NetApp

Collecting additional StorageGRID data

StorageGRID

NetApp
October 03, 2025

Table of Contents

Collecting additional StorageGRID data

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

- Using charts and reports
- Monitoring PUT and GET performance
- Monitoring object verification operations
- Monitoring events
- Reviewing audit messages
- Collecting log files and system data
- Manually triggering an AutoSupport message
- Viewing the Grid Topology tree
- Reviewing support metrics
- Running diagnostics
- Creating custom monitoring applications

Using charts and reports

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

Types of charts and graphs

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

The Grid Manager Dashboard includes pie (doughnut) charts to summarize available storage for the grid and each site.
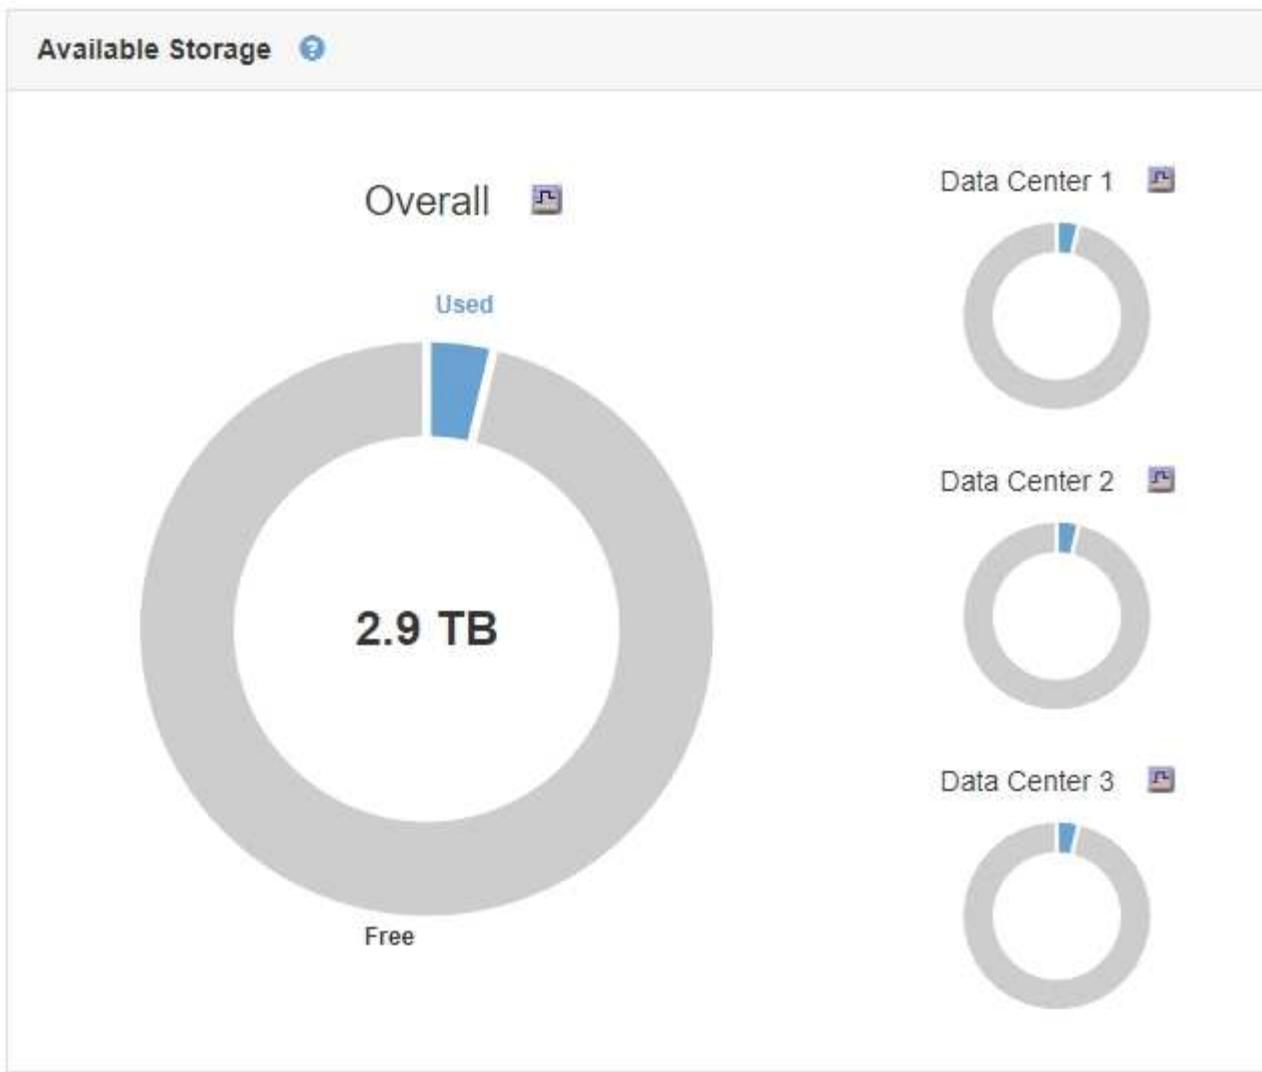
The Storage usage panel on the Tenant Manager Dashboard displays the following:

- A list of the largest buckets (S3) or containers (Swift) for the tenant
- A bar chart that represents the relative sizes of the largest buckets or containers
- The total amount of space used and, if a quota is set, the amount and percentage of space remaining

Dashboard

16 Buckets
View buckets

2 Platform services endpoints
View endpoints

0 Groups
View groups

1 User
View users

Storage usage ⦿

6.5 TB of 7.2 TB used   0.7 TB (10.1%) remaining

Bucket name	Space used	Number of objects
● Bucket-15	969.2 GB	913,425
● Bucket-04	937.2 GB	576,806
● Bucket-13	815.2 GB	957,389
● Bucket-06	812.5 GB	193,843
● Bucket-10	473.9 GB	583,245
● Bucket-03	403.2 GB	981,226
● Bucket-07	362.5 GB	420,726
● Bucket-05	294.4 GB	785,190
● 8 other buckets	1.4 TB	3,007,036

Total objects

8,418,886
objects

Tenant details

Name   Human Resources

ID   4955 9096 9804 4285 4354

View the instructions for Tenant Manager.

