="list-style-type: none"> a. Select Support > 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	<i>none</i>	<p>Too few Storage Nodes with the RSM service are running or available at a site.</p> <p>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>Administering StorageGRID</p>

Alert name	Related alarm	Description and recommended actions
Unable to communicate with node	<i>none</i>	<p>One or more services are unresponsive or cannot be reached by the metrics collection job.</p> <p>This alert indicates a problem connecting to the node or a service on the node. For example, the node might be powered down, there might be a network connectivity issue, or a service on the node might be stopped.</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. Determine if there a network connectivity issue between this node and the Admin Node. 3. Contact technical support.

Related references

[Commonly used Prometheus metrics](#) on page 110

Commonly used Prometheus metrics

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

The following table lists the most commonly used Prometheus metrics and provides a mapping of each metric to the equivalent attribute (used in the alarm system).

You can refer to this list to better understand the conditions in the default alert rules or to construct the conditions for custom alert rules. For a complete list of metrics, select **Help > API Documentation**.

Note: Metrics that include `_private_` in their names are intended for internal use only and are subject to change between StorageGRID releases without notice.

Note: Prometheus metrics are retained for 31 days.

Prometheus metric	Attribute	Description
storagegrid_appliance_failed_disks	BADD	For the storage controller in an appliance, the number of drives that are not optimal.
storagegrid_ilm_awaiting_background_objects	BQUZ	The total number of objects on this node awaiting ILM evaluation from the scan.
storagegrid_service_network_received_bytes	BREC	The total amount of data received by this service since installation.
storagegrid_service_network_transmitted_bytes	BTRA	The total amount of data sent by this service.
storagegrid_storage_utilization_metadata_bytes	CADL	The amount of object metadata on storage volume 0, in bytes.

Prometheus metric	Attribute	Description
storagegrid_storage_utilization_metadata_allowed_bytes	CEMS	The total space available on storage volume 0 for object metadata. Metadata Allowed Space (CEMS) is always less than the Metadata Reserved Space (CAWM) because a portion of the reserved metadata space is required for essential database operations, such as compaction and repair.
storagegrid_metadata_queries_average_latency_milliseconds	CQST	The average time required to run a query against the metadata store through this service.
storagegrid_ilm_awaiting_client_objects	CQUZ	The total number of objects on this node awaiting ILM evaluation from client operations (for example, ingest).
storagegrid_ilm_awaiting_client_evaluation_objects_per_second	EVRT	The current rate at which objects are evaluated against the ILM policy on this node.
storagegrid_http_sessions_incoming_attempted	HAIS	The total number of HTTP sessions that have been attempted to a Storage Node.
storagegrid_http_sessions_incoming_currently_established	HCCS	The number of HTTP sessions that are currently active (open) on the Storage Node.
storagegrid_http_sessions_incoming_failed	HEIS	The total number of HTTP sessions that failed to complete successfully, either due to a malformed HTTP request or a failure while processing an operation.
storagegrid_http_sessions_incoming_successful	HISC	The total number of HTTP sessions that have completed successfully.
storagegrid_content_objects_lost	LOST	The total number of objects this service detects as missing from the StorageGRID system.
storagegrid_ntp_chosen_time_source_offset_milliseconds	NTSO	Systematic offset of time provided by a chosen time source. Offset is introduced when the delay to reach a time source is not equal to the time required for the time source to reach the NTP client.
storagegrid_ilm_awaiting_total_objects	QUSZ	The total number of objects awaiting ILM evaluation.
storagegrid_service_restarts	RSTS	The total number of times the service has been restarted.
storagegrid_content_buckets_and_containers	SBKC	The total number of S3 buckets and Swift containers known by this Storage Node.
storagegrid_ilm_scan_objects_per_second	SCRT	The rate at which objects owned by this node are scanned and queued for ILM.
storagegrid_ilm_scan_period_estimated_minutes	SCTM	<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>

Prometheus metric	Attribute	Description
storagegrid_content_objects	SDOC	The total number of S3 and Swift data objects known by this Storage Node. Count is valid only for data objects created by client applications that interface with the system through S3 or Swift.
storagegrid_s3_operations_failed	SFAL	The total number of failed S3 operations (HTTP status codes 4xx and 5xx), excluding those caused by S3 authorization failure.
storagegrid_service_load	SLOD	The percentage of available CPU time currently being used by this service. Indicates how busy the service is. The amount of available CPU time depends on the number of CPUs for the server.
storagegrid_service_memory_usage_bytes	SMEM	The amount of memory (RAM) currently in use by this service. This value is identical to that displayed by the Linux top utility as RES.
storagegrid_storage_utilization_data_bytes	SPSD	An estimate of the total size of replicated and erasure coded object data on the Storage Node.
storagegrid_s3_data_transfers_bytes_ingested	SRXB	The total amount of data ingested from S3 clients to this Storage Node since the attribute was last reset.
storagegrid_storage_state_current	SSCR	The current state of the storage services. Attribute values are: <ul style="list-style-type: none"> 10 = Offline 15 = Maintenance 20 = Read-only 30 = Online
storagegrid_storage_status	SSTS	The current status of the storage services. Attribute values are: <ul style="list-style-type: none"> 0 = No Errors 10 = In Transition 20 = Insufficient Free Space 30 = Volume(s) Unavailable 40 = Error
storagegrid_s3_operations_successful	SSUC	The total number of successful S3 operations (HTTP status code 2xx).
storagegrid_storage_utilization_usable_space_bytes	STAS	The total amount of object storage space remaining. Calculated by adding together the amount of available space for all object stores on the Storage Node.
storagegrid_storage_utilization_total_space_bytes	STTS	The total amount of storage space allocated to all object stores.

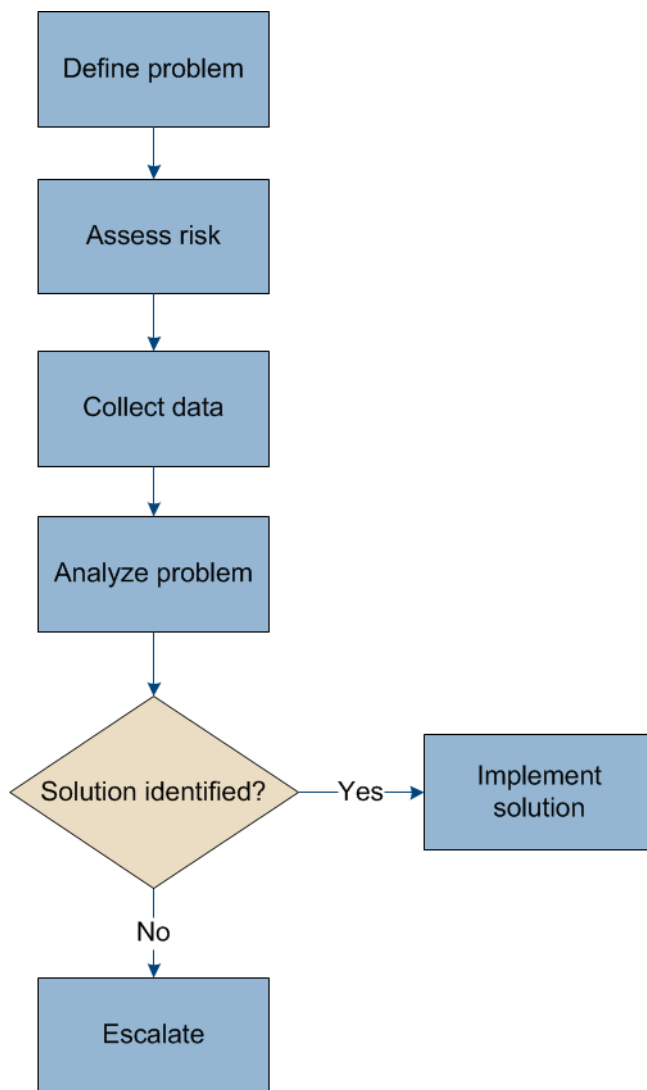
Prometheus metric	Attribute	Description
storagegrid_s3_data_transfers_bytes_retrieved	STXB	The total amount of data retrieved by S3 clients from this Storage Node since the attribute was last reset.
storagegrid_s3_operations_unauthorized	SUAU	The total number of failed S3 operations that are the result of an authorization failure.
storagegrid_service_cpu_seconds	SUTM	The cumulative amount of time that the CPU has been used by this service since installation.
storagegrid_service_runtime_seconds	SVRT	The total amount of time that the service has been running since installation.
storagegrid_service_uptime_seconds	SVUT	The total amount of time the service has been running since it was last restarted.
storagegrid_network_received_bytes	TRXB	The total amount of data received since installation.
storagegrid_network_transmitted_bytes	TTXB	The total amount of data sent since installation.
storagegrid_swift_operations_failed	WFAL	The total number of failed Swift operations (HTTP status codes 4xx and 5xx), excluding those caused by Swift authorization failure.
storagegrid_swift_data_transfers_bytes_ingested	WRXB	The total amount of data ingested from Swift clients to this Storage Node since the attribute was last reset.
storagegrid_swift_operations_successful	WSUC	The total number of successful Swift operations (HTTP status code 2xx).
storagegrid_swift_data_transfers_bytes_retrieved	WTXB	The total amount of data retrieved by Swift clients from this Storage Node since the attribute was last reset.
storagegrid_swift_operations_unauthorized	WUAU	The total number of failed Swift operations that are the result of an authorization failure (HTTP status codes 401, 403, 405).

Troubleshooting a StorageGRID system

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

Overview of problem determination

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



Steps

1. [Defining the problem](#) on page 115
2. [Assessing the risk and impact on the system](#) on page 115

3. [Collecting data](#) on page 115
4. [Analyzing data](#) on page 135
5. [Escalation information checklist](#) on page 136

Defining the problem

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

This table provides an example of the type of information that you would collect to define the 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. Many alarms have been triggered.
When did the problem start?	Object ingest was first denied at about 14:50 on January 8, 2019.
How did you first notice the problem?	Notified by client application. Also received email notifications.
Is the problem reproducible?	Problem is ongoing.
What is the frequency of the problem?	This is the first time this has happened.

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 alarms 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 the problem has been defined and its risk and impact assessed, collect data for analysis. You can use the following steps to collect data, depending on the nature of the problem.

Before you begin

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

Related information

[Administering StorageGRID](#)

Performing ingest and retrieval tests

To troubleshoot ingest and retrieval performance issues, you can perform a simple test that uses a workstation in the place of the actual client application and analyze the store/retrieve performance during the test compared to the performance normally seen with the client application.

Messages in the audit log indicate the total time required to execute 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, S3 GET, S3 HEAD, S3 Metadata Updated, S3 POST, S3 PUT
- Swift DELETE, Swift GET, Swift HEAD, Swift PUT

Related information

[Understanding audit messages](#)

Listing recent changes

It is important that you make a list of any recent changes made to the StorageGRID system or its environment.

List of recent changes
Was the StorageGRID system recently installed, expanded, or recovered?
Has any hardware been repaired or changed recently?
Have any object storage applications or users been added or removed?
Have there been any changes to the network infrastructure? For example, VLANs, routers, or DNS.
Have any long term grid tasks been recently triggered?
Is data migration taking place?
Have any changes been made to user authentication, such as adding a tenant or changing LDAP configuration?
Have any changes been made to the ILM policy?
Have any changes been made to storage compression or encryption?
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?

Checking connectivity status

Confirm that servers are online and connected to each other. In the Grid Topology tree, look for grid nodes whose state is unknown (blue) or are administratively down, that is, that have been purposely stopped (gray).

If all grid nodes are green, the primary Admin Node is connected to all other grid nodes. If grid nodes are blue, there is likely a connectivity issue.

Reviewing alarms

An alarm is triggered when the value of an attribute reaches a set alarm threshold value. Node icons that are not green indicate that an alarm has been triggered.

About this task

To check and review alarms, see the following table.

To	Do this
Get a list of all current alarms for the StorageGRID system	<p>Select Alarms. Then, in the Alarms section of the menu, select Current Alarms.</p> <p>The alarms are sorted by severity.</p> <p>Note: The alarms indicator immediately tells you the most serious status (state or alarm) affecting the system.</p>
Get a list of all alarms triggered over a period of time	<ol style="list-style-type: none"> 1. Select Alarms. Then, in the Alarms section of the menu, select Historical Alarms. 2. Do one of the following: <ul style="list-style-type: none"> • Click one of the time periods. • Enter a custom range and click Custom Query.
Find out how often alarms have been triggered for a particular attribute	<ol style="list-style-type: none"> 1. Select Support > Grid Topology. 2. Select <i>grid node > service or component > Alarms > History</i>. 3. Select the attribute from the list. 4. Do one of the following: <ul style="list-style-type: none"> • Click one of the time periods. • Enter a custom range and click Custom Query. The alarms are listed in reverse chronological order. 5. To return to the alarms history request form, click History.
Determine whether an alarm has been disabled globally	<ol style="list-style-type: none"> 1. Select Alarms. Then, in the Alarms section of the menu, select Global Alarms. 2. Enter filtering criteria, or select Disabled Defaults. If the alarm is disabled, the Enabled check box is unselected.

Related concepts

[Viewing node icons](#) on page 36

Related references

[Alarms reference](#) on page 166

Plotting trends

Reports (both chart and text) are an invaluable tool when troubleshooting. The fastest way to create a chart is to click the **Chart** button on the Overview tab of a component or service. This is known as an immediate report. You can also create charts from the Reports tab.

