



Collect additional StorageGRID data

StorageGRID 11.8

NetApp
March 19, 2024

Table of Contents

- Collect additional StorageGRID data 1
 - Use charts and graphs 1
 - Use text reports 12
 - Monitor PUT and GET performance 16
 - Monitor object verification operations 17
 - Monitor events 19
 - Review audit messages 23
- Collect log files and system data 24
- Manually trigger an AutoSupport package 25
- View the Grid Topology tree 26
- Review support metrics 27
- Run diagnostics 29
- Create custom monitoring applications 33

Collect additional StorageGRID data

Use charts and graphs

You can use charts and reports to monitor the state of the StorageGRID system and troubleshoot problems.




The Grid Manager is updated with each release and might not match the example screenshots on this page.

Types of charts

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

The Grid Manager dashboard includes cards that summarize available storage for the grid and each site.



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

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

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: Tenant02
ID: 3341 1240 0546 8283 2208

- Platform services enabled
- Can use own identity source
- S3 Select enabled

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

There are four types of graphs:

- **Grafana charts:** Shown on the Nodes page, Grafana charts are used to plot the values of Prometheus metrics over time. For example, the **NODES > Network** tab for a Storage Node includes a Grafana chart for network traffic.

DC1-S2 (Storage Node)

Overview

Hardware

Network

Storage

Objects

ILM

Tasks


1 hour

1 day

1 week

1 month

Custom



Network interfaces

Name	Hardware address	Speed	Duplex	Auto-negotiation	Link status
eth0	00:50:56:A7:E8:1D	10 Gigabit	Full	Off	Up

Network communication

Receive


Interface	Data	Packets	Errors	Dropped	Frame overruns	Frames
eth0	3.04 GB	20,403,428	0	24,899	0	0


Transmit


Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	3.65 GB	19,061,947	0	0	0	0




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

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



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



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


1 hour 1 day 1 week 1 month Custom


From: 2020-10-01 [calendar icon] 12 : 45 PM PDT

To: 2020-10-01 [calendar icon] 01 : 10 PM PDT Apply



Close

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









Related information
[View the Nodes page](#)

[View the Grid Topology tree](#)

[Review support metrics](#)

Chart legend

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

Example	Meaning
	Reported attribute values are plotted using dark green lines.
	Light green shading around dark green lines indicates that the actual values in that time range vary and have been "binned" for faster plotting. The dark line represents the weighted average. The range in light green indicates the maximum and minimum values within the bin. Light brown shading is used for area graphs to indicate volumetric data.
	Blank areas (no data plotted) indicate that the attribute values were unavailable. The background can be blue, gray, or a mixture of gray and blue, depending on the state of the service reporting the attribute.
	Light blue shading indicates that some or all of the attribute values at that time were indeterminate; the attribute was not reporting values because the service was in an unknown state.
	Gray shading indicates that some or all of the attribute values at that time were not known because the service reporting the attributes was administratively down.
	A mixture of gray and blue shading indicates that some of the attribute values at the time were indeterminate (because the service was in an unknown state), while others were not known because the service reporting the attributes was administratively down.

Display charts and graphs

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


Before you begin

You must be signed in to the Grid Manager using a [supported web browser](#).


Steps


1. Select **NODES**. Then, select a node, a site, or the entire grid.
2. Select the tab for which you want to view information.

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




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



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




Charts aren't available for all metrics and attributes.


Example 1: From the Objects tab for a Storage Node, you can select the chart icon  to see the total number of successful metadata store queries for the Storage Node.




Attribute: Total Queries - Successful
Quick Query: Last Hour
Vertical Scaling:
Raw Data:
Start Date: 2021/09/24 13:24:01
End Date: 2021/09/24 14:24:01




Close



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

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



1 hour 1 day 1 week 1 month Custom

From: 2020-10-01 [calendar icon] 12 : 45 PM PDT



To: 2020-10-01 [calendar icon] 01 : 10 PM PDT [Apply](#)





[Close](#)

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

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. Select the chart icon  next to the attribute.