Go to documentation ↗

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

There are four types of graphs:

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

DC1-SG1000-ADM (Admin Node)

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

1 hour   1 day   1 week   1 month   Custom


Load Balancer Request Traffic

Load Balancer Incoming Request Rate

Average Request Duration (Non-Error)

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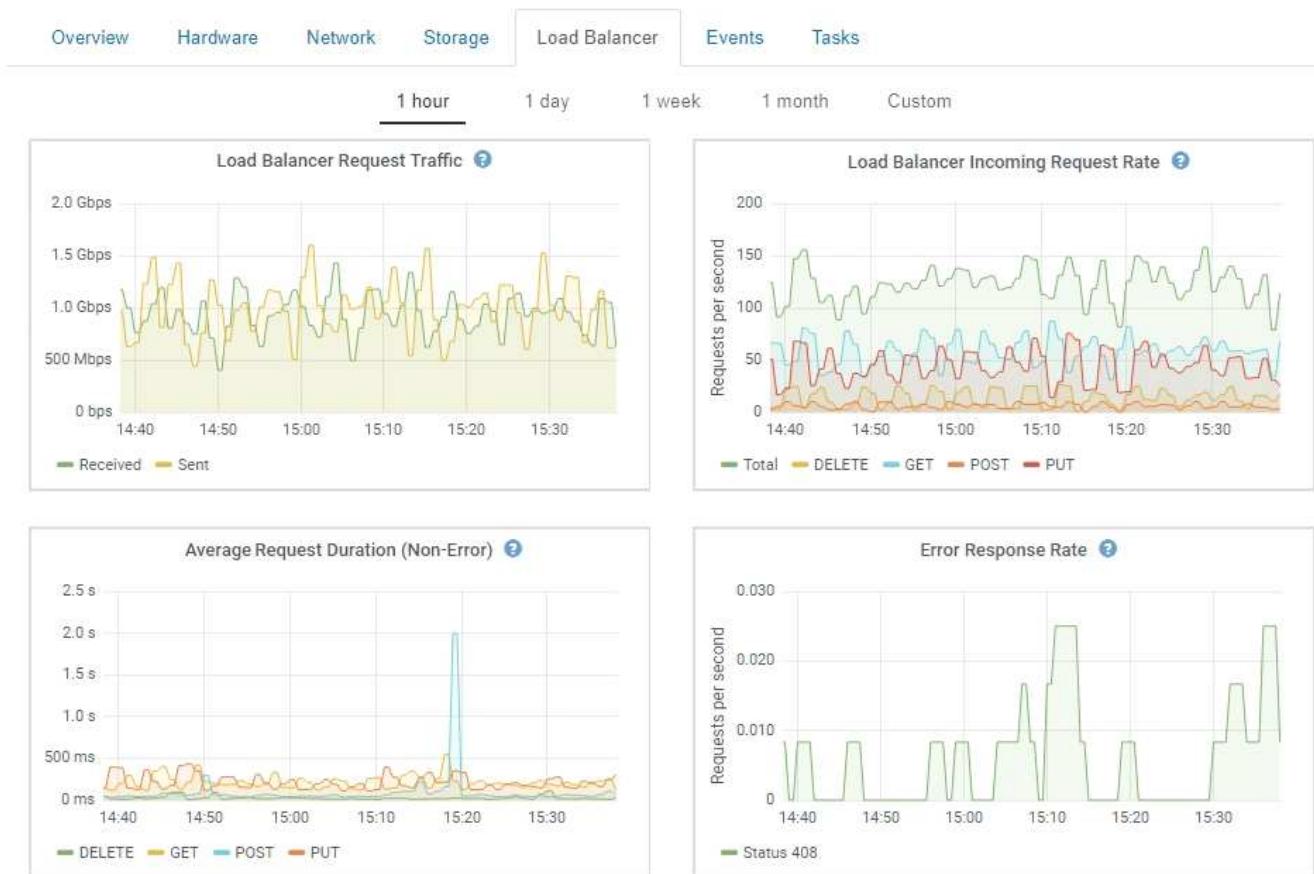
ⓘ   Grafana charts are also included on the pre-constructed dashboards available from the
     Support > **Tools** > **Metrics** page.

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

NTP Frequency Offset (ppm) vs Time
2010-07-18 16:32:15 PDT to 2010-07-18 17:32:15 PDT
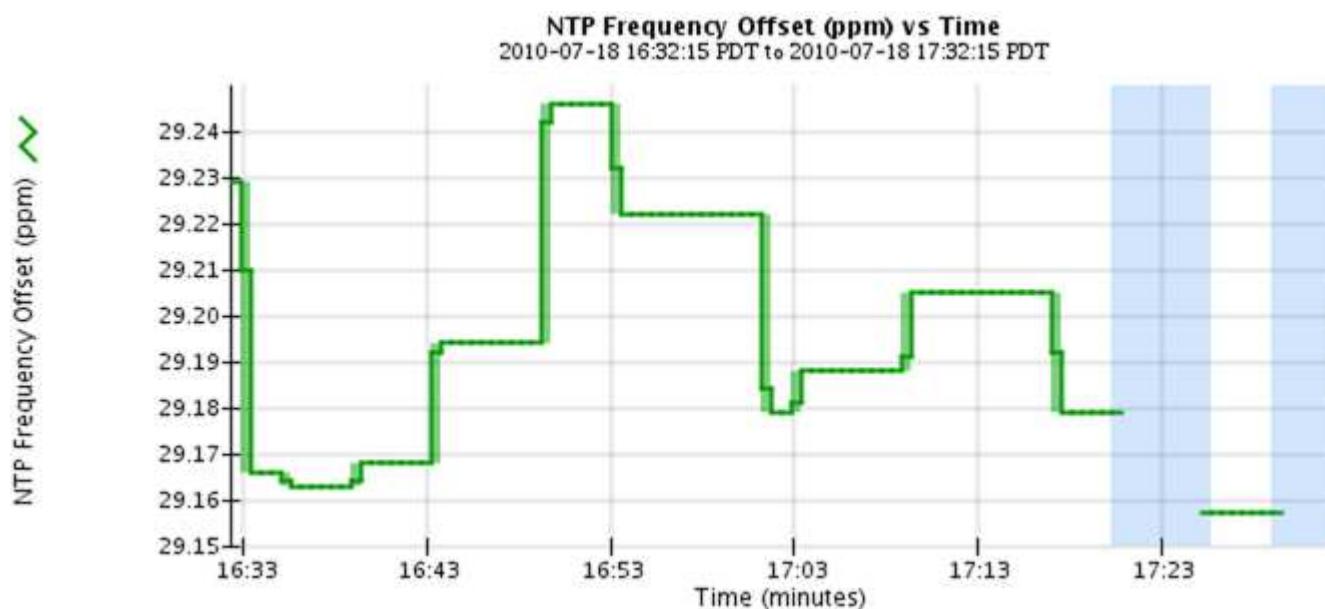
- **Area graphs**: Available from the Nodes page and from the **Support** > **Tools** > **Grid Topology** page (click the chart icon ⊓ after a data value), area graphs are used to plot volumetric attribute quantities, such as object counts or service load values. Area graphs are similar to line graphs, but include a light brown shading below the line. The changes in the value are plotted in regular data intervals (bins) over time.
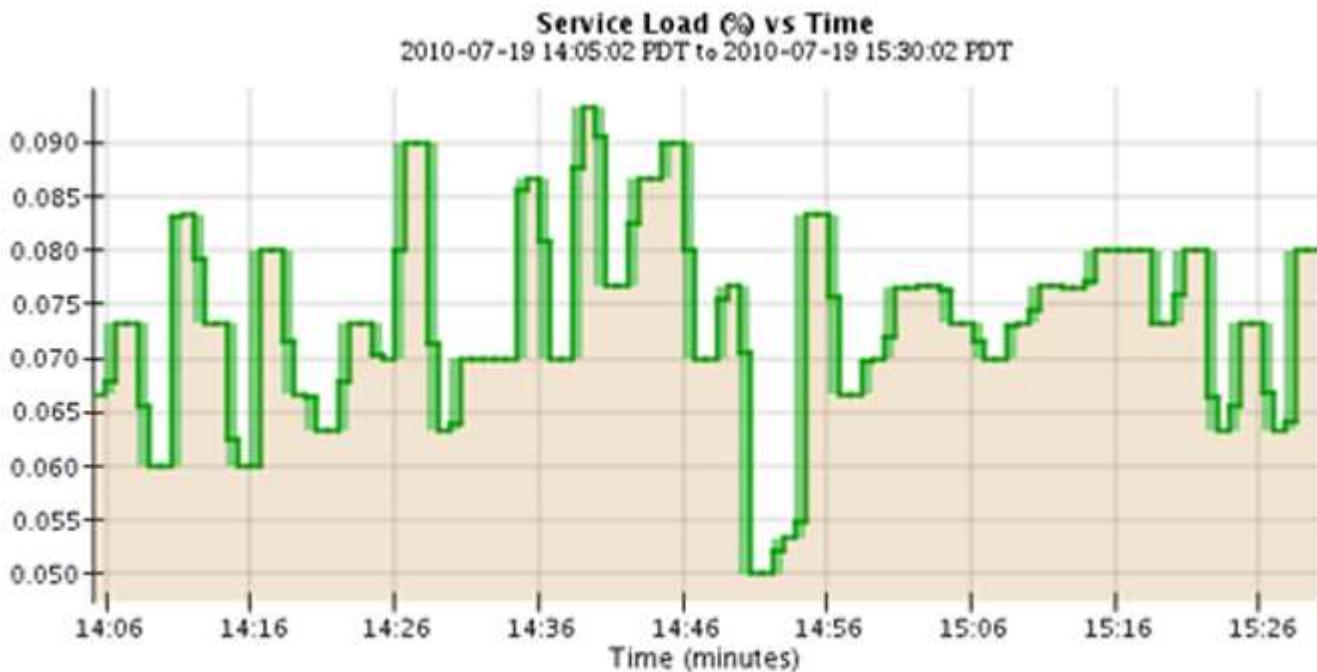