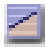
Related concepts

[Using reports](#) on page 52

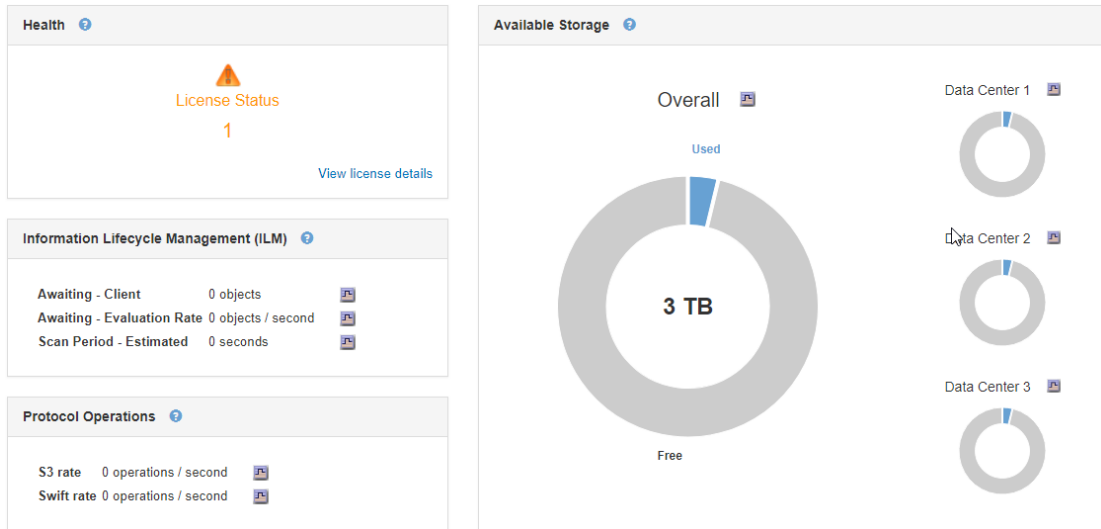
Establishing baselines

Baseline information is operational data during normal system operations that provides clues that can help you solve problems.

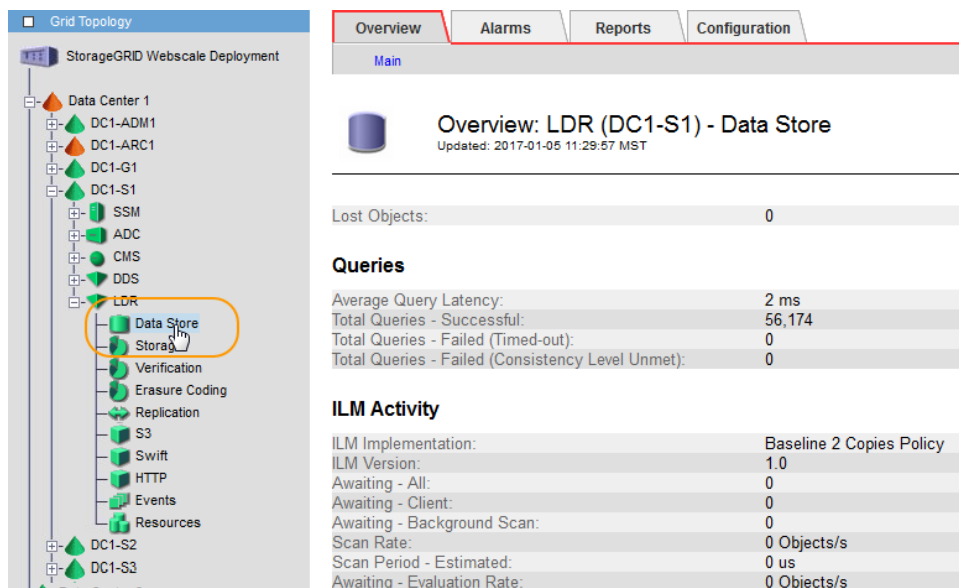
You can gather useful baseline information using the StorageGRID system:

Property	Value	How to obtain
Average storage consumption	_____/consumed/day GB _____/consumed/day %	Go to the Dashboard in the Grid Manager, and click  for one of the Available Storage charts. Find a period where the operation is fairly stable and estimate the daily storage consumption rate in bytes and in percentage. You can collect this information for the entire system or for a specific data center.
Rate of S3/Swift Operations	_____ operations/s	Go to the Dashboard in the Grid Manager. Under Protocol Operations, view the values for S3 rate and Swift rates.
Failed S3/Swift Operations	_____	Select Support > Grid Topology . Then, select Overview tab > Main . Under API Operations, view the value for S3 Operations - Failed or Swift Operations - Failed.
ILM Evaluation Rate	_____ objects/second	Select Nodes > deployment > ILM . Choose a range (1 hour, 1 day, 1 week, 1 month, 1 year, or Custom). Hover your cursor over the ILM Queue graph to view the Evaluation Rate.
Average Query Latency	_____ milliseconds	Select Nodes > Storage Node > Objects . In the Queries table, view the value for Average Latency.
Failed Queries	_____	Select Nodes > Storage Node > Objects . In the Queries table, view the value for Queries - Failed (timed-out).
Failed Queries for Consistency Level Unmet	_____	Select Nodes > Storage Node > Objects . In the Queries table, view the value for Queries - Failed (consistency level unmet).

The following example screenshot shows the Dashboard:



The following example screenshot shows how you can select each attribute under each service to view baseline values for that attribute on the Overview tab:



Monitoring events

You can monitor any events that are detected by a grid node to help with troubleshooting. The Last Event provides an area of focus when an error occurs.

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

Related tasks

[Resetting event counts](#) on page 121

Reviewing previous events

You can view previous events to help isolate issues that occurred in the past.

Steps

1. Select **Support > 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**.

Reports (Text): SSM (170-41) - Events

Attribute: Last Event Results Per Page: 20 Start Date: 2009/04/15 15:19:53
 Quick Query: Last 5 Minutes Update Raw Data: ☒ End Date: 2009/04/15 15:24:53

Text Results for Last Event
 2009-04-15 15:19:53 PDT To 2009-04-15 15:24:53 PDT

1 - 2 of 2

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

Related concepts

[Using reports](#) on page 52

Resetting event counts

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

Before you begin

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

Steps






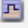

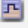





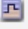

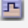

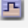

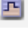

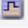

- Select **Nodes > Grid Node > Events**.
- Make sure that any event with a count greater than 0 has been resolved.
- Click **Reset event counts**.

[Overview](#)
[Hardware](#)
[Network](#)
[Storage](#)
[Events](#)

Events

Last Event

No Events

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

[Reset event counts](#)

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



Events
Updated: 2016-03-24 15:16:20 PDT

Custom Events (1 - 1 of 1)

Event	Actions
xfs internal error	   
shutdown	   

Show 10 Records Per Page Refresh

Previous « 1 » Next

Apply Changes 










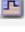








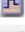


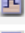

4. Click **Apply Changes**.
5. Select **Nodes**. Then, select **grid node > Events**.
6. Locate the entry for Custom Events in the **Events** table, and monitor the value for **Count**.
If the count increases, a custom event you are monitoring is being triggered on that grid node.


[Overview](#)
[Hardware](#)
[Network](#)
[Storage](#)
[Events](#)

Events

Last Event

No Events

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

[Reset event counts](#)


Resetting the count of custom events to zero

You can reset the count of custom events to zero.

About this task

Resetting a counter causes the alarm to be triggered by the next event. In contrast, when you acknowledge an alarm, that alarm is only re-triggered if the next threshold level is reached.

Steps

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

Overview


Alarms

Reports

Configuration

Main

Alarms



Configuration: SSM (DC2-ADM1) - Events

Updated: 2018-04-11 10:35:44 MDT

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**.

Collecting log files and system data

To help troubleshoot a problem, you can retrieve log files and system data (including configuration data) for your StorageGRID system. This information is retrieved using the Grid Manager.

Before you begin

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

About this task

You can use the Grid Manager to gather log files, system data, and configuration data from any grid node for the time period that you select. Data is collected and archived in a `.tar.gz` file that you can then download to your local computer.

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

Steps

1. Select **Support > Logs**.

Logs

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

The screenshot shows the 'Logs' collection interface. On the left, a tree view shows the grid structure: 'StorageGRID Webscale Deployment' (expanded) contains 'Data Center 1' (expanded), 'Data Center 2' (expanded), and 'Data Center 3' (expanded). Under 'Data Center 1', nodes DC1-ADM1, DC1-ARC1, DC1-G1, DC1-S1, DC1-S2, and DC1-S3 are listed. Under 'Data Center 2', nodes DC2-ADM1, DC2-S1, DC2-S2, and DC2-S3 are listed. Under 'Data Center 3', nodes DC3-S1, DC3-S2, and DC3-S3 are listed. On the right, the 'Log Start Time' is set to 2018-04-18 01:38 PM MDT, and the 'Log End Time' is set to 2018-04-18 05:38 PM MDT. Below these are a 'Notes' text area and a 'Provisioning Passphrase' text field. A 'Collect Logs' button is located at the bottom right.

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

As required, you can collect log files for the entire grid or an entire data center site.

3. Select a **Start Time** and **End Time** to set the time range of the data to be included in the log files.

If you select a very long time period or collect logs from all nodes in a large grid, the log archive could become too large to be stored on a node, or too large to be collected to the primary Admin Node for download. If this occurs, you must restart log collection with a smaller set of data.

4. Optionally type notes about the log files you are gathering in the **Notes** text box.

You can use these notes to give technical support information about the problem that prompted you to collect the log files. Your notes are added to a file called `info.txt`, along with other information about the log file collection. The `info.txt` file is saved in the log file archive package.

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

6. Click **Collect Logs**.

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

Logs

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

Log collection is in progress.

Last Collected

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

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

Notes

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

23%

Collecting logs: 10 of 13 nodes remaining

Download

Delete

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

You can use the Logs page to monitor the progress of log file collection for each grid node.

If you receive an error message about log size, try collecting logs for a shorter time period or for fewer nodes.

7. Click **Download** when log file collection is complete.

The `.tar.gz` file contains all log files from all grid nodes where log collection was successful. Inside the combined `.tar.gz` file, there is one log file archive for each grid node.

After you finish

You can re-download the log file archive package later if you need to.

Optionally, you can click **Delete** to remove it and free up disk space, though there is no requirement to do so. The current log file archive package is automatically removed the next time you collect log files.

Related concepts

[Log files](#) on page 191

Reviewing audit messages

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

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

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

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

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

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

Related information

[Understanding audit messages](#)

[Administering StorageGRID](#)

Triggering an AutoSupport message

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

Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- The StorageGRID system's email server must be correctly configured. See the instructions for administering StorageGRID.

About this task

The AutoSupport message includes the following information:

- System- and site-level attribute information
- All alarms raised in the last seven days
- The current status of all grid tasks (including historical data)
- Events information
- Admin Node database usage
- The number of lost or missing objects (zero or more)
- The system's ILM policy

Steps

1. Select **Support > AutoSupport**.

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

Updated: 2016-03-21 11:45:17 PDT

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

3. Click Send.**Related information**[Administering StorageGRID](#)**Reviewing support metrics**

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

Before you begin

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

About this task

The Metrics page allows you to access the Prometheus and Grafana user interfaces. Prometheus is open-source software for collecting metrics. Grafana is open-source software for metrics visualization.

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

Steps

1. As directed by technical support, select **Support > Metrics**.

The Metrics page appears.

Metrics

Access charts and metrics to help troubleshoot issues.

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

Prometheus

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

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

- <https://quicksilver.vtc.englab.netapp.com/metrics/graph>

Grafana

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

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

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

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.

Prometheus
Alerts
Graph
Status ▾
Help

☐ Enable query history

Execute

- insert metric at cursor - ▾

Graph

Console

Element	Value
no data	

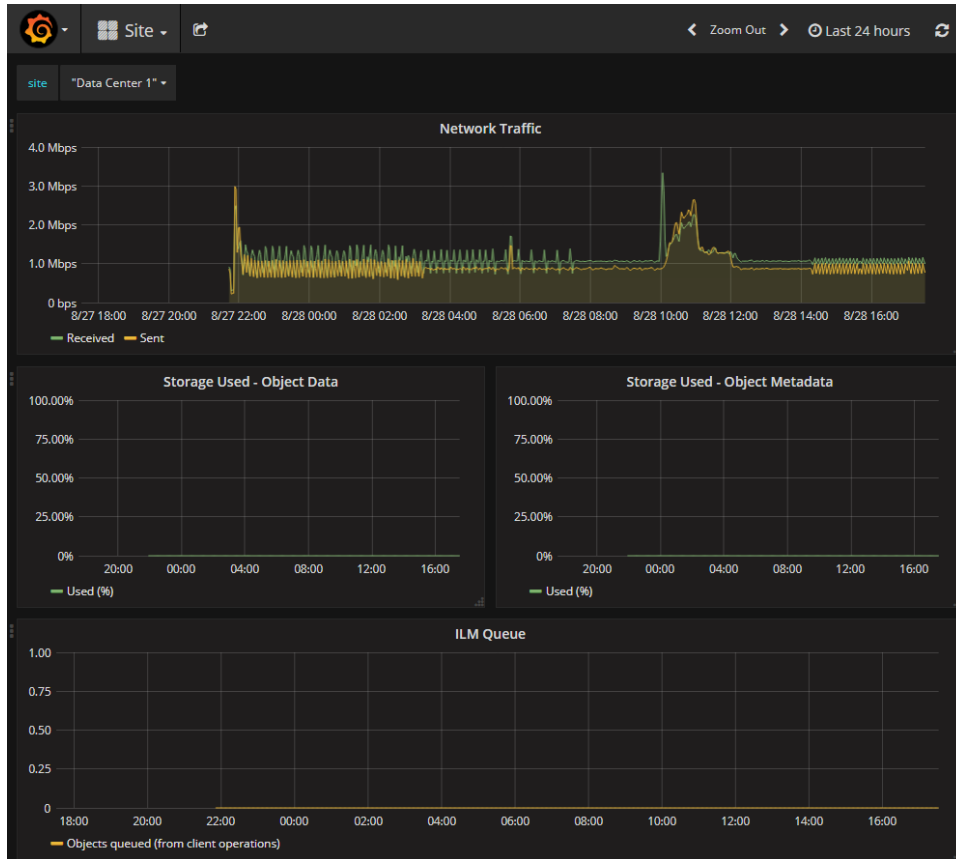
Remove Graph

Add Graph

Note: Metrics that include `_private_` in their names are intended for internal use only and are subject to change between StorageGRID releases without notice.