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

Generate charts

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

Before you begin

- You must be signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **grid node > component or service > Reports > Charts**.
3. Select the attribute to report on from the **Attribute** drop-down list.
4. To force the Y-axis to start at zero, clear the **Vertical Scaling** checkbox.
5. To show values at full precision, select the **Raw Data** checkbox, or to round values to a maximum of three

decimal places (for example, for attributes reported as percentages), clear the **Raw Data** checkbox.

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

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

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

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

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

8. Select **Update**.

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

Use text reports

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

Raw text reports

A raw text report displays details about the selected attribute:

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

Text Results for Services: Load - System Logging

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

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

Aggregate text reports

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

Each entry displays the following information:

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

Text Results for Attribute Send to Relay Rate

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

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

Generate text reports

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

Before you begin

- You must be signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).

About this task

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

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

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

Steps

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

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

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

- 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.

- Click **Update**.

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


Export text reports

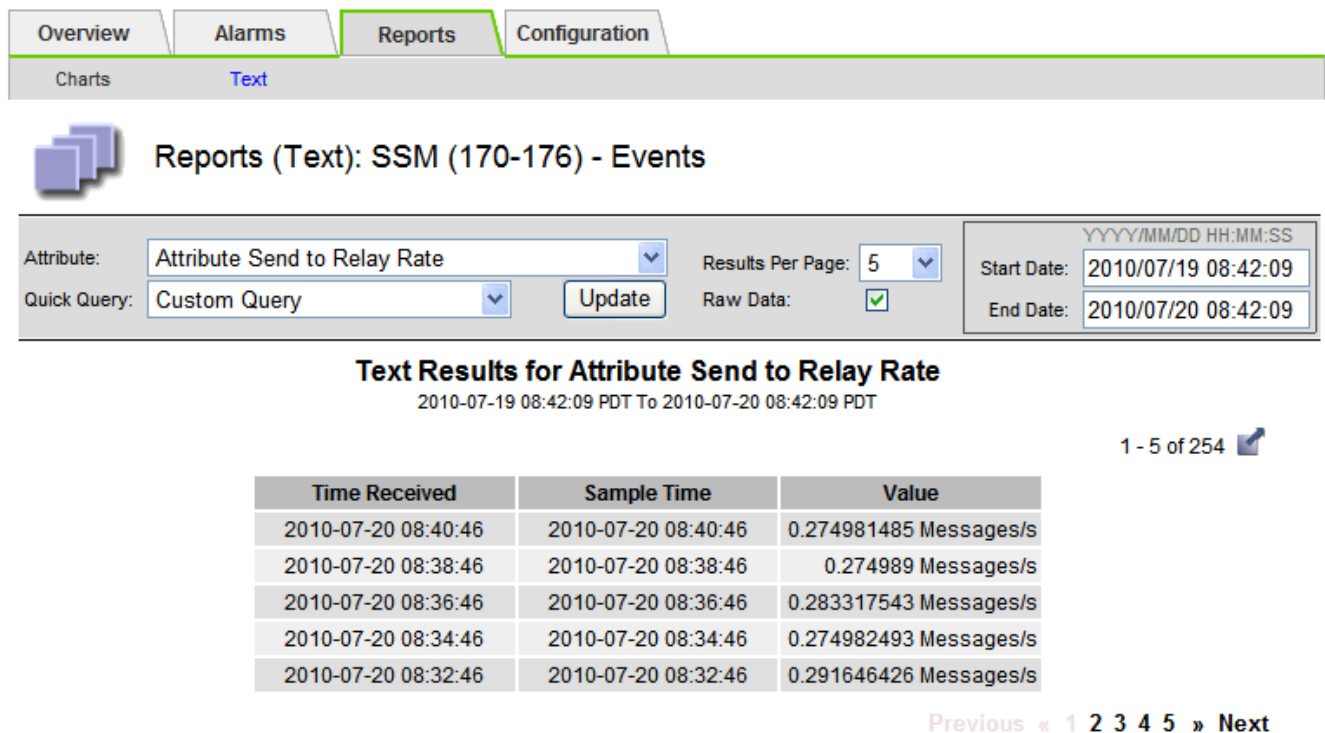
Exported text reports open a new browser tab, which enables you to select and copy the data.

About this task

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

Steps

- Select **SUPPORT > Tools > Grid topology**.
- Create a text report.
- Click ***Export*** .