- Some graphs are denoted with a different type of chart icon ᵢₗ and have a different format:

- **State graph**: Available from the **Support** > **Tools** > **Grid Topology** page (click the chart icon ⊓ after a

5

data value), state graphs are used to plot attribute values that represent distinct states such as a service state that can be online, standby, or offline. State graphs are similar to line graphs, but the transition is discontinuous; that is, the value jumps from one state value to another.

LDR State vs Time

2004-07-09 16:40:23 to 2004-07-09 17:17:11
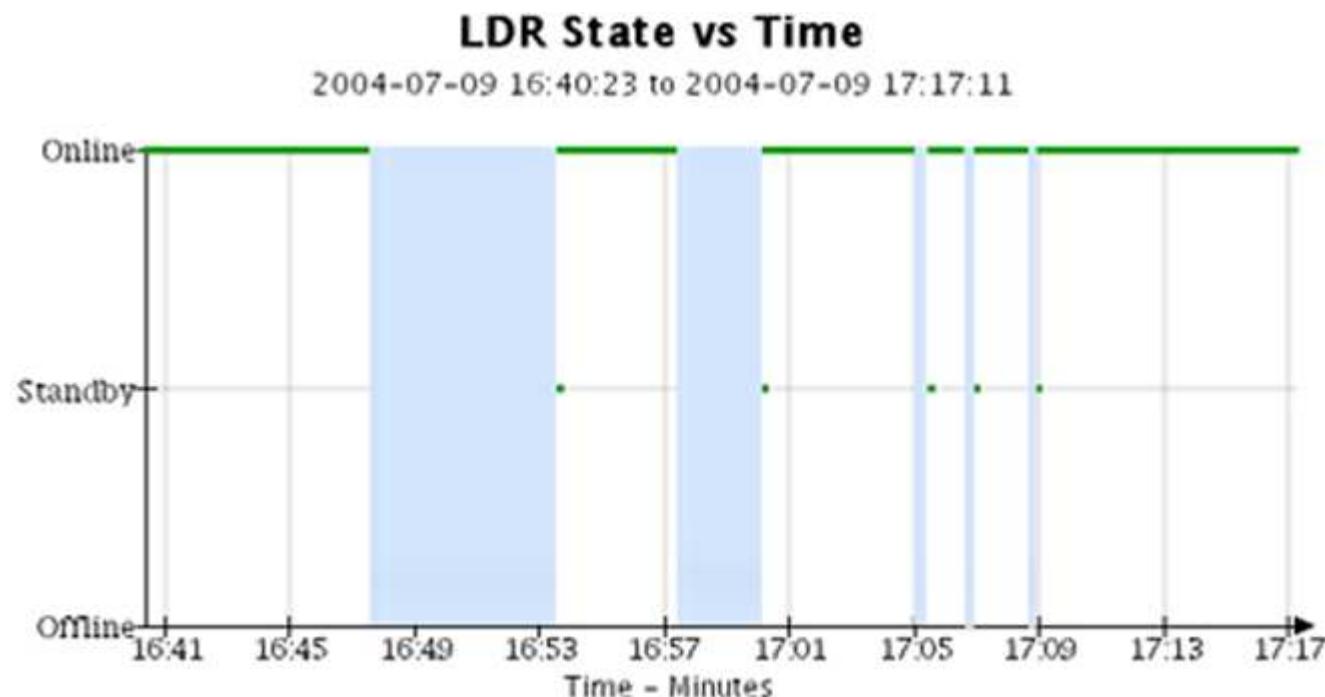
Related information

Viewing the Nodes page

Viewing the Grid Topology tree

Reviewing support metrics

Chart legend

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

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

Sample	Meaning
	Blank areas (no data plotted) indicate that the attribute values were unavailable. The background can be blue, gray, or a mixture of gray and blue, depending on the state of the service reporting the attribute.
	Light blue shading indicates that some or all of the attribute values at that time were indeterminate; the attribute was not reporting values because the service was in an unknown state.
	Gray shading indicates that some or all of the attribute values at that time were not known because the service reporting the attributes was administratively down.
	A mixture of gray and blue shading indicates that some of the attribute values at the time were indeterminate (because the service was in an unknown state), while others were not known because the service reporting the attributes was administratively down.