3. To access pre-constructed dashboards containing graphs of StorageGRID metrics over time, click the links in the **Grafana** section.

The Grafana interface for the link you selected appears.



Related references

[Commonly used Prometheus metrics](#) on page 110

Running foreground verification

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

Before you begin

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

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

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

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

About this task

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

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

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

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

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

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

Steps

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

Overview Alarms Reports **Configuration**

Main Alarms

Configuration: LDR (dc1-cs1-99-82) - Verification
Updated: 2015-08-19 14:07:04 PDT

Reset Missing Objects Count ☐

Foreground Verification

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

Background Verification

Verification Rate

Reset Corrupt Objects Count ☐

Apply Changes

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**.

If the count for the attribute **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 count for the attribute **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 > 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 valid`, 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

[Administering StorageGRID](#)

[Recovery and maintenance](#)

Confirming object data locations

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

Before you begin

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

Steps

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

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

Object Metadata Lookup

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

Identifier

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, including the name of the grid node and the full path to the disk location of the object.
- For erasure-coded object copies, the current storage location of each fragment, including the name of the grid node and the type of fragment (data or parity).
- 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, 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	F9251F03-F791-4843-A85F-09B8AE777D07
Name	base-8105145
Container	source
Account	t-1554403014
Size	1 byte
Creation Time	2019-04-04 14:46:54 MDT
Modified Time	2019-04-04 14:46:54 MDT

Replicated Copies

Node	Disk Path
99-98	/var/local/rangedb/0/p/1A/08/01DF98D08B4E950Fp
99-99	/var/local/rangedb/0/p/00/09/01DF98D08B4E950Fp

Raw Metadata

```
{
  "TYPE": "CTNT",
  "CHND": "F9251F03-F791-4843-A85F-09B8AE777D07",
  "NAME": "base-8105145",
  "CBID": "0x01DF98D08B4E950F",
  "PHND": "AB2A45E2-5708-11E9-8B1F-EFC900C06E1F",
  "PPTH": "source",
  "META": {
    "BASE": {
      "PAWS": "2",
      "ACCT": "39212516505481629340",
      "ctp": "binary/octet-stream"
    },
    "BYCB": {
      "SHSH": "MD5D 0x92EB5FFEE6AE2FEC3AD71C777531578F",
      "CSIZ": "1",
      "DCT7": "C7E"
    }
  }
}
```

Related information

[Administering StorageGRID](#)

[Implementing S3 client applications](#)

[Implementing Swift client applications](#)

Analyzing data


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

The analysis is problem-dependent, but in general:

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

Escalation information checklist

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

	Item	Notes
	Problem statement	What are the problem symptoms? What is the history of the problem?
	Impact assessment	What is the severity of the problem?
	Grid ID	Click Configuration > Grid Options > Overview to locate the Grid ID.
	Software version	Click Help > About to see the StorageGRID version.
	Customization	Summarize how the StorageGRID system is configured. For example, list whether it uses storage compression, storage encryption, or compliance, and summarize how ILM is configured. Does ILM make replicated or erasure coded objects? Does ILM ensure site redundancy? Do ingest rules use the Strict, Balanced, or Dual Commit ingest behaviors?
	Grid specification file	Export the latest version of the grid specification file for your StorageGRID system. From the Grid Manager, go to Configuration > Grid Options > Configuration , and click the export button.
	Baseline information	Collect baseline information regarding ingest operations, retrieval operations, and storage consumption.
	Recent changes	Summarize any recent changes made to the system or its environment.

Related concepts

[Establishing baselines](#) on page 118

Related references

[Defining the problem](#) on page 115

[Assessing the risk and impact on the system](#) on page 115

[Listing recent changes](#) on page 116

Related information

[Administering StorageGRID](#)

Troubleshooting Admin Nodes

There are several tasks you can perform to help determine the source of Admin Node related problems.

Troubleshooting sign-on errors

If you experience an error when you are signing in to a StorageGRID Admin Node, your system might have an issue with the identity federation configuration, a networking or hardware problem, an issue with Admin Node services, or an issue with the Cassandra database on connected Storage Nodes.

Before you begin

- You must have the `Passwords.txt` file.
- You must have specific access permissions.

About this task

Use these troubleshooting guidelines if you see any of the following error messages when attempting to sign in to an Admin Node:

- `Your credentials for this account were invalid. Please try again.`
- `Waiting for services to start...`
- `Internal server error. The server encountered an error and could not complete your request. Please try again. If the problem persists, contact Technical Support.`
- `Unable to communicate with server. Reloading page...`

Steps

1. Wait 10 minutes, and try signing in again.
If the error is not resolved automatically, go to the next step.
2. If your StorageGRID system has more than one Admin Node, try signing in to the Grid Manager from another Admin Node.
 - If you are able to sign in, you can use the **Dashboard, Nodes, Alarms, and Support > Grid Topology** options to help determine the cause of the error.
 - If you have only one Admin Node or you still cannot sign in, go to the next step.
3. Determine if the node's hardware is offline.
4. If single sign-on (SSO) is enabled for your StorageGRID system, refer to the steps for “Configuring single sign-on” in the guide for administering StorageGRID.
You might need to temporarily disable and re-enable SSO for a single Admin Node to resolve any issues.
Note: If SSO is enabled, you cannot sign on using a restricted port. You must use port 443.
5. Determine if the account you are using belongs to a federated user.
If the federated user account is not working, try signing in to the Grid Manager as a local user, such as root.

- If the local user can sign in:
 - a. Review any displayed alarms.
 - b. Select **Configuration > Identity Federation**.
 - c. Click **Test Connection** to validate your connection settings for the LDAP server.
 - d. If the test fails, resolve any configuration errors.
 - If the local user cannot sign in and you are confident that the credentials are correct, go to the next step.
6. Use Secure Shell (ssh) to log in to the Admin Node:
 - a. Enter the following command: `ssh admin@Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.

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

7. View the status of all services running on the grid node:

storagegrid-status

Make sure the `nms`, `mi`, `nginx`, and `mgmt api` services are all running.

The output is updated immediately if the status of a service changes.

```
$ storagegrid-status
Host Name          99-211
IP Address         10.96.99.211
Operating System Kernel 4.9.0      Verified
Operating System Environment Debian 9.6    Verified
StorageGRID Webscale Release 11.2.0    Verified
Networking         Verified
Storage Subsystem  Verified
Database Engine    5.5.9999+default Running
Network Monitoring 11.2.0      Running
Time Synchronization 1:4.2.8p10+dfsg Running
ams                11.2.0      Running
cmn                11.2.0      Running
nms                11.2.0      Running
ssm                11.2.0      Running
mi                11.2.0      Running
dynip             11.2.0      Running
nginx             1.10.3      Running
tomcat            8.5.14      Running
grafana           4.2.0      Running
mgmt api          11.2.0      Running
prometheus        11.2.0      Running
persistence       11.2.0      Running
ade exporter      11.2.0      Running
alertmanager      11.2.0      Running
attrDownPurge     11.2.0      Running
attrDownSamp1     11.2.0      Running
attrDownSamp2     11.2.0      Running
node exporter     0.13.0+ds   Running
sg snmp agent     11.2.0      Running
```

8. Confirm that the Apache web server is running:

```
# service apache2 status
```

9. Use Lumberjack to collect logs:

```
# /usr/local/sbin/lumberjack.rb
```

If the failed authentication happened in the past, you can use the `-start` and `-end` Lumberjack script options to specify the appropriate time range. Use `lumberjack -h` for details on these options.

The output to the terminal indicates where the log archive has been copied.

10. Review the following logs:

- `/var/local/log/bycast.log`
- `/var/local/log/bycast-err.log`
- `/var/local/log/nms.log`
- `**/*commands.txt`

11. If you could not identify any issues with the Admin Node, issue either of the following commands to determine the IP addresses of the three Storage Nodes that run the ADC service at your site. Typically, these are the first three Storage Nodes that were installed at the site.

```
# cat /etc/hosts
```

```
# vi /var/local/gpt-data/specs/grid.xml
```

Admin Nodes use the ADC service during the authentication process.

12. From the Admin Node, log in to each of the ADC Storage Nodes, using the IP addresses you identified.

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

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

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

```
storagegrid-status
```

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

14. Repeat steps [9](#) and [10](#) to review the logs on the Storage Nodes.

15. If you are unable to resolve the issue, contact technical support.

Provide the logs you collected to technical support.

Related concepts

[Log files](#) on page 191

Related information

[Administering StorageGRID](#)

Checking the status of an unavailable Admin Node

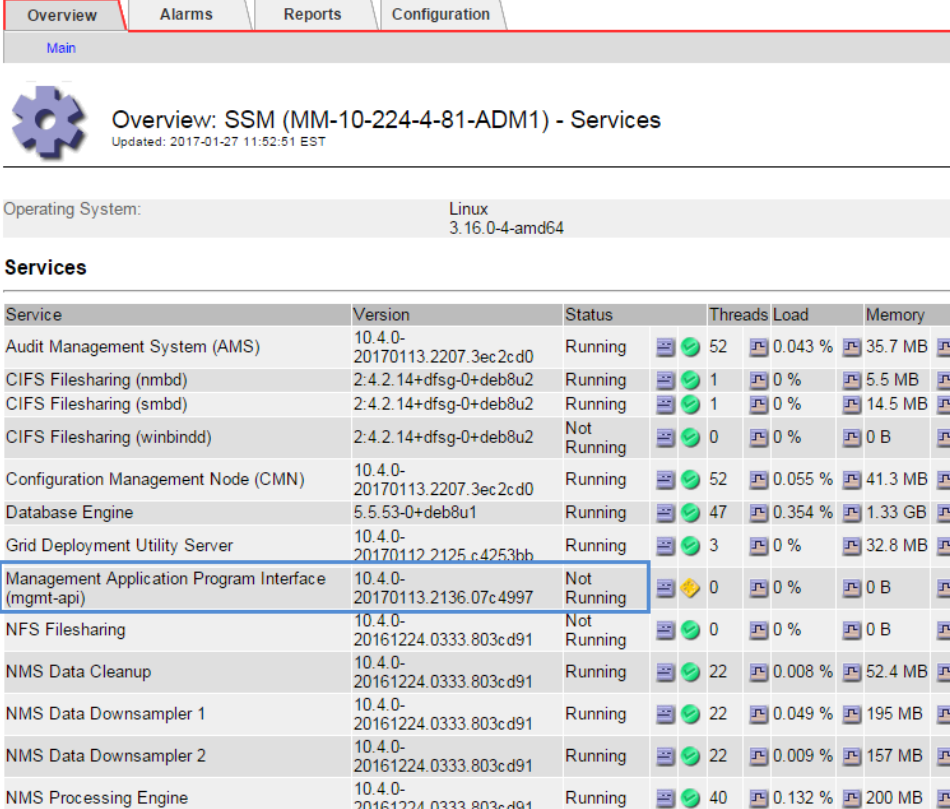
If the StorageGRID system includes multiple Admin Nodes, you can use another Admin Node to check the status of an unavailable Admin Node.

Before you begin

You must have specific access permissions.

Steps

1. From an available Admin Node, sign in to the Grid Manager using a supported browser.
2. Select **Support > Grid Topology**.
3. Select **Site > unavailable Admin Node > SSM > Services > Overview > Main**.
4. Look for services that have a status of Not Running and that might also be displayed in blue.



The screenshot shows the Grid Manager interface with the 'Overview' tab selected. The page title is 'Overview: SSM (MM-10-224-4-81-ADM1) - Services' with a sub-header 'Updated: 2017-01-27 11:52:51 EST'. Below this, the 'Operating System' is listed as 'Linux 3.16.0-4-amd64'. The 'Services' section contains a table with the following data:

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	Running	0	0 %	0 B
NMS Data Cleanup	10.4.0-20161224.0333.803cd91	Running	22	0.008 %	52.4 MB
NMS Data Downsampler 1	10.4.0-20161224.0333.803cd91	Running	22	0.049 %	195 MB
NMS Data Downsampler 2	10.4.0-20161224.0333.803cd91	Running	22	0.009 %	157 MB
NMS Processing Engine	10.4.0-20161224.0333.803cd91	Running	40	0.132 %	200 MB

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

Related information

[Administering StorageGRID](#)

Troubleshooting Storage Nodes

There are several tasks you can perform to help determine the source of Storage Node related problems.

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. Select **Support > Grid Topology**. Then select **site > Storage Node > LDR > Storage > Overview > Main**.

Overview: LDR (dc1-s4-100-116) - Storage
Updated: 2016-01-30 15:24:26 PST

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

Utilization

Total Space:	160 GB	
Total Usable Space:	155 GB	
Total Usable Space (Percent):	96.6 %	
Total Data:	5.34 GB	
Total Data (Percent):	3.335 %	

Replication

Block Reads:	3,392	
Block Writes:	3,561	
Objects Retrieved:	1,197	
Objects Committed:	675	
Objects Deleted:	673	
Delete Service State:	Enabled	

Object Stores

ID	Total	Available	Stored Data	Stored (%)	Cached Data	Cached (%)	Health
0000	53.4 GB	48 GB	5.34 GB	10.002 %	32.3 MB	0.06 %	No Errors
0001	53.4 GB	53.4 GB	727 KB	0.001 %	34 MB	0.064 %	No Errors
0002	53.4 GB	53.4 GB	862 KB	0.002 %	80.5 MB	0.151 %	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 place the Storage Node in a read-only state so that the StorageGRID system can use it for data retrieval while you prepare for a full recovery of the server.