The screenshot shows the 'Reports' tab in the StorageGRID interface. The report title is 'Reports (Text): SSM (170-176) - Events'. The report is for the attribute 'Attribute Send to Relay Rate' using a 'Custom Query'. The time range is from 2010/07/19 08:42:09 to 2010/07/20 08:42:09. The report shows 5 results per page, with raw data checked. The table below displays the results.

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

Navigation: Previous « 1 2 3 4 5 » Next

The Export Text Report window opens displaying the report.

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

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

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

Monitor PUT and GET performance

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

About this task

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

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

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

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

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

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

- You can [download S3tester from github](#).

Monitor object verification operations

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

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Maintenance or Root access permission](#).

About this task

Two [verification processes](#) work together to ensure data integrity:

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

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



The **Unidentified corrupt object detected** alert is triggered if the system detects a corrupt object that can't be corrected automatically.

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

Object existence check verifies whether all expected replicated copies of objects and erasure-coded fragments exist on a Storage Node. Object existence check provides a way to verify the integrity of storage devices, especially if a recent hardware issue could have affected data integrity.

You should review the results from background verifications and object existence checks regularly. Investigate any instances of corrupt or missing object data immediately to determine the root cause.

Steps

1. Review the results from background verifications:
 - a. Select **NODES > Storage Node > Objects**.
 - b. Check the verification results:
 - To check replicated object data verification, look at the attributes in the Verification section.

Verification

Status: ?	No errors	
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	

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

Erasure coding verification

Status: ?	Idle	
Next scheduled: ?	2021-10-08 10:45:19 MDT	
Fragments verified: ?	0	
Data verified: ?	0 bytes	
Corrupt copies: ?	0	
Corrupt fragments: ?	0	
Missing fragments: ?	0	

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

- Review the results from object existence check jobs:
 - Select **MAINTENANCE** > **Object existence check** > **Job history**.
 - Scan the Missing object copies detected column. If any jobs resulted in 100 or more missing object copies and the **Objects lost** alert has been triggered, contact technical support.

Object existence check

Perform an object existence check if you suspect storage volumes have been damaged or are corrupt. You can verify that objects defined by your ILM policy, still exist on the volumes.

Active job | **Job history**

<input type="checkbox"/>	Job ID [?]	Status [⬇]	Nodes (volumes) [?]	Missing object copies detected [?]
<input type="checkbox"/>	15816859223101303015	Completed	DC2-S1 (3 volumes)	0
<input type="checkbox"/>	12538643155010477372	Completed	DC1-S3 (1 volume)	0
<input type="checkbox"/>	5490044849774982476	Completed	DC1-S2 (1 volume)	0
<input type="checkbox"/>	3395284277055907678	Completed	DC1-S1 (3 volumes) DC1-S2 (3 volumes) DC1-S3 (3 volumes) and <u>7 more</u>	0

Monitor events

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

Event messages are also listed in the `/var/local/log/bycast-err.log` log file. See the [Log files reference](#).

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

Steps

1. Review the system events for each grid node:
 - a. Select **SUPPORT > Tools > Grid topology**.
 - b. Select **site > grid node > SSM > Events > Overview > Main**.
2. Generate a list of previous event messages to help isolate issues that occurred in the past:

- a. Select **SUPPORT** > **Tools** > **Grid topology**.
- b. Select **site** > **grid node** > **SSM** > **Events** > **Reports**.
- c. Select **Text**.

The **Last Event** attribute is not shown in the [charts view](#). To view it:

- d. Change **Attribute** to **Last Event**.
- e. Optionally, select a time period for **Quick Query**.
- f. Select **Update**.