Displaying charts and graphs

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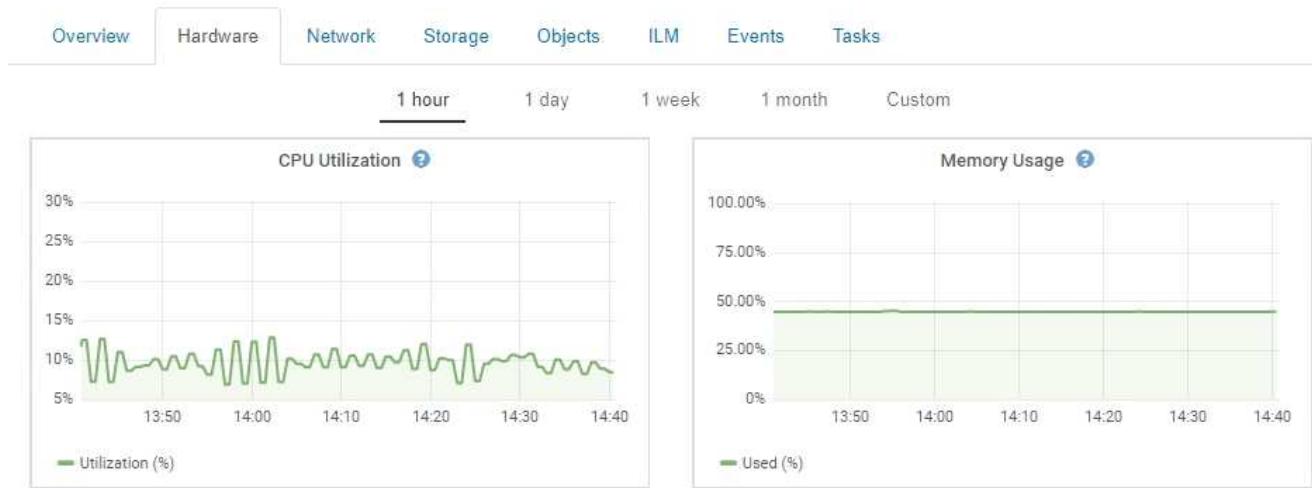
What you'll need

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

Steps

1. Select **Nodes**. Then, select a node, a site, or the entire grid.

2. Select the tab for which you want to view information.

   Some tabs include one or more Grafana charts, which are used to plot the values of Prometheus metrics over time. For example, the **Nodes** > **Hardware** tab for a node includes two Grafana charts.

DC1-S1 (Storage Node)

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

1 hour   1 day   1 week   1 month   Custom

3. Optionally, hover your cursor over the chart to see more detailed values for a particular point in time.
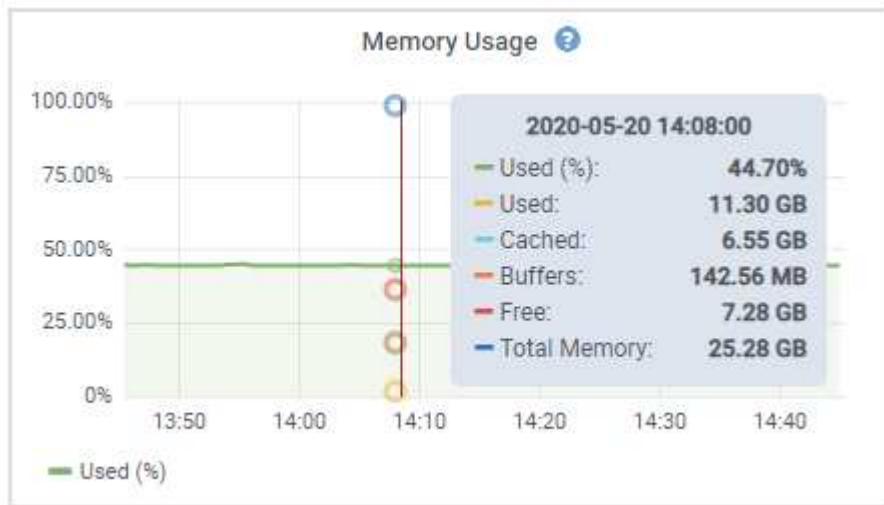
4. As required, you can often display a chart for a specific attribute or metric. From the table on the Nodes page, click the chart icon 📈 or 📊 to the right of the attribute name.

ⓘ   Charts are not available for all metrics and attributes.

Example 1: From the Objects tab for a Storage Node, you can click the chart icon 📈 to see the average latency for a metadata query over time.

Queries

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

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

Attribute:	Average Query Latency	▼		Vertical Scaling:	✔		YYYY/MM/DD HH:MM:SS
Quick Query:	Last Hour	▼	Update	Raw Data:	☐	Start Date:	2020/05/20 14:57:46
						End Date:	2020/05/20 15:57:46