Related information

[Recovery and maintenance](#)

Troubleshooting SAVP Total Usable Space (Percent) alarm

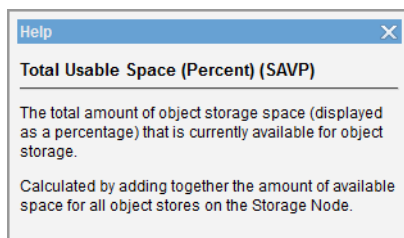
If the SAVP Total Usable Space alarm appears, you can investigate the cause.

Before you begin

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

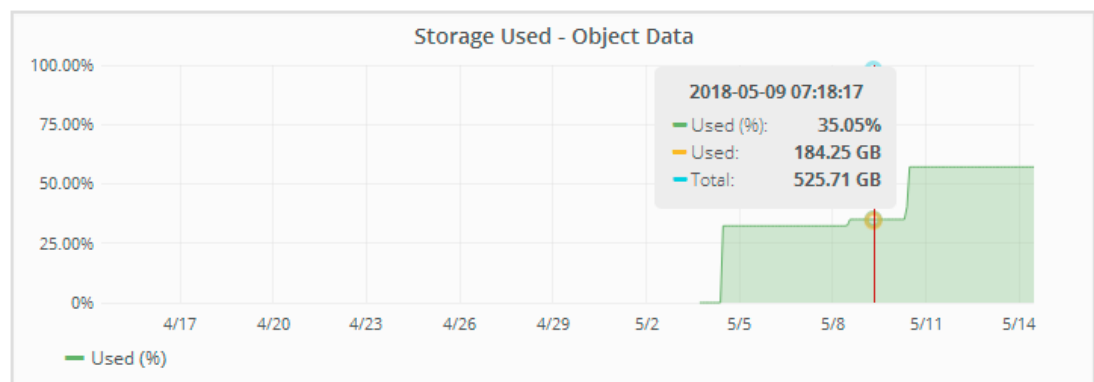
Steps

1. Select **Alarms**. Then, use the options in the Alarms section of the menu.
2. Notice the SAVP alarm.
3. Click the attribute name to display more information.



4. Determine when the Storage Node is likely to reach capacity:
 - a. Select **Nodes > Storage Node > Storage**.
 - b. Choose a range (1 hour, 1 day, 1 week, 1 month, 1 year, or Custom).
 - c. Hover over the **Storage Used - Object Data** graph. Then, slide your cursor to the right to determine how quickly storage is being used over the chosen range of time.

Note: The graph shows how much storage is being used for object data. You can use the Storage Used - Object Metadata graph to determine how much storage is being used for object metadata.



5. If an inadequate amount of space remains on this Storage Node and other Storage Nodes, add storage to the system by adding Storage Nodes in an expansion procedure.

Related references

[Alarms reference](#) on page 166

Related information

[Expanding a StorageGRID system](#)

Troubleshooting Total Events (SMTT) alarms

If a Total Events (SMTT) alarm appears on a node, it could mean that platform services messages cannot be delivered or that a Cassandra out-of-memory error has occurred.

Follow the appropriate instructions to determine the source of the alarm and how to address the issue.

Platform services message cannot be delivered

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 more information, see “Troubleshooting platform services” in the instructions for administering StorageGRID.

Steps

1. To view the alarm, select **Nodes** > *site* > *grid node* > **Events**.
2. View Last Event at the top of the table.
Event messages are also listed in `/var/local/log/bycast-err.log`.
3. Follow the guidance provided in the SMTT alarm contents to correct the issue.
4. Click **Reset event counts**.
5. Notify the tenant of the objects whose platform services messages have not been delivered.
6. Instruct the tenant to trigger the failed replication or notification by updating the object's metadata or tags.

Related concepts

[Log files](#) on page 191

Related tasks

[Resetting event counts](#) on page 121

Related information

[Administering StorageGRID](#)

Cassandra database out-of-memory error

A Total Events (SMTT) alarm is triggered when the Cassandra database has an out-of-memory error. If this error occurs, contact technical support to work through the issue.

About this task

If an out-of-memory error occurs for the Cassandra database, a heap dump is created, a Total Events (SMTT) alarm is triggered, and the Cassandra Heap Out Of Memory Errors count is incremented by one.

Steps

1. To view the event, select **Nodes > *grid node* > Events**.
2. Verify that the Cassandra Heap Out Of Memory Errors count is 1 or greater.
3. Go to `/var/local/core/`, compress the `Cassandra.hprof` file, and send it to technical support.
4. Make a backup of the `Cassandra.hprof` file, and delete it from the `/var/local/core/` directory.

This file can be as large as 24 GB, so you should remove it to free up space.

5. Once the issue is resolved, click **Reset event counts**.

Note: To reset event counts, you must belong to a group that has the Grid Topology Page Configuration permission.

Related tasks

[Resetting event counts](#) on page 121

Troubleshooting Storage Status (SSTS) alarms

The Storage Status (SSTS) alarm is triggered if a Storage Node has insufficient free space remaining for object storage.

Before you begin

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

About this task

The SSTS (Storage Status) alarm is triggered at the Notice level when the amount of free space on every volume in a Storage Node falls below the value of the Storage Volume Soft Read Only Watermark (**Configuration > Storage Options > Overview**).



Storage Options Overview

Updated: 2019-10-09 13:09:30 MDT

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark	30 GB
Storage Volume Soft Read-Only Watermark	10 GB
Storage Volume Hard Read-Only Watermark	5 GB
Metadata Reserved Space	3,000 GB

For example, suppose the Storage Volume Soft Read-Only Watermark is set to 10 GB, which is its default value. The SSTS alarm is triggered if less than 10 GB of usable space remains on each storage volume in the Storage Node. If any of the volumes has 10 GB or more of available space, the alarm is not triggered.

If an SSTS alarm has been triggered, you can follow these steps to better understand the issue.

Steps

1. Select **Alarms**. Then, in the Alarms section of the menu, select **Current Alarms**.
2. From the Service column, select the data center, node, and service that are associated with the SSTS alarm.

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

Overview

Alarms

Reports

Configuration

Main

History

Alarms: LDR (DC1-S3-101-195) - Storage

Updated: 2019-10-09 12:52:43 MDT

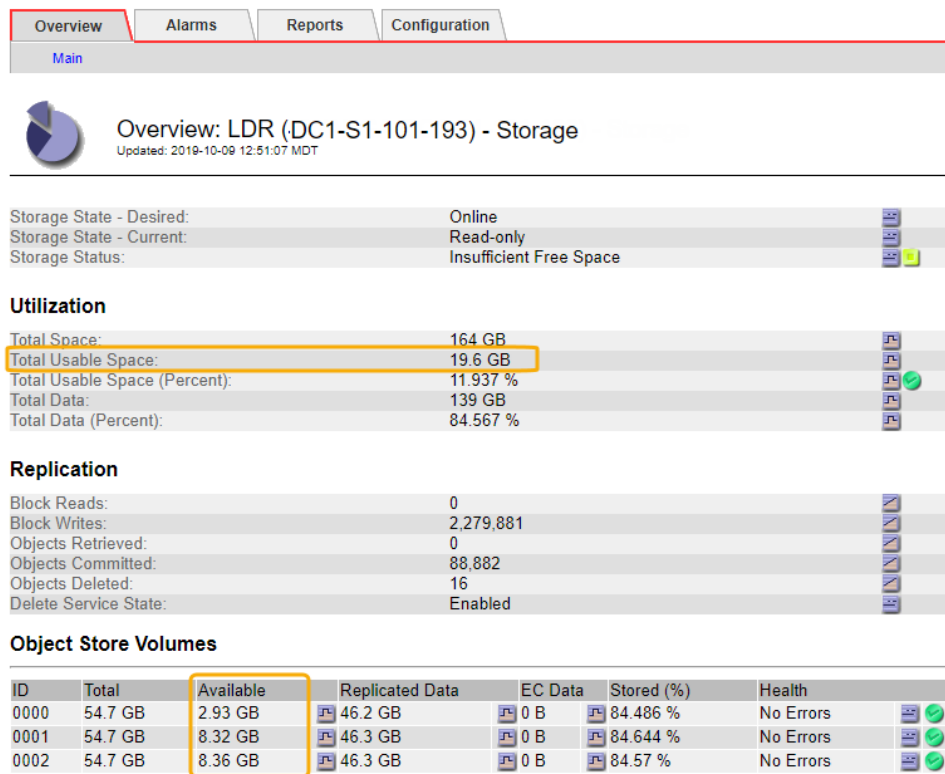
Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Notice	SSTS (Storage Status)	Insufficient Free Space	2019-10-09 12:42:51 MDT	Insufficient Free Space	Insufficient Free Space		<input type="checkbox"/>
Notice	SAVP (Total Usable Space (Percent))	Under 10 %	2019-10-09 12:43:21 MDT	7.95 %	7.95 %		<input type="checkbox"/>
Normal	SHLH (Health)						<input type="checkbox"/>

Apply Changes

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

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

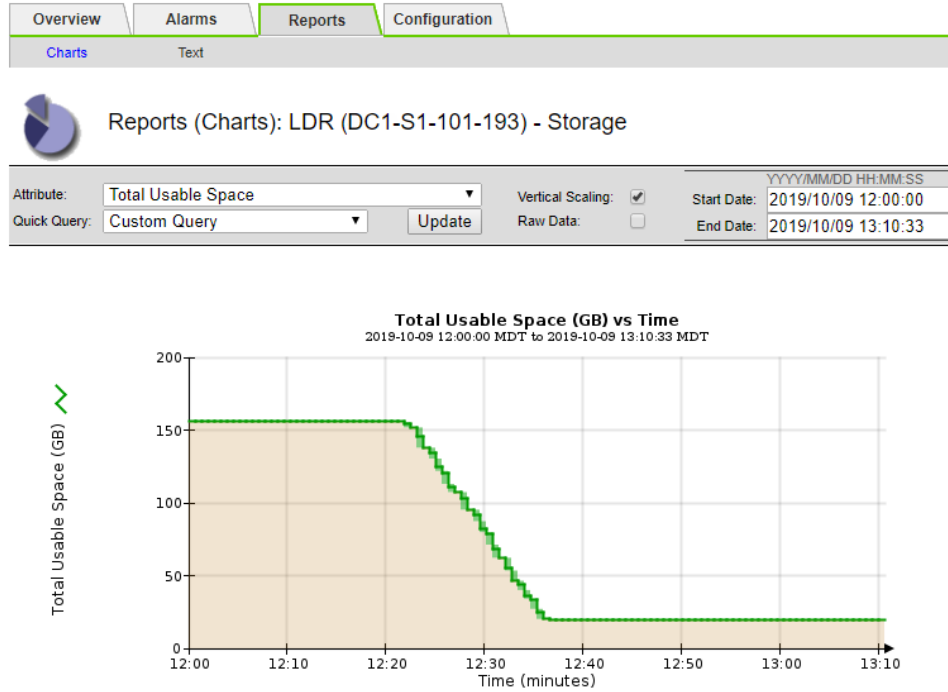
3. To determine how much usable space is actually available, select **LDR > Storage > Overview**, and find the Total Usable Space (STAS) attribute.



In this example, only 19.6 GB of the 164 GB of space on this Storage Node remains available. Note that the total value is the sum of the **Available** values for the three object store volumes. The SSTS alarm was triggered because each of the three storage volumes had less than 10 GB of available space.

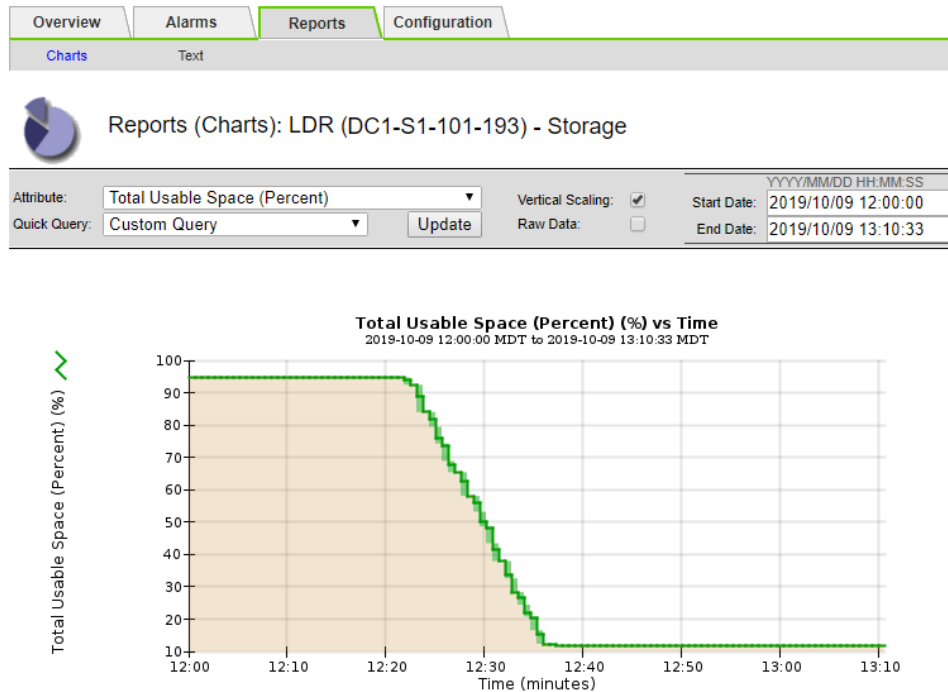
- To understand how storage has been used over time, select the **Reports** tab, and plot Total Usable Space over the last few hours.

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



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

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



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

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

Related information

[Expanding a StorageGRID system](#)

[Administering StorageGRID](#)

Troubleshooting Low object data storage alerts

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

Before you begin

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

About this task

The **Low object data storage** is triggered when the total amount of replicated and erasure coded object data on a Storage Node meets one of the conditions configured in the alert rule.

By default, a major alert is triggered when this condition evaluates as true:

```
(storagegrid_storage_utilization_data_bytes /
(storagegrid_storage_utilization_data_bytes +
storagegrid_storage_utilization_usable_space_bytes)) >=0.90
```

In this condition:

- `storagegrid_storage_utilization_data_bytes` is an estimate of the total size of replicated and erasure coded object data for a Storage Node.
- `storagegrid_storage_utilization_usable_space_bytes` is the total amount of object storage space remaining for a Storage Node.