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 }

Create custom syslog events

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



About this task

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

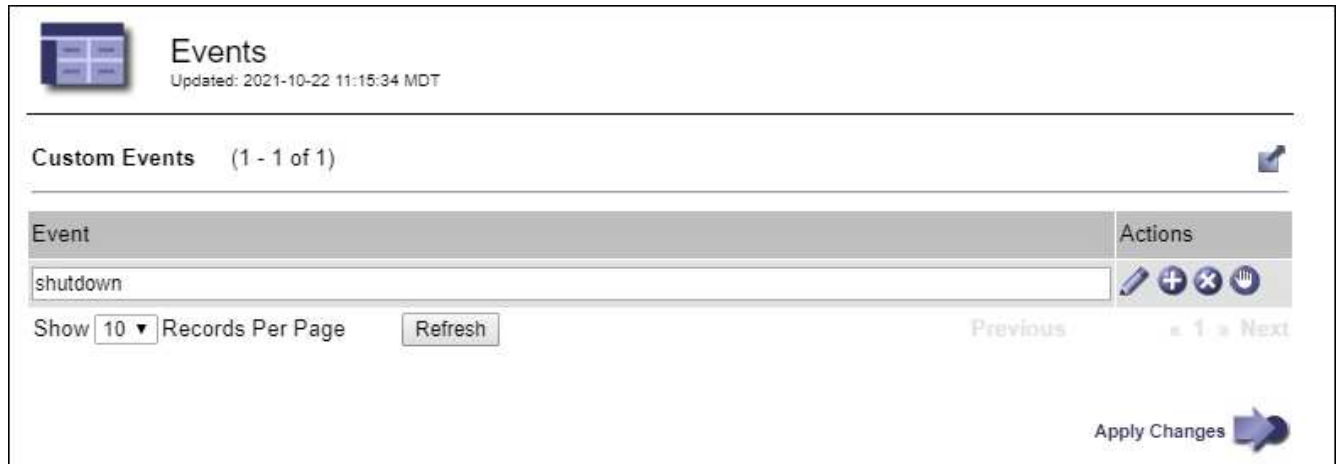
- After a custom event is created, every occurrence of it is monitored.
- To create a custom event based on keywords in the `/var/local/log/messages` files, the logs in those files must be:
 - Generated by the kernel
 - Generated by daemon or user program at the error or critical level

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

Steps

1. Select **SUPPORT** > **Alarms (legacy)** > **Custom events**.
2. Click **Edit**  (or **Insert**  if this is not the first event).

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



The screenshot shows a web interface titled "Events" with a sub-header "Custom Events (1 - 1 of 1)". Below this is a table with two columns: "Event" and "Actions". The "Event" column contains the text "shutdown". The "Actions" column contains icons for edit, add, delete, and refresh. Below the table, there is a "Show 10 Records Per Page" dropdown menu, a "Refresh" button, and navigation links for "Previous" and "Next" (with "1" in the middle). At the bottom right, there is an "Apply Changes" button with a right-pointing arrow icon.

4. Select **Apply Changes**.

5. Select **SUPPORT > Tools > Grid topology**.

6. Select **grid node > SSM > Events**.

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

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

Overview
Alarms
Reports
Configuration

Main

Overview: SSM (DC1-ADM1) - Events

Updated: 2021-10-22 11:19:18 MDT

System Events

Log Monitor State: Connected

Total Events: 0

Last Event: No Events

Description	Count	✕
Abnormal Software Events	0	✕
Account Service Events	0	✕
Cassandra Errors	0	✕
Cassandra Heap Out Of Memory Errors	0	✕
Chunk Service Events	0	✕
Custom Events	0	✕
Data-Mover Service Events	0	✕
File System Errors	0	✕
Forced Termination Events	0	✕
Grid Node Errors	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	✕
Out Of Memory Errors	0	✕
Replicated State Machine Service Events	0	✕
SCSI Errors	0	✕


Reset the count of custom events to zero

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

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

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **grid node > SSM > Events > Configuration > Main**.
3. Select the **Reset** checkbox for Custom Events.

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. Select **Apply Changes**.

Review audit messages

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

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

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

Each Admin Node stores audit messages in text files. The audit share contains the active file (audit.log) as well as compressed audit logs from previous days. Each node in the grid also stores a copy of the audit information generated on the node.

For easy access to audit logs, you can [configure audit client access for NFS](#). You can also access audit log files directly from the command line of the Admin Node.

Optionally, you can change the destination of audit logs and send audit information to an external syslog server. Local logs of audit records continue to be generated and stored when an external syslog server is configured. See [Configure audit messages and log destinations](#).

For details on the audit log file, the format of audit messages, the types of audit messages, and the tools available to analyze audit messages, see [Review audit logs](#).