Average Query Latency (Micros) vs Time
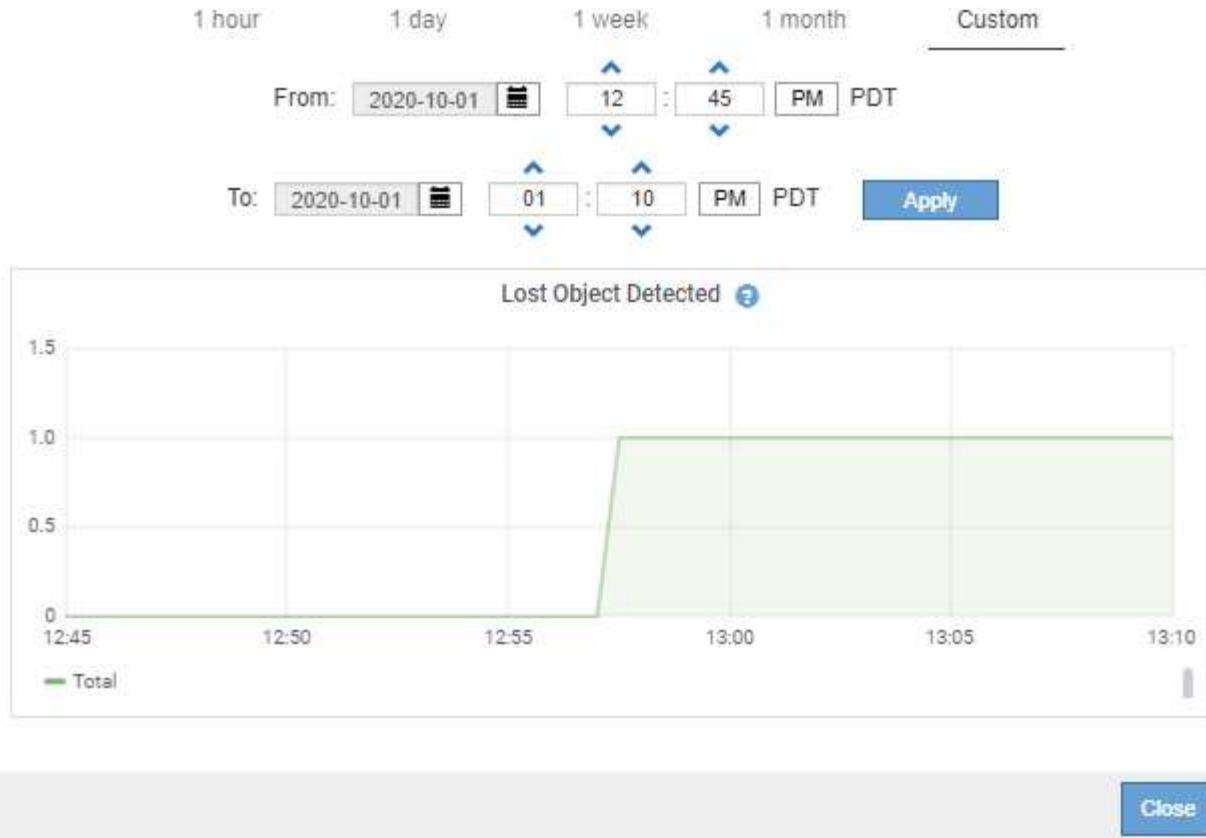#### 2020-05-20 14:57:46 MDT to 2020-05-20 15:57:46 MDT
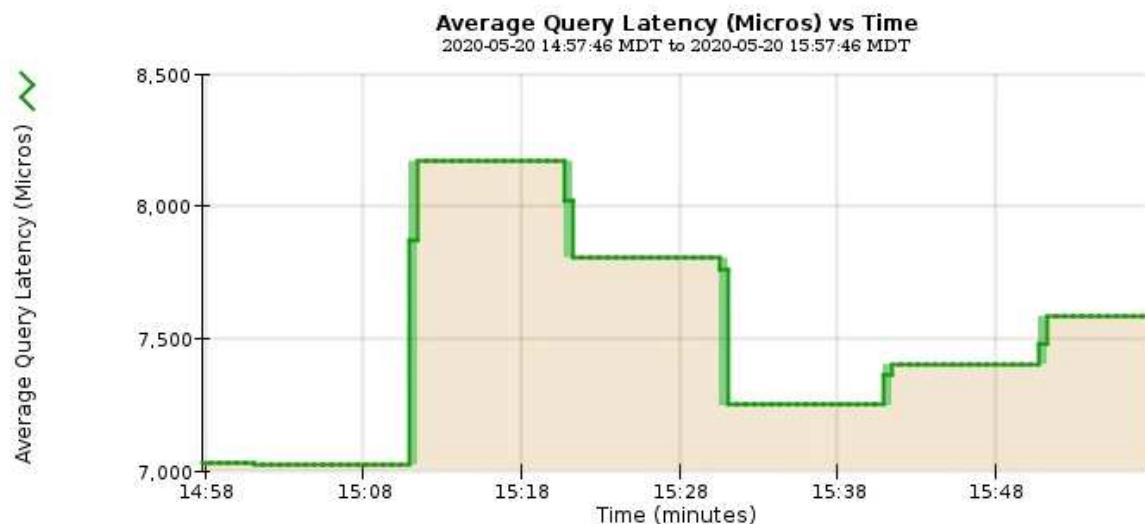
Close

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

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

5. To display charts for attributes that are not shown on the Node page, select **Support** > **Tools** > **Grid Topology**.

6. 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

7. Click the chart icon ⌐ next to the attribute.

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

Generating charts

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

What you'll need
- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.

Steps
1. Select **Support** > **Tools** > **Grid Topology**.
2. Select *grid node* > *component or service* > **Reports** > **Charts**.
3. Select the attribute to report on from the **Attribute** drop-down list.
4. To force the Y-axis to start at zero, deselect the **Vertical Scaling** check box.

5. To show values at full precision, select the **Raw Data** check box, or to round values to a maximum of three decimal places (for example, for attributes reported as percentages), deselect the **Raw Data** check box.

6. Select the time period to report on from the **Quick Query** drop-down list.

   Select the Custom Query option to select a specific time range.

   The chart appears after a few moments. Allow several minutes for tabulation of long time ranges.

7. If you selected Custom Query, customize the time period for the chart by entering the **Start Date** and **End Date**.

   Use the format $YYYY/MM/DDHH:MM:SS$ in local time. Leading zeros are required to match the format. For example, 2017/4/6 7:30:00 fails validation. The correct format is: 2017/04/06 07:30:00.

8. Click **Update**.

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

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

Types of text reports

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

Raw text reports

A raw text report displays details about the selected attribute:

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

Text Results for Services: Load - System Logging
2010-07-18 15:58:39 PDT To 2010-07-19 15:58:39 PDT

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

Aggregate text reports

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

Each entry displays the following information:

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

Text Results for Attribute Send to Relay Rate
2010-07-11 16:02:46 PDT To 2010-07-19 16:02:46 PDT

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

Generating text reports

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

What you'll need

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

About this task

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

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

Gray text indicates the service was administratively down during the time it was sampled. Blue text indicates the service was in an unknown state.

Steps

1. Select **Support** > **Tools** > **Grid Topology**.
2. Select *grid node* > *component or service* > **Reports** > **Text**.
3. Select the attribute to report on from the **Attribute** drop-down list.
4. Select the number of results per page from the **Results per Page** drop-down list.
5. To round values to a maximum of three decimal places (for example, for attributes reported as percentages), unselect the **Raw Data** check box.
6. Select the time period to report on from the **Quick Query** drop-down list.

   Select the Custom Query option to select a specific time range.

The report appears after a few moments. Allow several minutes for tabulation of long time ranges.

7. If you selected Custom Query, you need to customize the time period to report on by entering the **Start Date** and **End Date**.

   Use the format `YYYY/MM/DDHH:MM:SS` in local time. Leading zeros are required to match the format. For example, 2017/4/6 7:30:00 fails validation. The correct format is: 2017/04/06 07:30:00.