If a major or minor **Low object data storage** alert is triggered, you should perform an expansion procedure as soon as possible.

Steps

1. Select **Alarms**. Then, in the Alerts (Preview) section of the menu, select **Alerts**.

The Alerts page appears.

2. From the table of alerts, expand the **Low object data storage** alert group, if required, and select the alert you want to view.

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

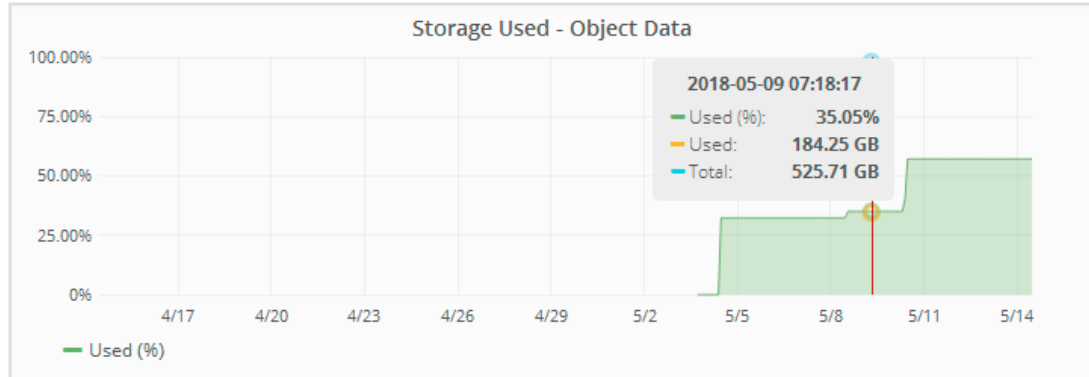
3. Review the details in the dialog box, and note the following:

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

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

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

A pop-up displays Used (%), Used, and Total capacities. The Used value is the `storagegrid_storage_utilization_data_bytes` metric.



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

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

7. As soon as possible, perform an expansion procedure to add storage capacity.

You can add storage volumes (LUNs) to existing Storage Nodes, or you can add new Storage Nodes.

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

Related tasks

[Troubleshooting Storage Status \(STS\) alarms](#) on page 144

Related information

[Expanding a StorageGRID system](#)

[Administering StorageGRID](#)

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, an alarm is triggered, as follows:

- For replicated copies, if another copy cannot be retrieved, the object is considered lost, and a LOST (Lost Objects) alarm and the **Objects lost** alert 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 LOST (Lost Objects) alarm and the **Objects lost** alert are triggered, as they are for replicated object data.

You should investigate all LOST (Lost Objects) alarms and **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 attribute under the DDS or the LDR service and clear the LOST (Lost Objects) alarm. Clearing the alarm will prevent known LOST alarm instances from masking any new LOST alarm instances.

Related tasks

[Investigating lost objects](#) on page 150

[Resetting lost and missing object counts](#) on page 155

Investigating lost objects

When the LOST (Lost Objects) alarm and the **Objects lost** alert are triggered, you must investigate immediately. Collect information about the affected objects and contact technical support.

Before you begin

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

About this task

The LOST (Lost Objects) alarm and the **Objects lost** alert indicate that StorageGRID believes that there are no copies of an object in the grid. Data might have been permanently lost.

You must investigate lost object alarms 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 Lost Objects attribute might be seen on the following pages:

- Select **Nodes**. Then, select **Storage Node > Objects**. The **Lost Objects** entry in the Object Counts table indicates the total number of objects this grid node detects as missing from the StorageGRID system. This value is the sum of the Lost Objects counters of the Data Store component within the LDR and DDS services.
- Select **Support > Grid Topology**. Then, select **site > Storage Node > LDR > Data Store > Overview > Main**.
- Select **Support > Grid Topology**. Then, select **site > Storage Node > DDS > Data Store > Overview > Main**.

This procedure shows the Lost Objects attribute on the **LDR > Data Store** page.

Steps

1. Select **Support > Grid Topology**.
2. Select **site > Storage Node > LDR > Data Store > Overview > Main**.
3. Review the Lost Objects attribute to see how many lost objects have been identified.


Overview

Alarms

Reports

Configuration

Main




Overview: LDR (DC2-S1) - Data Store





Updated: 2017-06-12 17:27:16 PDT

Lost Objects:








1





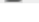
Queries

Average Query Latency:	2 ms	
Total Queries - Successful:	90,926	
Total Queries - Failed (Timed-out):	0	
Total Queries - Failed (Consistency Level Unmet):	0	

ILM Activity

ILM Implementation:	Baseline 2 Copies Policy	
ILM Version:	1.0	
Awaiting - All:	0	
Awaiting - Client:	0	
Awaiting - Background Scan:	0	
Scan Rate:	0 Objects/s	
Scan Period - Estimated:	0 us	
Awaiting - Evaluation Rate:	0 Objects/s	
Repairs Attempted:	0	

Object Transfer

Active Transfers:	0 Objects	
Transfer Rate:	0 Transfers/s	
Total Transfers:	0	

4. From an Admin Node, access the audit log to determine the identifier (CBID) of the object that triggered the LOST (Lost Objects) 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 CBID value included in the message.

```
Admin: # grep OLST audit.log
2012-01-14T11:03:27.362483 [AUDT:[CBID(UI64):0x498D8A1F681F05B3]
[UUID(CSTR):"6213A021-91FC-49C0-AF44-EC6BF377D264"]
[NOID(UI32):12088241][VOLI(UI64):2][RSLT(FC32):NONE][AVER(UI32):10]
[ATYP(FC32):OLST][ATIM(UI64):1350613602969243]
[ATID(UI64):16956755694216746320][ANID(UI32):13959984]]
```

5. Use the `ObjectByCBID` command to find the object by its identifier (CBID), and then determine if data is at risk.

- a. Telnet to localhost 1402 to access the LDR console.
- b. Enter: `/proc/OBRP/ObjectByCBID -h hexadecimal_CBID_value`

In the following example, the object with CBID 0xFE1C42ABD3CD2AC0 has a UUID, but it has no locations listed.

```
ade 21511404: / > /proc/OBRP/ObjectByCBID -h 0xFE1C42ABD3CD2AC0

{
  "OID": "00006FFD00198494009DC7E0C02DEA4CC7BCFB513B11B81B8A",
  "TYPE(Object Type)": "Data object",
  "CHND(Content handle)": "9DC7E0C0-2DEA-4CC7-BCFB-513B11B81B8A",
  "NAME": "lost/testau.dat",
  "CBID": "0xFE1C42ABD3CD2AC0",
  "PHND(Parent handle, UUID)": "402BC3FE-1BB4-11E7-8FCB-18EB00C226D9",
  "PPTH(Parent path)": "LOST",
  "META": {
    "BASE(Protocol metadata)": {
      "ISIA(Source client ip address)": "10.55.72.90",
      "PHTP(HTTP protocol handler version)": "1",
      "PAWS(S3 protocol version)": "1",
      "ACCT(S3 account ID)": "10699577065449838288",
      "*ctp(HTTP content MIME type)": "application/octet-stream"
    },
    "AWS3": {
      "USDM(User-defined metadata)": "{ \"s3b-last-modified\": \"20161117T230402Z\" }"
    }
  },
  "BYCB(System metadata)": {
    "SHSH(Supplementary Plaintext hash)": "MD5D0xC9B110581DAC712BFAE0D1D8EF36CB7E",

    "CSIZ(Plaintext object size)": "8204",
    "BSIZ(Content block size)": "8886",
    "CVER(Content block version)": "196612",
    "CFLG(Content block flags)": "256",
    "CTME(Object store begin timestamp)": "2017-04-10T20:01:58.399632",

    "CTYP(Compression algorithm type)": "NONE",
    "CHSH(Object hash)": "SHA10x7973967630676847CEB60C4C0D9384075F81A3C6",

    "MTME(Object store modified timestamp)": "2017-04-10T20:01:58.406157"
  },
  "CMSM": {
    "OWNR(ILM owner node ID)": "13895688",
    "LATM(Object last access time)": "2017-04-10T20:01:58.399632"
  }
}
```

- c. Review the output of `/proc/OBRP/ObjectByCBID`, and take the appropriate action:

Metadata	Conclusion
No object found ("ERROR": "") or an object was found with no UUID metadata	If the object is not found, the message "ERROR":"" is returned. If the object is not found, or if there is no UUID metadata, it is safe to ignore the alarm. The lack of an object, or the absence of a UUID, indicates that the object was intentionally deleted.
UUID is present Locations > 0	If there is a UUID and there are locations listed in the output, the Lost Objects alarm was a false positive. There are other object locations in the grid. You can reset the Lost Objects alarm.

Metadata	Conclusion
UUID is present Locations = 0	<p>If there is a UUID but there are no locations listed in the output, the object is potentially missing.</p> <p>If the ILM policy does not include an ILM rule with only one active content placement instruction, contact technical support. You could also try to find and restore the object yourself.</p> <p>Technical support might ask you to determine if there is a storage recovery procedure in progress. That is, has a <code>repair-data</code> command been issued on any Storage Node and is the recovery still in progress? See “Restoring object data to a storage volume” in the recovery and maintenance instructions.</p>

Related information

[Recovery and maintenance](#)

[Understanding audit messages](#)

Searching for and restoring potentially lost objects

It might be possible to find and restore objects that have triggered a Lost Objects (LOST) alarm and a **Object lost** alert, that you have identified as potentially lost.

Before you begin

- You must have the CBID 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. In some cases, you might be able to find and restore a lost replicated object if you take prompt action.

Attention: Contact technical support for assistance with this procedure.

Steps

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 hexadecimal-cbid-value audit_file_name > output_file_name
```

For example:

```
Admin: # grep 0x2E2C7E93FD5E4ED4 audit.log >
messages_about_lost_object.txt
```

- d. Use grep to extract the Object Rules Met (ORLM) audit messages for the lost object. Enter:

```
grep ORLM output_file_name
```

For example:

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

An ORLM audit message looks like this sample message.

```
2019-09-15.txt.gz:2019-09-15T13:52:54.648789 [AUDT:[CBID(UI64):
0x2E2C7E93FD5E4ED4][RULE(CSTR):"Make 2 Copies"][STAT(FC32):DONE]
[CSIZ(UI64):10][UUID(CSTR):"2A8F30F0-11D2-49E4-B649-7FA875AC018B"]
[LOCS(CSTR):"CLDI 12099341, CLDI 12018810"][RSLT(FC32):SUCS]
[AVER(UI32):10][ATIM(UI64):1568555574000][ATYP(FC32):ORLM]
[ANID(UI32):12099341][AMID(FC32):OBDI][ATID(UI64):
15559490536956943145]]
```

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

If present, the value of CLDI in LOCS is the node ID of the LDR and the volume ID 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 > 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. Look for the object on the Storage Node where the audit message indicates it might be found:

- 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 directories:

```
cd /var/local/rangedb
```

- c. Look for the object copy. Enter:

```
find . -name "hexadecimal-cbid-value"
```

For example, enter:

```
DC-SN1: # find . -name "0x2E2C7E93FD5E4ED4"
```

Note: This `find` command might take a long time (days) to complete.

If the object exists on one of the node's storage volumes, the `find` command returns the file path to the object from your current location. You can use this truncated file path to restore the object.

3. Restore the lost object:

a. Telnet to localhost 1402 to access the LDR console.

b. Enter:

```
cd /proc/CMSI
```

c. Enter:

```
Object_Found file_path_of_object
```

The full file path of the object includes `/var/local/rangedb`, but use the truncated file path you found in the previous step. For example, enter:

```
Object_Found 1/p/12/0E/0x2E2C7E93FD5E4ED4p
```

Issuing the `Object_Found` command notifies the grid of the object's location. It also triggers the active ILM policy, which will make 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 the lowest subdirectory of any `/var/local/rangedb` of the online Storage Node. Then issue the `Object_Found` command using the truncated file path to the object.

4. Reset the count of lost objects in the Grid Manager.

Related tasks

[Investigating lost objects](#) on page 150

[Resetting lost and missing object counts](#) on page 155

Related information

[Understanding audit messages](#)

Resetting lost and missing object counts

After investigating the StorageGRID system and verifying that all recorded lost objects are permanently lost or that it is a false alarm, you can reset the value of the Lost Objects attribute to zero.

Before you begin

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

About this task

You might see the Lost Objects attribute on the following pages:

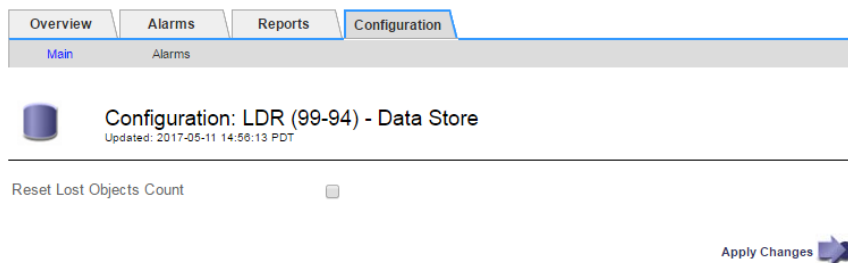
- Select **Nodes**. Then select **Storage Node > Objects**.

- Select **Support > Grid Topology**. Then select *site > Storage Node > LDR > Data Store > Overview > Main*.
- Select **Support > Grid Topology**. Then select *site > Storage Node > DDS > Data Store > Overview > Main*.

This example shows the **LDR > Data Store** page.

Steps

1. Select **Support > Grid Topology**.
2. Select *Site > Storage Node > LDR > Data Store > Configuration* for the Storage Node that has raised the (LOST) Lost Objects alarm.
3. Select **Reset Lost Objects Count**.



4. Click **Apply Changes**.
The Lost Objects attribute is reset to 0 and the LOST (Lost Objects) alarm clears. It can take a few moments for the Lost Objects attribute to reset and the alarm to clear.
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 a Lost Objects alarm is triggered, technical support might want to access the quarantined objects.
 - g. Click **Apply Changes**.

It can take a few moments for the attributes to reset after you click **Apply Changes**.

Related information

[Administering StorageGRID](#)

Troubleshooting SVST (Services: Status - Cassandra) alarm

The SVST alarm indicates that you might need to rebuild the Cassandra database for a Storage Node. Cassandra is used as the metadata store for StorageGRID.

Before you begin

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

About this task

If Cassandra is stopped for more than 15 days (for example, the Storage Node is powered off), Cassandra will not start when the node is brought back online. You must rebuild the Cassandra database for the affected DDS service.

Attention: If two or more of the Cassandra database services are down for more than 15 days, contact technical support, and do not proceed with the steps below.

Steps

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

This example shows that the SVST alarm was triggered.

Overview


Alarms

Reports

Configuration


Main

History

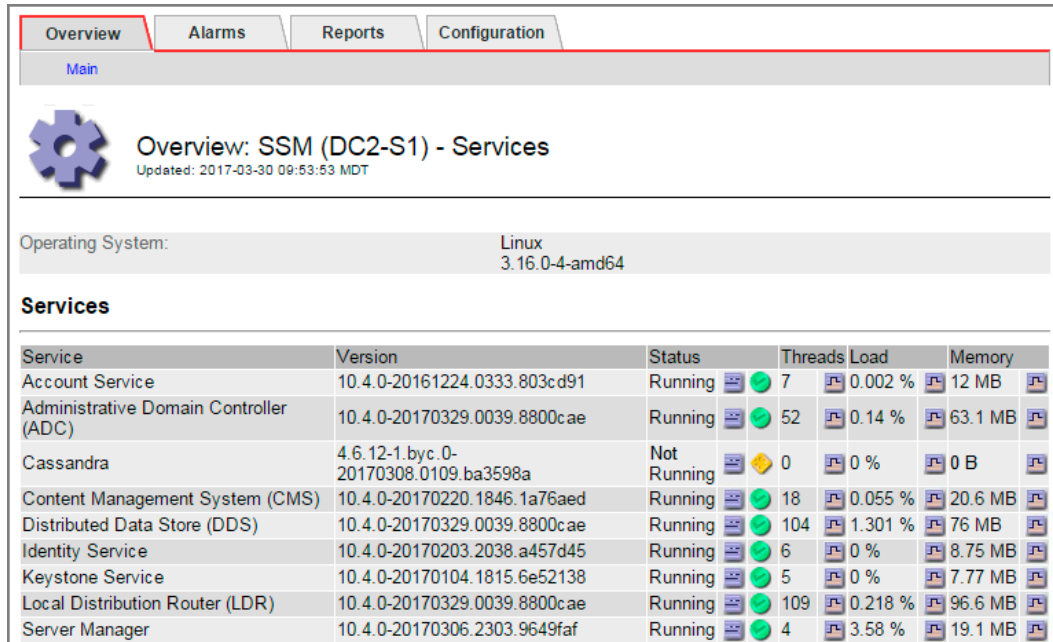


Alarms: SSM (DC1-S3) - Services

Updated: 2014-08-14 16:29:36 PDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
 Minor	SVST (Services: Status - Cassandra)	Not Running	2014-08-14 14:56:26 PDT	Not Running	Not Running		<input type="checkbox"/>

The SSM Services Main page also indicates that Cassandra is not running.



The screenshot shows the 'Overview: SSM (DC2-S1) - Services' page. The page has tabs for Overview, Alarms, Reports, and Configuration. The Overview tab is selected. Below the tabs, there is a gear icon and the text 'Overview: SSM (DC2-S1) - Services' with a timestamp 'Updated: 2017-03-30 09:53:53 MDT'. Below this, there is a section for 'Operating System' showing 'Linux 3.16.0-4-amd64'. The main section is 'Services', which contains a table with the following data:

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:

- 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 `#`.

- Enter:


```
/etc/init.d/cassandra status
```

- If Cassandra is not running, restart it:


```
/etc/init.d/cassandra restart
```

4. If Cassandra does not restart, determine how long Cassandra has been down. If Cassandra has been down for longer than 15 days, you must rebuild the Cassandra database.

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

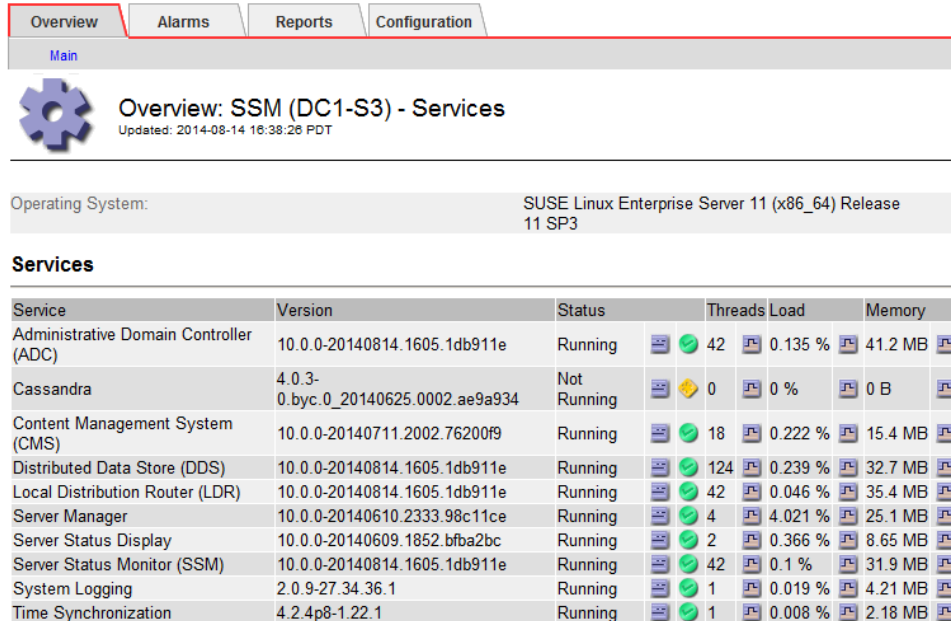
You can determine how long Cassandra has been down by charting it or by reviewing the `servermanager.log` file.

5. To chart Cassandra:

- Select **Support > Grid Topology**. Then select *site* > **Storage Node** > **SSM** > **Services** > **Reports** > **Charts**.
- Select **Attribute > Service: Status - Cassandra**.
- For **Start Date**, enter a date that is at least 16 days before the current date. For **End Date**, enter the current date.

- d. Click **Update**.
- e. If the chart shows Cassandra as being down for more than 15 days, rebuild the Cassandra database.

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



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

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

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

b. Enter:

```
cat /var/local/log/servermanager.log
```

The contents of the `servermanager.log` file are displayed.

If Cassandra has been down for longer than 15 days, the following message is displayed in the `servermanager.log` file:

```
"2014-08-14 21:01:35 +0000 | cassandra | cassandra not
started because it has been offline for longer than
its 15 day grace period - rebuild cassandra"
```

- c. Make sure the timestamp of this message is the time when you attempted restarting Cassandra as instructed in step 3 on page 158.

There can be more than one entry for Cassandra; you must locate the most recent entry.
- d. If Cassandra has been down for longer than 15 days, you must rebuild the Cassandra database.

For instructions, see “Recovering from a single Storage Node down more than 15 days” in the recovery and maintenance instructions.

- e. Contact technical support if alarms do not clear after Cassandra is rebuilt.

Related information

[Recovery and maintenance](#)

Other StorageGRID troubleshooting tips

There are several tips that can help you determine the source of a problem with your StorageGRID system.

Choices

- [User interface issues](#) on page 160
- [Time synchronization](#) on page 161
- [Network connectivity](#) on page 161
- [Linux: Node status is "orphaned"](#) on page 161
- [Linux: Enabling IPv6 support in the kernel](#) on page 162
- [Linux: Changing trigger values for CPU Load Average](#) on page 164

User interface issues

You might see issues with the Grid Manager or the Tenant Manager after upgrading to a new version of StorageGRID software.

Web interface does not respond as expected

The Grid Manager or the Tenant Manager might not respond as expected after StorageGRID software is upgraded. For example, after you use the Grid Manager to acknowledge an alarm and click **Apply Changes**, the change might not be saved.

If you experience issues with the web interface:

- Make sure you are using a supported browser.
 - Note:** Browser support has changed for StorageGRID 11.3. 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 references

[Web browser requirements](#) on page 6

Related information

[Administering StorageGRID](#)

Time synchronization

You might see issues with time synchronization in your grid.

Inaccurate NTP servers

If you encounter time synchronization problems, verify that you have specified at least four external NTP sources, each providing a Stratum 3 or better reference, and that all external NTP sources are operating normally and are accessible by your StorageGRID nodes.

Note: When specifying the external NTP source for a production-level StorageGRID installation, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

Related information

[Recovery and maintenance](#)

Network connectivity

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

Promiscuous mode

If you are deploying StorageGRID nodes on Linux VMs and you observe that the node containers are unable to communicate over the network even though the host networking is working, verify that *promiscuous mode* is enabled in the hypervisor. Promiscuous mode is disabled by default for VMware hypervisors.

Note: This information only applies to containers deployed on Linux virtual nodes; VMware virtual machine nodes do not require promiscuous mode.

Linux: Node status is “orphaned”

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

About this task

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

- Check logs for errors and messages.
- Attempt to start the node again.
- If necessary, use Docker commands to stop the existing node container.
- Restart the node.

Steps

1. Check logs for both the service daemon and the orphaned node for obvious errors or messages about exiting unexpectedly.
2. Log in to the host as root or using an account with sudo permission.
3. Attempt to start the node again by running the following command:


```
$ sudo storagegrid node start node-name
```

Example

```
$ 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).

Example

```
sudo docker stop --time 900 storagegrid-DC1-S1-172-16-1-172
```

5. Restart the node:

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

Example

```
storagegrid node start DC1-S1-172-16-1-172
```

Linux: Enabling IPv6 support in the kernel

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)

Overview
Hardware
Network
Storage
Objects
ILM
Events

Node Information ⓘ

Name	DC1-S1	
Type	Storage Node	
Software Version	11.1.0 (build 20180606.2152.b3bbe9d)	
IP Addresses	10.96.106.102 Show less ^	

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

- Select **Support > 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.

Note: These instructions apply only if you have deployed nodes on Linux hosts. They do not apply to VMware virtual nodes or to StorageGRID appliances, which have kernel IPv6 support enabled by default.

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
```

Example

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

The result should be 0.

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

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

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

Example

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

The result should be 1.

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

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

5. Exit the container:
`exit`

Example

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

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

Example