Collect log files and system data

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

Before you begin

- You must be signed in to the Grid Manager on the primary Admin Node using a [supported web browser](#).
- You have [specific access permissions](#).
- You must have the provisioning passphrase.

About this task

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

Optionally, you can change the destination of audit logs and send audit information to an external syslog server. Local logs of audit records continue to be generated and stored when an external syslog server is configured. See [Configure audit messages and log destinations](#).

Steps

1. Select **SUPPORT > Tools > Logs**.

The screenshot shows the 'Collect Logs' configuration page in the Grid Manager. On the left, a tree view shows the StorageGRID hierarchy. Under 'StorageGRID', 'DC1' is expanded, and 'DC1-S1' is selected. Under 'DC2', 'DC2-S1' is selected. On the right, the 'Log Start Time' is set to 2021-12-03 06:31 AM MST, and the 'Log End Time' is set to 2021-12-03 10:31 AM MST. The 'Log Types' section has 'Application Logs' checked, while 'Audit Logs', 'Network Trace', and 'Prometheus Database' are unchecked. There is a 'Notes' text area and a 'Provisioning Passphrase' field with masked characters. A 'Collect Logs' button is 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. Select the types of logs you want to collect.
 - **Application Logs:** Application-specific logs that technical support uses most frequently for troubleshooting. The logs collected are a subset of the available application logs.
 - **Audit Logs:** Logs containing the audit messages generated during normal system operation.
 - **Network Trace:** Logs used for network debugging.
 - **Prometheus Database:** Time series metrics from the services on all nodes.
5. Optionally, enter notes about the log files you are gathering in the **Notes** text box.

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

6. Enter the provisioning passphrase for your StorageGRID system in the **Provisioning Passphrase** text box.
7. Select **Collect Logs**.

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

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

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

8. Select **Download** when log file collection is complete.

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

After you finish

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

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

Manually trigger an AutoSupport package

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

Before you begin

- You must be signed in to the Grid Manager using a [supported web browser](#).
- You must have the Root access or Other grid configuration permission.

Steps

1. Select **SUPPORT > Tools > AutoSupport**.
2. On the **Actions** tab, select **Send User-Triggered AutoSupport**.

StorageGRID attempts to send an AutoSupport package to the NetApp Support Site. 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 package again.

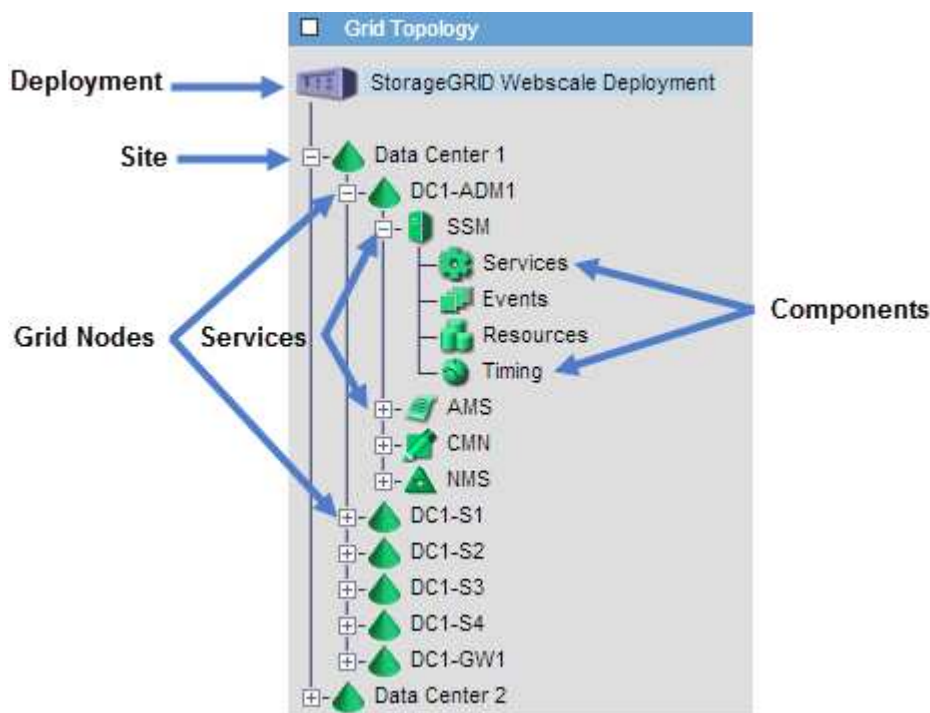


After sending an User-triggered AutoSupport package, refresh the AutoSupport page in your browser after 1 minute to access the most recent results.