8. Click **Update**.

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

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

Exporting text reports

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

About this task

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

Steps

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

2. Create a text report.

3. Click *Export*.

Overview	Alarms	Reports	Configuration
Charts	Text		

Reports (Text): SSM (170-176) - Events
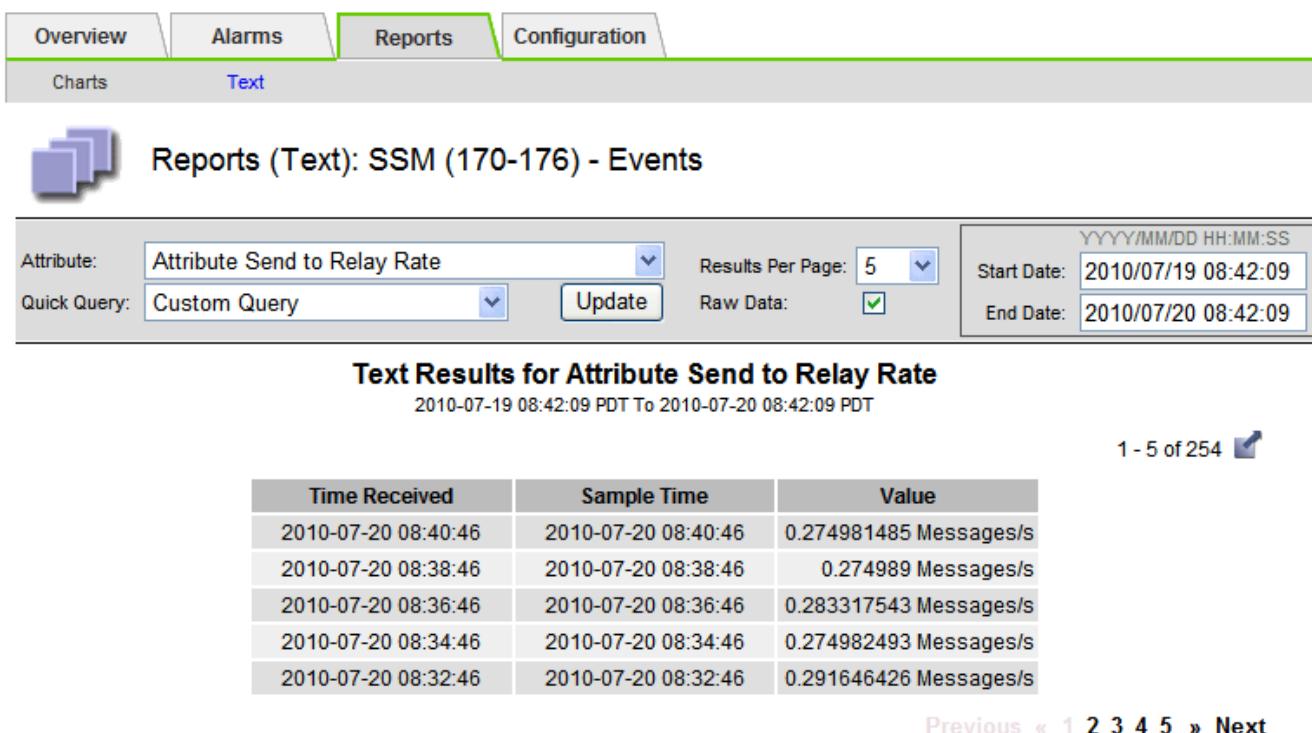
Attribute:	Attribute Send to Relay Rate	Results Per Page:	5		YYYY/MM/DD HH:MM:SS
				Start Date:	2010/07/19 08:42:09
Quick Query:	Custom Query	Update	Raw Data: ☑	End Date:	2010/07/20 08:42:09

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.

Monitoring PUT and GET performance

You can monitor the performance of certain operations, such as object store and retrieve, to help identify changes that might require further investigation.

About this task

To monitor PUT and GET performance, you can run S3 and Swift commands directly from a workstation or by using the open-source S3tester application. Using these methods allows you to assess performance independently of factors that are external to StorageGRID, such as issues with a client application or issues with an external network.

When performing tests of PUT and GET operations, use the following guidelines:

- Use object sizes comparable to the objects that you typically ingest into your grid.
- Perform operations against both local and remote sites.

Messages in the audit log indicate the total time required to run certain operations. For example, to determine the total processing time for an S3 GET request, you can review the value of the TIME attribute in the SGET audit message. You can also find the TIME attribute in the audit messages for the following operations:

- **S3**: DELETE, GET, HEAD, Metadata Updated, POST, PUT
- **Swift**: DELETE, GET, HEAD, PUT

When analyzing results, look at the average time required to satisfy a request, as well as the overall throughput that you can achieve. Repeat the same tests regularly and record the results, so that you can identify trends

that may require investigation.

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

Related information

[Review audit logs](#)

Monitoring object verification operations

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

What you'll need

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

About this task

There are two verification processes that work together to ensure data integrity:

- **Background verification** runs automatically, continuously checking the correctness of object data.

  Background verification automatically and continuously checks all Storage Nodes to determine if there are corrupt copies of replicated and erasure-coded object data. If problems are found, the StorageGRID system automatically attempts to replace the corrupt object data from copies stored elsewhere in the system. Background verification does not run on Archive Nodes or on objects in a Cloud Storage Pool.

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

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

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

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

Steps

1. Select **Nodes**.

2. Select *Storage Node* > **Objects**.

3. To check the verification results:

   - To check replicated object data verification, look at the attributes in the Verification section.

Verification

Status	No Errors
Rate Setting	Adaptive
Percent Complete	0.00%
Average Stat Time	0.00 microseconds
Objects Verified	0
Object Verification Rate	0.00 objects / second
Data Verified	0 bytes
Data Verification Rate	0.00 bytes / second
Missing Objects	0
Corrupt Objects	0
Corrupt Objects Unidentified	0
Quarantined Objects	0

(i) 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

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

Related information

Verifying object integrity

Monitoring events

You can monitor events that are detected by a grid node, including custom events that you have created to track events that are logged to the syslog server. The Last Event message shown in the Grid Manager provides more information about the most recent

event.

Event messages are also listed in the `/var/local/log/bycast-err.log` log file.

The SMTT (Total events) alarm can be repeatedly triggered by issues such as network problems, power outages or upgrades. This section has information on investigating events so that you can better understand why these alarms have occurred. If an event occurred because of a known issue, it is safe to reset the event counters.

Reviewing events from the Nodes page

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

1. Select **Nodes**.
2. Select *grid node* > **Events**.
3. At the top of the page, determine if an event is shown for **Last Event**, which describes the last event detected by the grid node.

   The event is relayed verbatim from the grid node and includes any log messages with a severity level of ERROR or CRITICAL.

4. Review the table to see if the Count for any event or error is not zero.
5. After resolving issues, click **Reset event counts** to return the counts to zero.

Reviewing events from the Grid Topology page

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

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

Related information

[Resetting event counts](#)

[Log files reference](#)

Reviewing previous events

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

1. Select **Support** > **Tools** > **Grid Topology**.
2. Select *site* > *grid node* > **SSM** > **Events** > **Reports**.
3. Select **Text**.

   The **Last Event** attribute is not shown in the Charts view.

4. Change **Attribute** to **Last Event**.
5. Optionally, select a time period for **Quick Query**.
6. Click **Update**.

Related information

Using charts and reports

Resetting event counts

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

What you'll need

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

Steps

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

Events ❓

Last Event   No Events

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

Reset event counts 🔗

Creating custom syslog events

Custom events allow you to track all kernel, daemon, error and critical level user events logged to the syslog server. A custom event can be useful for monitoring the occurrence of system log messages (and thus network security events and hardware faults).