```
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
```

Linux: Changing trigger values for CPU Load Average

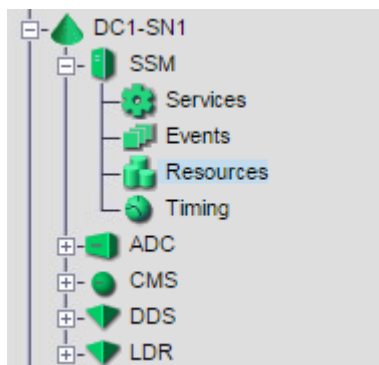
If you are using StorageGRID with Linux hosts and you are running multiple containers on a single host, you can change the trigger values for the CPU Load Average alarm to better reflect the host utilization.

About this task

To change the alarm trigger values, you must perform these steps on each node.

Steps

1. Select **Support > Grid Topology**. Then select *site* > *grid node* > **SSM > Resources**.

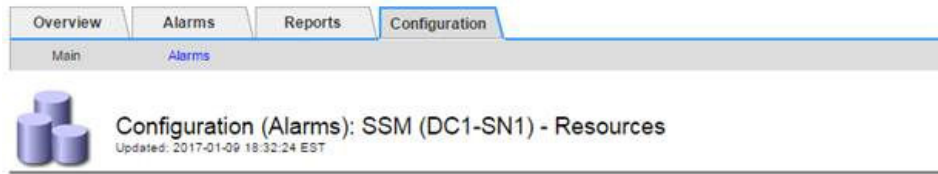



2. Scroll to the bottom of the **Processors** section and note the highest numbered processor.

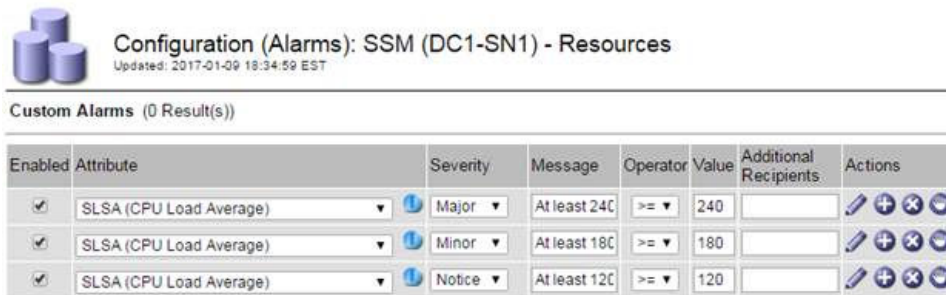
Processors

Processor Number	Vendor	Type	Cache
1	GenuineIntel	Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz	30 MiB
2	GenuineIntel	Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz	30 MiB
3	GenuineIntel	Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz	30 MiB

3. Select the **Configuration** tab, and select **Alarms**.



4. Scroll down to the SLSA (CPU Load Average) alarms.
By default, three alarm levels are defined: Major, Minor, and Notice.
5. For each of the three SLSA (CPU Load Average) alarms, click the edit icon  and uncheck the check box.
6. Click **Apply Changes** at the bottom of the page.
7. In the **Custom Alarms** section of the **Configuration** tab, click the edit icon, and define new SLSA (CPU Load Average) alarms to replace the ones you disabled in the previous steps.
When creating the new alarms, choose values that are appropriate for number of cores available.



8. When you have finished creating the custom alarms, click **Apply Changes** at the bottom of the page.

Alarms reference

The following table lists all Default StorageGRID alarms. Responses are assigned according to the severity of the alarm. As required, you can customize the alarm settings to fit your system management approach. Select **Alarms**. Then, in the Alarms section of the menu, select **Global Alarms**.

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>
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 > 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>

Code	Name	Service	Recommended action
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>
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 > 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>
AITU	Retrieve Status	BARC	<p>If the value of Retrieve Status is Target Error, check the targeted external archival storage system for errors.</p> <p>If the value of Archive Retrieve Status is Session Lost, check the targeted external archival storage system to ensure it is online and operating correctly. Check the network connection with the target.</p> <p>If the value of Archive Retrieve Status is Unknown Error, contact technical support.</p>
ALIS	Inbound Attribute Sessions	ADC	<p>If the number of inbound attribute sessions on an attribute relay grows too large, it can be an indication that the StorageGRID system has become unbalanced. Under normal conditions, attribute sessions should be evenly distributed amongst ADC services. An imbalance can lead to performance issues.</p> <p>If the problem persists, contact technical support.</p>
ALOS	Outbound Attribute Sessions	ADC	<p>The ADC service has a high number of attribute sessions, and is becoming overloaded. If this alarm is triggered, contact technical support.</p>

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>
AMQS	Audit Messages Queued	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BDDS	<p>If audit messages cannot be immediately forwarded to an audit relay or repository, the messages are stored in a disk queue. If the disk queue becomes full, outages can occur.</p> <p>To allow you to respond in time to prevent an outage, AMQS alarms are triggered when the number of messages in the disk queue reaches the following thresholds:</p> <ul style="list-style-type: none"> • Notice: More than 100,000 messages • Minor: At least 500,000 messages • Major: At least 2,000,000 messages • Critical: At least 5,000,000 messages <p>If an AMQS alarm is triggered, check the load on the system—if there have been a significant number of transactions, the alarm should resolve itself over time. In this case, you can ignore the alarm.</p> <p>If the alarm persists and increases in severity, view a chart of the queue size. If the number is steadily increasing over hours or days, the audit load has likely exceeded the audit capacity of the system. Reduce the client operation rate or decrease the number of audit messages logged by changing the audit level to Error or Off. See “Changing audit message levels” in <i>Understanding audit messages</i>.</p> <p>Understanding audit messages</p>
AOTE	Store State	BARC	<p>Only available for Archive Node's with a Target Type of Tivoli Storage Manager (TSM).</p> <p>If the value of Store State is Waiting for Target, check the external archival storage system and ensure that it is operating correctly. If the Archive Node has just been added to the StorageGRID system, ensure that the Archive Node's connection to the targeted external archival storage system is configured correctly.</p> <p>If the value of Store State is Offline, check the value of Store Status. Correct any problems before moving the Store State back to Online.</p>

Code	Name	Service	Recommended action
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>
APMS	Storage Multipath Connectivity	SSM	<p>If the multipath state alarm appears as “Degraded” (select Support > Grid Topology, then select <i>site > grid node > SSM > Events</i>), do the following:</p> <ol style="list-style-type: none"> 1. Plug in or replace the cable that does not display any indicator lights. 2. Wait one to five minutes. Do not unplug the other cable until at least five minutes after you plug in the first one. Unplugging too early can cause the root volume to become read-only, which requires that the hardware be restarted. 3. Return to the SSM > Resources page, and verify that the “Degraded” Multipath status has changed to “Nominal” in the Storage Hardware section.
ARCE	ARC State	ARC	<p>The ARC service has a state of Standby until all ARC components (Replication, Store, Retrieve, Target) have started. It then transitions to Online.</p> <p>If the value of ARC State does not transition from Standby to Online, check the status of the ARC components.</p> <p>If the value of ARC State is Offline, restart the service. 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 > Grid Topology. Then select <i>site > grid node > ARC > Retrieve > Configuration > Main</i>, select Reset Request Failure Count and click Apply Changes.</p>
ARRS	Repository Status	NMS	<p>The NMS service is unexpectedly not gathering attribute information from the StorageGRID system.</p> <p>If the problem persists, contact technical support.</p>
ARRV	Verification Failures	ARC	<p>To diagnose and correct this problem, contact technical support.</p> <p>Once the problem that triggered this alarm is addressed, reset the failures count. Select Support > Grid Topology. Then select <i>site > grid node > ARC > Retrieve > Configuration > Main</i>, select Reset Verification Failure Count and click Apply Changes.</p>
ARVF	Store Failures	ARC	<p>This alarm can occur as a result of errors with the targeted external archival storage system. Check the external archival storage system for errors, and ensure that it is operating correctly.</p> <p>Once the problem that triggered this alarm is addressed, reset the failures count. Select Support > Grid Topology. Then select <i>site > grid node > ARC > Retrieve > Configuration > Main</i>, select Reset Store Failure Count, and click Apply Changes.</p>

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

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>
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. Check the ILM Activity in the Dashboard, or select Support > Grid Topology, then select <i>site > grid node > DDS > Resources > Overview > Main</i>.</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>
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 su 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 SVST (Services: Status - Cassandra) alarm on page 157</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.
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><i>Monitoring object metadata capacity for each Storage Node on page 44</i></p> <p><i>Expanding a StorageGRID system</i></p>
CLBA	CLB Status	CLB	<p>If an alarm is triggered, select Support > Grid Topology, then select <i>site > grid node > CLB > Overview > Main</i> 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 > Grid Topology, then select <i>site</i> > <i>grid node</i> > 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>
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>
CQST	Average Query Latency	LDR, DDS	<p>This alarm is triggered when the average time required to run a query against the metadata store through the service exceeds the value set in the Grid Manager.</p> <p>To resolve this alarm, check for hardware and workload changes around the time the query latency increased. For example, hardware issues such as multiple failed disks and workload changes such as a sudden increase in ingests, can lead to an increase in query latency.</p>

Code	Name	Service	Recommended action
DNST	DNS Status	SSM	After installation completes, a DNST alarm is triggered in the SSM service. After the DNS is configured and the new server information reaches all grid nodes, the alarm is canceled.
ECCD	Corrupt Fragments Detected	LDR	<p>An alarm is triggered when the background verification process detects a corrupt erasure coded fragment. If a corrupt fragment is detected, an attempt is made to rebuild the fragment.</p> <p>Reset the Corrupt Fragments Detected and Copies Lost attributes to zero and monitor them to see if counts go up again. If counts do go up, there may be a problem with the Storage Node's underlying storage. A copy of erasure coded object data is not considered missing until such time that the number of lost or corrupt fragments breaches the erasure code's fault tolerance; therefore, it is possible to have corrupt fragment and to still be able to retrieve the object.</p> <p>If the problem persists, contact technical support.</p>
ECST	Verification Status	LDR	<p>This alarm indicates the current status of the background verification process for erasure coded object data on this Storage Node.</p> <p>A major alarm is triggered if there is an error in the background verification process.</p>
FOPN	Open File Descriptors	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BSSM, BDDS	FOPN can become large during peak activity. If it does not diminish during periods of slow activity, contact technical support.
HCCS	Currently Established Incoming Sessions	LDR	This alarm is triggered if there are 10,000 or more HTTP sessions currently active (open) on the Storage Node. If a node has too many connections, you might see connection failures. You should reduce the workload.

Code	Name	Service	Recommended action
HSTE	HTTP State	BLDR	<p>HSTE and HSTU are related to the HTTP protocol for all LDR traffic, including S3, Swift, and other internal StorageGRID traffic. An alarm indicates that one of the following situations has occurred:</p> <ul style="list-style-type: none"> • The HTTP protocol has been taken offline manually. • The Auto-Start HTTP attribute has been disabled. • The LDR service is shutting down. <p>The Auto-Start HTTP attribute is enabled by default. If this setting is changed, HTTP could remain offline after a restart.</p> <p>If necessary, wait for the LDR service to restart.</p> <p>Select Support > 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>
HSTU	HTTP Status	BLDR	
HTAS	Auto-Start HTTP	LDR	Specifies whether to start HTTP services automatically on start-up. This is a user-specified configuration option.
IRSU	Inbound Replication Status	BLDR, BARC	An alarm indicates that inbound replication has been disabled. Confirm configuration settings: Select Support > Grid Topology . Then select site > grid node > LDR > Replication > Configuration > Main .
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>

Code	Name	Service	Recommended action
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>
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>Lost and missing object data on page 149</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>Administering StorageGRID</p>
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 on page 76</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 on page 76</p>

Code	Name	Service	Recommended action
MISS	NMS Interface Engine Status	BNMS	An alarm is triggered if the NMS interface engine on the Admin Node that gathers and generates interface content is disconnected from the system. Check Server Manager to determine if the server individual application is down.
NANG	Network Auto Negotiate Setting	SSM	Check the network adapter configuration. The setting must match preferences of your network routers and switches. An incorrect setting can have a severe impact on system performance.
NDUP	Network Duplex Setting	SSM	Check the network adapter configuration. The setting must match preferences of your network routers and switches. An incorrect setting can have a severe impact on system performance.
NLNK	Network Link Detect	SSM	Check the network cable connections on the port and at the switch. Check the network router, switch, and adapter configurations. Restart the server. If the problem persists, contact technical support.
NRER	Receive Errors	SSM	These errors can clear without being manually reset. If errors do not clear, check the network hardware. Check that the adapter hardware and driver are correctly installed and configured to work with your network routers and switches. When the underlying problem is resolved, reset the counter: Select Support > Grid Topology . Then select <i>site > grid node > SSM > Resources > Configuration > Main</i> . Select Reset Receive Error Count and click Apply Changes .
NRLY	Available Audit Relays	BADC, BARC, BCLB, BCMN, BLDR, BNMS, BDDS	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. Restore connectivity to an ADC service as soon as possible. If the problem persists, contact technical support.
NSCA	NMS Status	NMS	If the value of NMS Status is DB Connectivity Error, restart the service. If the problem persists, contact technical support.

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

Code	Name	Service	Recommended action
NTLR	Repair Completion Status	DDS	<p>If a nodetool repair task for Cassandra stalls, the normal background process of checking for and repairing potential database inconsistencies cannot complete and is retried every hour.</p> <p>Check the Cassandra log at <code>/var/local/log/cassandra/system.log</code> for errors, and correct any issues that you discover. For example, the Storage Node could be isolated due to network issues.</p> <p>Contact technical support if you cannot identify or resolve the issue that prevents nodetool repair from completing.</p>
NTOF	NTP Time Offset	SSM	<p>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.</p>
NTSD	Chosen Time Source Delay	SSM	<p>These values give an indication of the reliability and stability of the time source that NTP on the local server is using as its reference.</p> <p>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.</p>
NTSJ	Chosen Time Source Jitter		
NTSO	Chosen Time Source Offset		
NTSU	NTP Status	SSM	<p>If the value of NTP Status is Not Running, contact technical support.</p>

Code	Name	Service	Recommended action
OCOR	Corrupt Objects Detected	LDR	<p>The total number of corrupt replicated objects that the most recently run background verification process has detected on the Storage Node. Any corrupt object should be investigated. More than 10 indicates a major problem.</p> <p>Note that this value is persistent: it is not updated once the corrupt objects have been restored.</p> <p>If corrupt objects are detected, change the Verification Priority to High. This speeds up verification and determining the magnitude of the problem.</p> <ol style="list-style-type: none"> 1. Select Support > Grid Topology. 2. Select <i>site</i> > <i>Storage Node</i> > LDR > Verification > Configuration > Main. 3. Select Verification Priority > High. 4. Click Apply Changes. <p>After the underlying problem is resolved, reset the counter to clear the alarm.</p> <ol style="list-style-type: none"> 1. Select Support > Grid Topology. 2. Select <i>site</i> > <i>Storage Node</i> > LDR > Verification > Configuration > Main. 3. Select Reset Corrupt Objects Count. 4. Click Apply Changes.
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>

Code	Name	Service	Recommended action
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 > Grid Topology. 2. Select <i>site > Storage Node > LDR > Verification > Configuration > Main</i>. 3. Select Delete Quarantined Objects. 4. Click Apply Changes. <p>The quarantined objects are removed, and the count is reset to zero.</p>
ORSU	Outbound Replication Status	BLDR, BARC	<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 > Grid Topology. Then select <i>site > grid node > LDR > Replication > Configuration</i>.</p> <p>An alarm is triggered if the LDR service is unavailable for replication. Select Support > Grid Topology. Then select <i>site > grid node > LDR > Storage</i>.</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.</p> <p>If this alarm is triggered, see the maintenance instructions for your appliance.</p>
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). 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 > Grid Topology. Then select site > grid node > ARC > Target > Configuration > Main, select Tivoli Storage Manager State > Online, and click Apply Changes.</p>