View the Grid Topology tree

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

To access the Grid Topology tree, select **SUPPORT > Tools > Grid topology**.



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

StorageGRID attributes

Attributes report values and statuses for many of the functions of the StorageGRID system. Attribute values are available for each grid node, each site, and the entire grid.

StorageGRID attributes are used in several places in the Grid Manager:

- **Nodes page:** Many of the values shown on the Nodes page are StorageGRID attributes. (Prometheus metrics are also shown on the Nodes pages.)
- **Alarms:** When attributes reach defined threshold values, StorageGRID alarms (legacy system) are triggered at specific severity levels.
- **Grid Topology tree:** Attribute values are shown in the Grid Topology tree (**SUPPORT > Tools > Grid topology**).
- **Events:** System events occur when certain attributes record an error or fault condition for a node, including errors such as network errors.

Attribute values

Attributes are reported on a best-effort basis and are approximately correct. Attribute updates can be lost under some circumstances, such as the crash of a service or the failure and rebuild of a grid node.

In addition, propagation delays might slow the reporting of attributes. Updated values for most attributes are sent to the StorageGRID system at fixed intervals. It can take several minutes before an update is visible in the system, and two attributes that change more or less simultaneously can be reported at slightly different times.

Review support metrics

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

Before you begin

- You must be signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).

About this task

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



The tools available on the Metrics page are intended for use by technical support. Some features and menu items within these tools are intentionally non-functional and are subject to change. See the list of [commonly used Prometheus metrics](#).

Steps

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

An example of the Metrics page is shown here:

Metrics

Access charts and metrics to help troubleshoot issues.

 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://...>

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

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


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



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

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

The Grafana interface for the link you selected appears.



Run diagnostics

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




- [Review support metrics](#)
- [Commonly used Prometheus metrics](#)

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).