About this task

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

- After a custom event is created, every occurrence of it is monitored. You can view a cumulative Count value for all custom events on the **Nodes** > *grid node* > **Events** page.

- To create a custom event based on keywords in the `/var/log/messages` or `/var/log/syslog` files, the logs in those files must be:

  ◦ Generated by the kernel

  ◦ Generated by daemon or user program at the error or critical level

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

Steps

1. Select **Configuration** > **Monitoring** > **Events**.

2. Click **Edit** ✏ (or **Insert** ➕ if this is not the first event).

3. Enter a custom event string, for example, shutdown

4. Click **Apply Changes**.

5. Select **Nodes**. Then, select *grid node* > **Events**.

6. Locate the entry for Custom Events in the Events table, and monitor the value for **Count**.

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

Overview	Hardware	Network	Storage	Events

Events ❓

Last Event   No Events

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

Reset event counts 🗗

Resetting the count of custom events to zero

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

About this task

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

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

4. Click **Apply Changes**.

Reviewing audit messages

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

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

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

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

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

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

client access, see the instructions for administering StorageGRID.

Related information

[Review audit logs](#)

[Administer StorageGRID](#)

Collecting log files and system data

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

What you'll need

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

About this taak

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

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

Steps

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

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

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

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

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

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

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

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

6. Click **Collect Logs**.

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

Logs

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

> Log collection is in progress.

Last Collected

Log Start Time	2017-05-17 05:01:00 PDT
Log End Time	2017-05-18 09:01:00 PDT
Notes	Issues began approximately 7am on the 17th, then multiple alarms propagated throughout the grid.

23%   Collecting logs: 10 of 13 nodes remaining

[Download] [Delete]

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

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

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

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

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

After you finish

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

Optionally, you can click **Delete** to remove the log file archive package and free up disk space. The current log file archive package is automatically removed the next time you collect log files.

Manually triggering an AutoSupport message

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

What you'll need

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

Steps

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

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

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

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

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

Related information

Configuring email server settings for alarms (legacy system)

Viewing the Grid Topology tree

The Grid Topology tree provides access to detailed information about StorageGRID system elements, including sites, grid nodes, services, and components. In most cases, you only need to access the Grid Topology tree when instructed in the documentation or when working with technical support.

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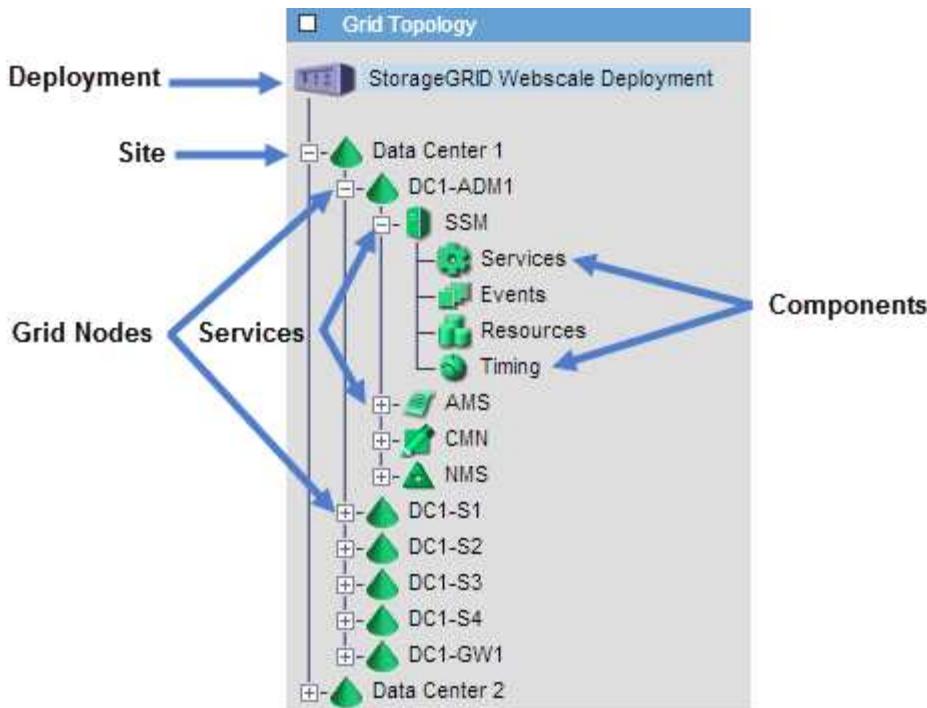
To expand or collapse the Grid Topology tree, click ⊞ or ⊟ at the site, node, or service level. To expand or collapse all items in the entire site or in each node, hold down the **<Ctrl>** key and click.

Reviewing support metrics

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

What you'll need

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

About this task

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

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

Steps

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

   The Metrics page appears.

Metrics

Access charts and metrics to help troubleshoot issues.