RDTU	Tivoli Storage Manager Status	BARC	<p>Only available for Archive Nodes with a Target Type of Tivoli Storage Manager (TSM). If the value of Tivoli Storage Manager Status is Configuration Error and the Archive Node has just been added to the StorageGRID system, ensure that the TSM middleware server is correctly configured.</p> <p>If the value of Tivoli Storage Manager Status is Connection Failure, or Connection Failure, Retrying, check the network configuration on the TSM middleware server, and the network connection between the TSM middleware server and the StorageGRID system.</p> <p>If the value of Tivoli Storage Manager Status is Authentication Failure, or Authentication Failure, Reconnecting, the StorageGRID system can connect to the TSM middleware server, but cannot authenticate the connection. Check that the TSM middleware server is configured with the correct user, password, and permissions, and restart the service.</p> <p>If the value of Tivoli Storage Manager Status is Session Failure, an established session has been lost unexpectedly. Check the network connection between the TSM middleware server and the StorageGRID system. Check the middleware server for errors.</p> <p>If the value of Tivoli Storage Manager Status is Unknown Error, contact technical support.</p>

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 > Grid Topology, then select <i>site > grid node > LDR > Replication > Configuration > Main</i>. Select Reset Inbound Replication Failure Count, and click Apply Changes.</p>
RIRQ	Inbound Replications – Queued	BLDR, BARC	<p>Alarms can occur during periods of high load or temporary network disruption. After system activity reduces, this alarm should clear. If the count for queued replications continues to increase, look for network problems and verify that the source and destination LDR and ARC services are online and available.</p>
RORQ	Outbound Replications – Queued	BLDR, BARC	<p>The outbound replication queue contains object data being copied to satisfy ILM rules and objects requested by clients.</p> <p>An alarm can occur as a result of a system overload. Wait to see if the alarm clears when system activity declines. If the alarm recurs, add capacity by adding Storage Nodes.</p>
SAVP	Total Usable Space (Percent)	LDR	<p>If usable space reaches a low threshold, options include expanding the StorageGRID system or move object data to archive through an Archive Node.</p> <p>Troubleshooting SAVP Total Usable Space (Percent) alarm on page 142</p>
SCAS	Status	CMN	<p>If the value of Status for the active grid task is Error, look up the grid task message. Select Support > Grid Topology. Then select <i>site > grid node > CMN > Grid Tasks > Overview > Main</i>. The grid task message displays information about the error (for example, “check failed on node 12130011”).</p> <p>After you have investigated and corrected the problem, restart the grid task. Select Support > Grid Topology. Then select <i>site > grid node > CMN > Grid Tasks > Configuration > Main</i>, and select Actions > Run.</p> <p>If the value of Status for a grid task being aborted is Error, retry aborting the grid task.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
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>Administering StorageGRID</p>
SCHR	Status	CMN	<p>If the value of Status for the historical grid task is Aborted, investigate the reason and run the task again if required.</p> <p>If the problem persists, contact technical support.</p>
SCSA	Storage Controller A	SSM	<p>An alarm is triggered if there is an issue with storage controller A in a StorageGRID appliance.</p> <p>If necessary, replace the component.</p>
SCSB	Storage Controller B	SSM	<p>An alarm is triggered if there is an issue with storage controller B in a StorageGRID appliance.</p> <p>If necessary, replace the component.</p> <p>Some appliance models do not have a storage controller B.</p>
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 > Grid Topology. Then select <i>site > grid node > SSM > Resources > Reports > Charts</i>.</p> <p>If the load on the system is not heavy and the problem persists, contact technical support.</p> <p>Note: If you use Linux and run multiple containers on a single host, you might want to change the trigger values for the CPU Load Average alarm to better reflect the host utilization. See Linux: Changing trigger values for CPU Load Average on page 164.</p>
SMST	Log Monitor State	SSM	<p>If the value of Log Monitor State is not Connected for a persistent period of time, contact technical support.</p>
SMTT	Total Events	SSM	<p>If the value of Total Events is greater than zero, check if there are known events (such as network failures) that can be the cause. Unless these errors have been cleared (that is, the count has been reset to 0), Total Events alarms can be triggered.</p> <p>When an issue is resolved, reset the counter to clear the alarm. Select Nodes > site > grid node > Events > Reset event counts.</p> <p>Note: To reset event counts, you must belong to a group that has 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 > 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>
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 > 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 Storage Status (SSTS) alarms on page 144</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 > 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 (/etc/init.d/<service> stop). • There is an issue with the MySQL database and Server Manager shuts down the MI service. • A grid node has been added, but not started. • During installation, a grid node has not yet connected to the Admin Node. <p>If a service is listed as Not Running, restart the service (/etc/init.d/<service> restart).</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> <p><i>Troubleshooting SVST (Services: Status - Cassandra) alarm on page 157</i></p>
TMEM	Installed Memory	SSM	<p>Nodes running with less than 24 GiB of installed memory can lead to performance problems and system instability. The amount of memory installed on the system should be increased to at least 24 GiB.</p>
TPOP	Pending Operations	ADC	<p>A queue of messages can indicate that the ADC service is overloaded. Too few ADC services can be connected to the StorageGRID system. In a large deployment, the ADC service can require adding computational resources, or the system can require additional ADC services.</p>
UMEM	Available Memory	SSM	<p>If the available RAM gets low, determine whether this is a hardware or software issue. If it is not a hardware issue, or if available memory falls below 50 MB (the default alarm threshold), contact technical support.</p>
VMFI	Entries Available	SSM	<p>This is an indication that additional storage is required. Contact technical support.</p>

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	<p>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.</p>
VPRI	Verification Priority	BLDR, BARC	<p>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.</p>
VSTU	Object Verification Status	BLDR	<p>Select Support > Grid Topology. Then select site > grid node > LDR > Storage > Overview > Main.</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>

Log files

The following sections list the logs used to capture events, diagnostic messages, and error conditions. You might be asked to collect log files and forward them to technical support to assist with troubleshooting.

Attention: These tables are for reference only. The logs are intended for advanced troubleshooting by technical support. Advanced techniques that involve reconstructing the problem history using the audit logs and the application log files are beyond the scope of this guide.

To access these logs, you can collect log files and system data (**Support > Logs**). Or, if the primary Admin Node is unavailable or unable to reach a specific node, you can access the logs for each grid node, as follows:

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

Related tasks

[Collecting log files and system data](#) on page 124

StorageGRID software logs

You can use StorageGRID logs to troubleshoot issues.

General StorageGRID logs

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

Server Manager logs

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

Logs for StorageGRID services

File name	Notes	Found on
/var/local/log/acct.errlog		Storage Nodes running the ADC service
/var/local/log/adc.errlog	Contains the Standard Error (stderr) stream of the corresponding services. There is one log file per service. These files are generally empty unless there are problems with the service.	Storage Nodes running the ADC service
/var/local/log/ams.errlog		Admin Nodes
/var/local/log/arc.errlog		Archive Nodes
/var/local/log/chunk.errlog		Storage Nodes
/var/local/log/clb.errlog		Gateway Nodes
/var/local/log/cmn.errlog		Admin Nodes
/var/local/log/cms.errlog	This log file might be present on systems that have been upgraded from an older version of StorageGRID. It contains legacy information.	Storage Nodes
/var/local/log/cts.errlog	This log file is only created if the Target Type is Cloud Tiering - Simple Storage Service (S3) .	Archive Nodes
/var/local/log/dds.errlog		Storage Nodes
/var/local/log/dmv.errlog		Storage Nodes
/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/idnt.errlog		Storage Nodes running the ADC service
/var/local/log/hagroups.log		Admin Nodes and Gateway Nodes

File name	Notes	Found on
/var/local/log/kstn.errlog		Storage Nodes running the ADC service
/var/local/log/ldr.errlog		Storage Nodes
/var/local/log/miscd/*.log	Contains logs for the MISCd service (Information Service Control Daemon), which provides an interface for querying and managing services on other nodes and for managing environmental configurations on the node such as querying the state of services running on other nodes.	All nodes
/var/local/log/nginx/*.log	Contains logs for the nginx service, which acts as an authentication and secure communication mechanism for various grid services (such as Prometheus and Dynip) to be able to talk to services on other nodes over HTTPS APIs.	All nodes
/var/local/log/nginx-gw/*.log	Contains logs for the restricted admin ports on Admin Nodes and for the Load Balancer service, which provides load balancing of S3 and Swift traffic from clients to Storage Nodes.	Admin Nodes and Gateway Nodes
/var/local/log/persistence*	Contains logs for the Persistence service, which manages files on the root disk that need to persist across a reboot.	All nodes
/var/local/log/prometheus.log	For all nodes, contains the node exporter service log and the ade-exporter metrics service log. For Admin Nodes, also contains logs for the Prometheus and Alert Manager services.	All nodes
/var/local/log/raft.log	Contains the output of the library used by the RSM service for the Raft protocol.	Storage Nodes with RSM service
/var/local/log/rms.errlog	Contains logs for the Replicated State Machine Service (RSM) service, which is used for S3 platform services.	Storage Nodes with RSM service
/var/local/log/ssm.errlog		All nodes
/var/local/log/update-snmpp-firewall.*	Contain logs related to the firewall ports being managed for SNMP.	All nodes
/var/local/log/update-utcn.log	Contains logs related to Untrusted Client Network mode on this node.	All nodes

NMS logs

File name	Notes	Found on
/var/local/log/nms.log	<ul style="list-style-type: none"> • Captures notifications from the Grid Manager and the Tenant Manager. • Captures events related to the operation of the NMS service, for example, alarm processing, email notifications, and configuration changes. • Contains XML bundle updates resulting from configuration changes made in the system. • Contains error messages related to the attribute downsampling done once a day. • Contains Java web server error messages, for example, page generation errors and HTTP Status 500 errors. 	Admin Nodes
/var/local/log/nms.errlog	<p>Contains error messages related to MySQL database upgrades.</p> <p>Contains the Standard Error (stderr) stream of the corresponding services. There is one log file per service. These files are generally empty unless there are problems with the service.</p>	

Related concepts

[About the bycast.log](#) on page 196

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

File name	Notes	Found on
/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/dsierror.log	Contains error information for TSM Client APIs.	Archive Nodes
Cassandra	/var/local/log/cassandra/system.log	Contains error 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
MySQL	/var/local/log/mysql.err /var/local/log/mysql-slow.log	Log files generated by MySQL. The file <code>mysql.err</code> captures database errors and events such as startups and shutdowns. The file <code>mysql-slow.log</code> (the slow query log) captures the SQL statements that took more than 10 seconds to execute.	Admin Nodes
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 > Grid Topology . Then select Topology > Site > Node > SSM > Events .	All nodes

Category	File name	Notes	Found on
NTP	/var/local/log/ntp.log	Log file for NTP error messages.	All nodes
	/var/lib/ntp/var/log/ntpstats/	The directory that contains NTP timing statistics. loopstats records loop filter statistics information. peerstats records peer statistics information.	
Samba	/var/local/log/samba/	The Samba log directory includes a log file for each Samba process (smb, nmb, and winbind) and every client hostname/IP.	Admin Node configured to export the audit share over CIFS

About the bycast.log

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

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

File rotation for bycast.log

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

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

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

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

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

Related tasks

[Collecting log files and system data](#) on page 124

Messages in bycast.log

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

This is an example of an ADE message:

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

ADE messages contain the following information:

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

Message severities in bycast.log

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

For example:

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

Error codes in bycast.log

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

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

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

Error code	Meaning
DONE	Done
SUNV	Service unavailable

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

Error number	Error code	Meaning
001	EPERM	Operation not permitted
002	ENOENT	No such file or directory
003	ESRCH	No such process
004	EINTR	Interrupted system call
005	EIO	I/O error
006	ENXIO	No such device or address
007	E2BIG	Argument list too long
008	ENOEXEC	Exec format error
009	EBADF	Bad file number
010	ECHILD	No child processes
011	EAGAIN	Try again
012	ENOMEM	Out of memory
013	EACCES	Permission denied
014	EFAULT	Bad address
015	ENOTBLK	Block device required
016	EBUSY	Device or resource busy
017	EEXIST	File exists
018	EXDEV	Cross-device link
019	ENODEV	No such device
020	ENOTDIR	Not a directory
021	EISDIR	Is a directory
022	EINVAL	Invalid argument
023	ENFILE	File table overflow
024	EMFILE	Too many open files
025	ENOTTY	Not a typewriter
026	ETXTBSY	Text file busy
027	EFBIG	File too large
028	ENOSPC	No space left on device
029	ESPIPE	Illegal seek
030	EROFS	Read-only file system

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

Error number	Error code	Meaning
066	EREMOTE	Object is remote
067	ENOLINK	Link has been severed
068	EADV	Advertise error
069	ESRMNT	Srmount error
070	ECOMM	Communication error on send
071	EPROTO	Protocol error
072	EMULTIHOP	Multihop attempted
073	EDOTDOT	RFS specific error
074	EBADMSG	Not a data message
075	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	.lib section in a.out corrupted
082	ELIBMAX	Attempting to link in too many shared libraries
083	ELIBEXEC	Cannot exec a shared library directly
084	EILSEQ	Illegal byte sequence
085	ERESTART	Interrupted system call should be restarted
086	ESTRPIPE	Streams pipe error
087	EUSERS	Too many users
088	ENOTSOCK	Socket operation on non-socket
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

Error number	Error code	Meaning
101	ENETUNREACH	Network is unreachable
102	ENETRESET	Network dropped connection because of reset
103	ECONNABORTED	Software caused connection abort
104	ECONNRESET	Connection reset by peer
105	ENOBUFS	No buffer space available
106	EISCONN	Transport endpoint is already connected
107	ENOTCONN	Transport endpoint is not connected
108	ESHUTDOWN	Cannot send after transport endpoint shutdown
109	ETOOMANYREFS	Too many references: cannot splice
110	ETIMEDOUT	Connection timed out
111	ECONNREFUSED	Connection refused
112	EHOSTDOWN	Host is down
113	EHOSTUNREACH	No route to host
114	EALREADY	Operation already in progress
115	EINPROGRESS	Operation now in progress
116		
117	EUCLEAN	Structure needs cleaning
118	ENOTNAM	Not a XENIX named type file
119	ENAVAIL	No XENIX semaphores available
120	EISNAM	Is a named type file
121	EREMOTEIO	Remote I/O error
122	EDQUOT	Quota exceeded
123	ENOMEDIUM	No medium found
124	EMEDIUMTYPE	Wrong medium type
125	ECANCELED	Operation Canceled
126	ENOKEY	Required key not available
127	EKEYEXPIRED	Key has expired
128	EKEYREVOKED	Key has been revoked
129	EKEYREJECTED	Key was rejected by service
130	EOWNERDEAD	For robust mutexes: Owner died
131	ENOTRECOVERABLE	For robust mutexes: State not recoverable

Copyright

Copyright © 2020 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

doccomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277