About this task

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

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

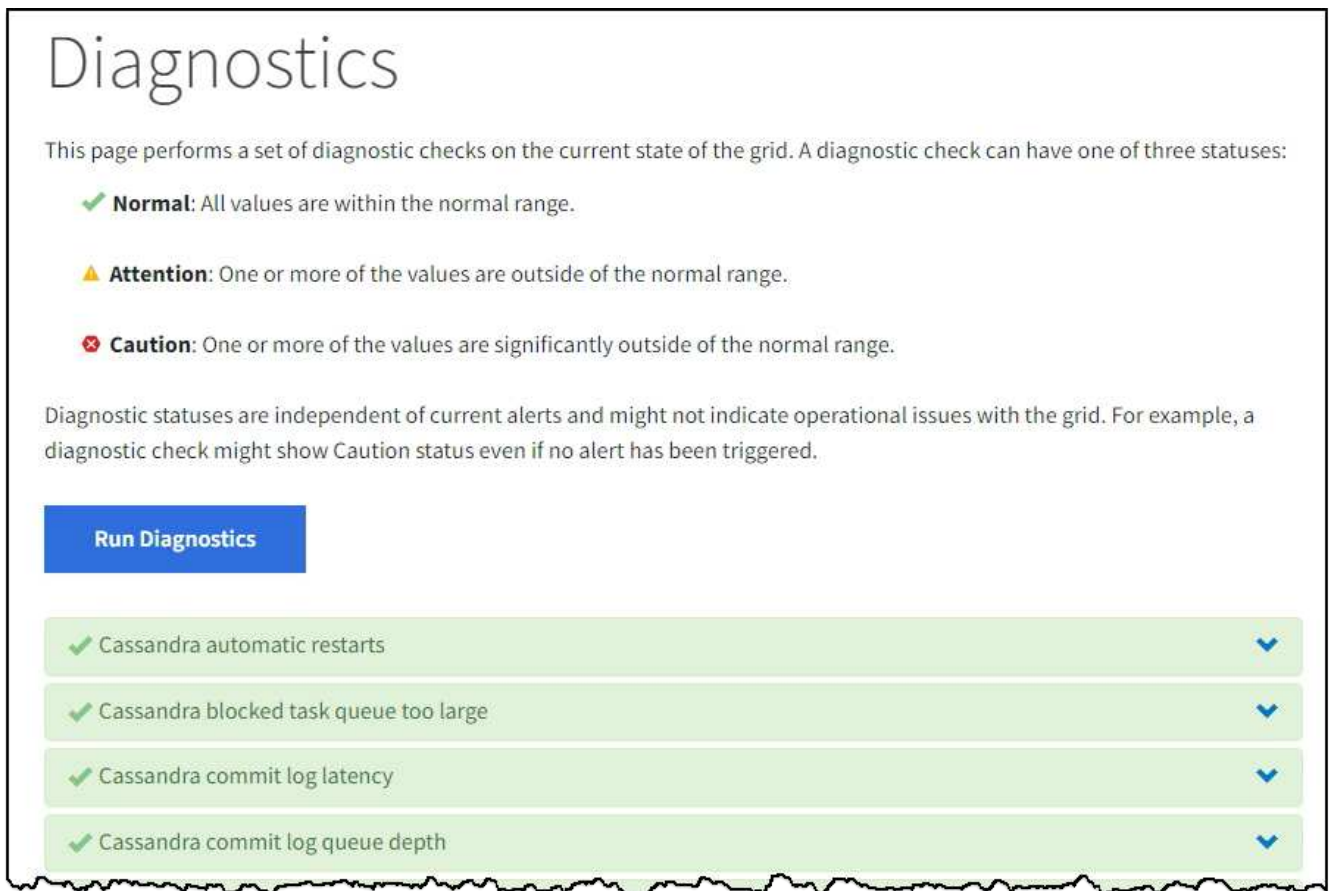
Diagnostic statuses are independent of current alerts and might not indicate operational issues with the grid. For example, a diagnostic check might show Caution status even if no alert has been triggered.

Steps

1. Select **SUPPORT > Tools > Diagnostics**.

The Diagnostics page appears and lists the results for each diagnostic check. The results are sorted by severity (Caution, Attention, and then Normal). Within each severity, the results are sorted alphabetically.

In this example, all diagnostics have a Normal status.



The screenshot shows a page titled "Diagnostics". Below the title, there is a brief explanation: "This page performs a set of diagnostic checks on the current state of the grid. A diagnostic check can have one of three statuses:". This is followed by the same three status definitions as in the first block: Normal (green checkmark), Attention (yellow warning triangle), and Caution (red X). Below this, another explanatory sentence states: "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." A blue button labeled "Run Diagnostics" is positioned above a list of four diagnostic checks. Each check is displayed in a light green row with a green checkmark icon on the left and a blue downward arrow on the right. The checks listed are: "Cassandra automatic restarts", "Cassandra blocked task queue too large", "Cassandra commit log latency", and "Cassandra commit log queue depth".

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

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

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



You can't change these thresholds.

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

✓ **CPU utilization** ^

Checks the current CPU utilization on each node.

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

Status ✓ Normal

Prometheus query `sum by (instance) (sum by (instance, mode) (irate(node_cpu_seconds_total{mode!="idle"}[5m])) / count by (instance, mode)(node_cpu_seconds_total{mode!="idle"}))`
[View in Prometheus](#)

Thresholds

- ⚠ Attention >= 75%
- ✖ Caution >= 95%

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

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

This link is not displayed for all diagnostics.

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



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



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

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

Enable query history

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


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

Execute

- insert metric at cursor -

Graph Console

1h + << Until >> Res. (s) stacked



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

Remove Graph

Add Graph

Create custom monitoring applications

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

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

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

To view the metrics API operations, including the complete list of the metrics that are available, go to the Grid Manager. From the top of the page, select the help icon and select **API documentation > metrics**.

metrics Operations on metrics



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

The details of how to implement a custom monitoring application are beyond the scope of this documentation.

Copyright information

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

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

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

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

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

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

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

Trademark information

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