> 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://   /metrics/graph

Grafana

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

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

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

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

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

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

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

   The Grafana interface for the link you selected appears.
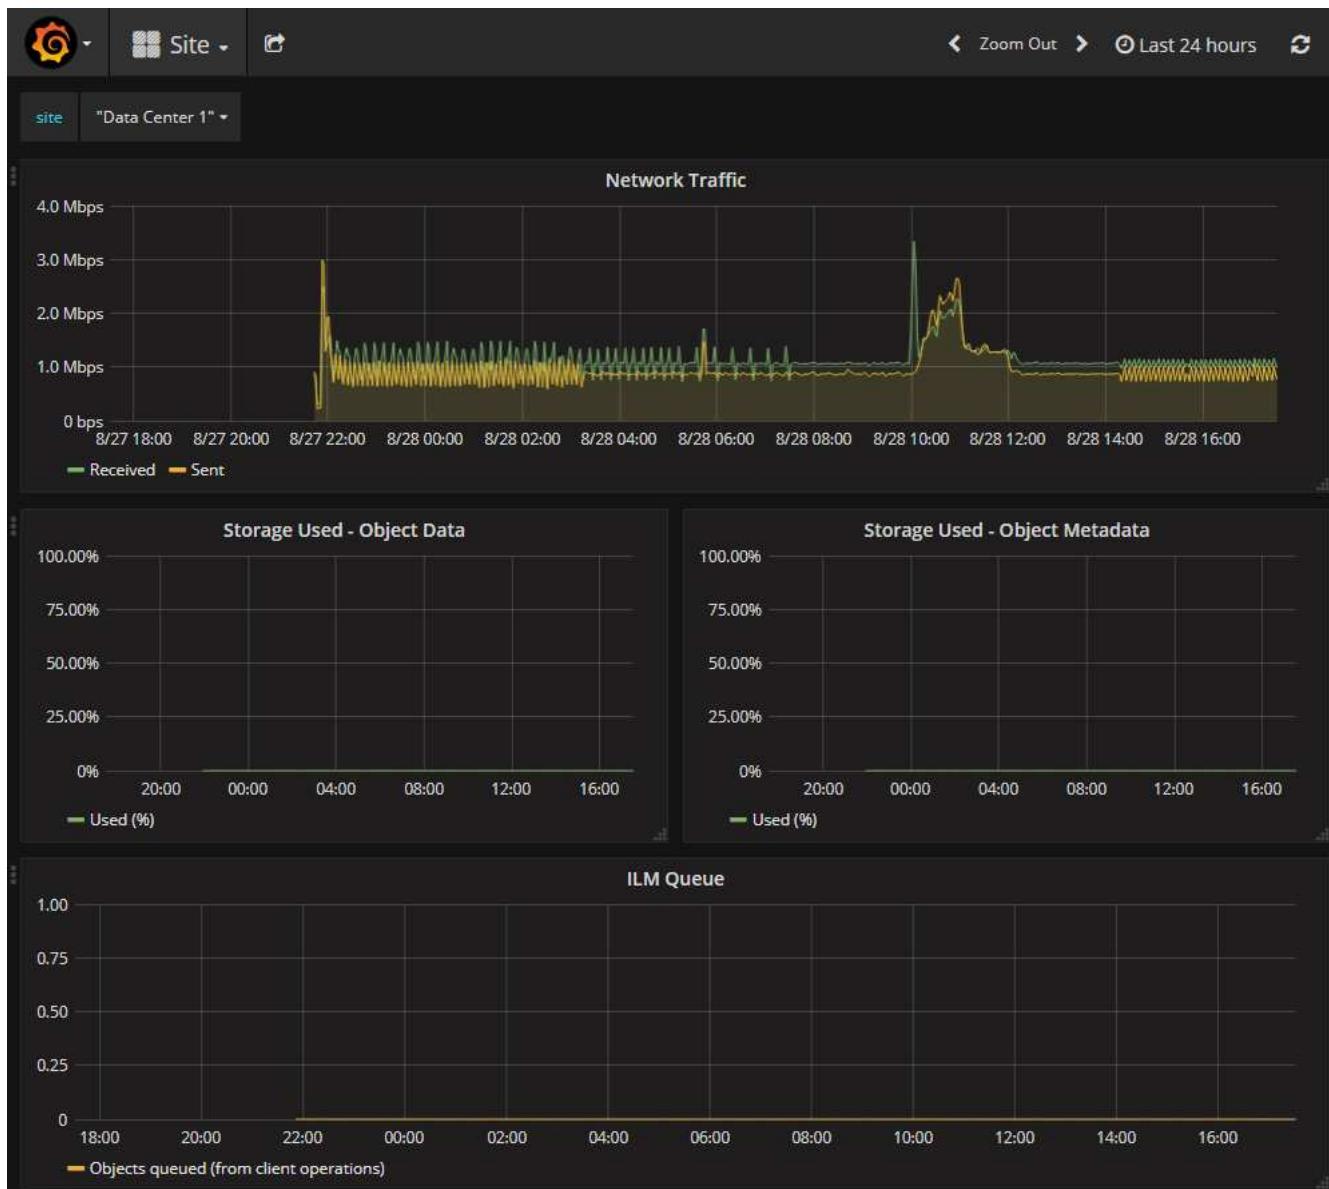
Related information

[Commonly used Prometheus metrics](#)

Running diagnostics

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

What you'll need

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

About this task

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

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

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

Steps

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

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

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

   Details about the diagnostic and its current results appear. The following details are listed:

   - **Status**: The current status of this diagnostic: Normal, Attention, or Caution.
   - **Prometheus query**: If used for the diagnostic, the Prometheus expression that was used to generate the status values. (A Prometheus expression is not used for all diagnostics.)
   - **Thresholds**: If available for the diagnostic, the system-defined thresholds for each abnormal diagnostic status. (Threshold values are not used for all diagnostics.)

     > ℹ️ You cannot change these thresholds.

   - **Status values**: A table showing the status and the value of the diagnostic throughout the StorageGRID system. In this example, the current CPU utilization for every node in a StorageGRID system is shown. All node values are below the Attention and Caution thresholds, so the overall status of the diagnostic is Normal.

CPU utilization

Checks the current CPU utilization on each node.

To view charts of CPU utilization and other per-node metrics, access the **Node** Grafana dashboard.

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

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

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

   This link is not displayed for all diagnostics.

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

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

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

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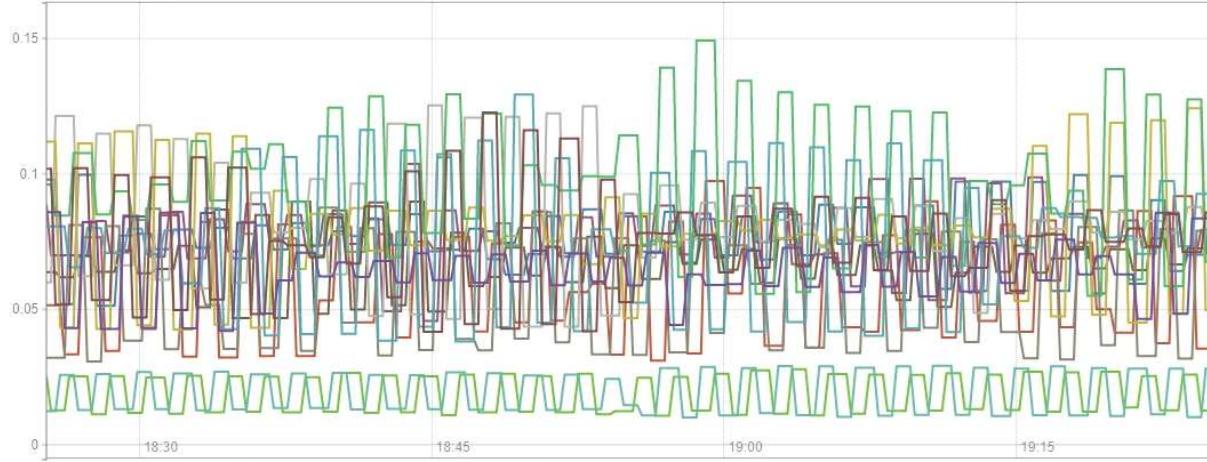
Related information

[Reviewing support metrics](#)

[Commonly used Prometheus metrics](#)

Creating custom monitoring applications

You can build custom monitoring applications and dashboards using the StorageGRID metrics available from the Grid Management API.

If you want to monitor metrics that are not displayed on an existing page of the Grid Manager, or if you want to create custom dashboards for StorageGRID, you can use the Grid Management API to query StorageGRID metrics.

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

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

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

Related information

Administer StorageGRID