



Troubleshoot a StorageGRID system

StorageGRID

NetApp
June 13, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-116/monitor/troubleshooting-storagegrid-system.html> on June 13, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Troubleshoot a StorageGRID system	1
Troubleshoot a StorageGRID system	1
Overview of problem determination	1
Troubleshoot object and storage issues	9
Confirm object data locations	9
Object store (storage volume) failures	12
Verify object integrity	14
Troubleshoot lost and missing object data	21
Troubleshoot the Low object data storage alert	34
Troubleshoot Low read-only watermark override alerts	35
Troubleshoot the Storage Status (SSTS) alarm	39
Troubleshoot delivery of platform services messages (SMTT alarm)	43
Troubleshoot metadata issues	44
Troubleshoot the Low metadata storage alert	44
Troubleshoot the Services: Status - Cassandra (SVST) alarm	46
Troubleshoot Cassandra Out of Memory errors (SMTT alarm)	50
Troubleshoot certificate errors	50
Troubleshoot Admin Node and user interface issues	52
Troubleshoot sign-on errors	52
Troubleshoot user interface issues	55
Check the status of an unavailable Admin Node	56
Troubleshoot network, hardware, and platform issues	57
Troubleshoot “422: Unprocessable Entity” errors	57
Troubleshoot the Grid Network MTU mismatch alert	58
Troubleshoot the Network Receive Error (NRER) alarm	59
Troubleshoot time synchronization errors	61
Linux: Network connectivity issues	61
Linux: Node status is “orphaned”	62
Linux: Troubleshoot IPv6 support	63
Troubleshoot an external syslog server	65


Troubleshoot a StorageGRID system

Troubleshoot a StorageGRID system

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

Overview of problem determination

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



Define the problem

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

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

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

Assess the risk and impact on the system

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

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

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

Collect data

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

Type of data to collect	Why collect this dat	Instructions
Create timeline of recent changes	Changes to your StorageGRID system, its configuration, or its environment can cause new behavior.	<ul style="list-style-type: none"> • Create a timeline of recent changes
Review alerts and alarms	<p>Alerts and alarms can help you quickly determine the root cause of a problem by providing important clues as to the underlying issues that might be causing it.</p> <p>Review the list of current alerts and alarms to see if StorageGRID has identified the root cause of a problem for you.</p> <p>Review alerts and alarms triggered in the past for additional insights.</p>	<ul style="list-style-type: none"> • View current alerts • View legacy alarms • View resolved alerts • Review historical alarms and alarm frequency (legacy system)
Monitor events	Events include any system error or fault events for a node, including errors such as network errors. Monitor events to learn more about issues or to help with troubleshooting.	<ul style="list-style-type: none"> • Monitor events
Identify trends using charts and text reports	Trends can provide valuable clues about when issues first appeared, and can help you understand how quickly things are changing.	<ul style="list-style-type: none"> • Use charts and graphs • Use text reports
Establish baselines	Collect information about the normal levels of various operational values. These baseline values, and deviations from these baselines, can provide valuable clues.	<ul style="list-style-type: none"> • Establish baselines
Perform ingest and retrieval tests	To troubleshoot performance issues with ingest and retrieval, use a workstation to store and retrieve objects. Compare results against those seen when using the client application.	<ul style="list-style-type: none"> • Monitor PUT and GET performance
Review audit messages	Review audit messages to follow StorageGRID operations in detail. The details in audit messages can be useful for troubleshooting many types of issues, including performance issues.	<ul style="list-style-type: none"> • Review audit messages
Check object locations and storage integrity	If you are having storage problems, verify that objects are being placed where you expect. Check the integrity of object data on a Storage Node.	<ul style="list-style-type: none"> • Monitor object verification operations • Confirm object data locations • Verify object integrity

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

Create a timeline of recent changes

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

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

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

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

Examples of significant recent changes

Here are some examples of potentially significant changes:

- Was the StorageGRID system recently installed, expanded, or recovered?
- Has the system been upgraded recently? Was a hotfix applied?
- Has any hardware been repaired or changed recently?
- Has the ILM policy been updated?
- Has the client workload changed?
- Has the client application or its behavior changed?
- Have you changed load balancers, or added or removed a high availability group of Admin Nodes or Gateway Nodes?

- Have any tasks been started that might take a long time to complete? Examples include:
 - Recovery of a failed Storage Node
 - Storage Node decommissioning
- Have any changes been made to user authentication, such as adding a tenant or changing LDAP configuration?
- Is data migration taking place?
- Were platform services recently enabled or changed?
- Was compliance enabled recently?
- Have Cloud Storage Pools been added or removed?
- Have any changes been made to storage compression or encryption?
- Have there been any changes to the network infrastructure? For example, VLANs, routers, or DNS.
- Have any changes been made to NTP sources?
- Have any changes been made to the Grid, Admin, or Client Network interfaces?
- Have any configuration changes been made to the Archive Node?
- Have any other changes been made to the StorageGRID system or its environment?

Establish baselines

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

Property	Value	How to obtain
Average storage consumption	GB consumed/day Percent consumed/day	<p>Go to the Grid Manager. On the Nodes page, select the entire grid or a site and go to the Storage tab.</p> <p>On the Storage Used - Object Data chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate how much storage is consumed each day</p> <p>You can collect this information for the entire system or for a specific data center.</p>

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

Property	Value	How to obtain
Objects queued from client operations	Objects/second	Select NODES > grid > ILM . On the ILM Queue chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate a baseline value for Objects queued (from client operations) for your system.
Average query latency	Milliseconds	Select NODES > Storage Node > Objects . In the Queries table, view the value for Average Latency.

Analyze data


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


The analysis is problem-dependent, but in general:

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

Escalation information checklist

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

	Item	Notes
	Problem statement	What are the problem symptoms? When did the problem start? Does it happen consistently or intermittently? If intermittently, what times has it occurred? Define the problem
	Impact assessment	What is the severity of the problem? What is the impact to the client application? <ul style="list-style-type: none"> • Has the client connected successfully before? • Can the client ingest, retrieve, and delete data?
	StorageGRID System ID	Select MAINTENANCE > System > License . The StorageGRID System ID is shown as part of the current license.

	Item	Notes
	Software version	From the top of the Grid Manager, select the help icon and select About to see the StorageGRID version.
	Customization	<p>Summarize how your StorageGRID system is configured. For example, list the following:</p> <ul style="list-style-type: none"> • Does the grid use storage compression, storage encryption, or compliance? • Does ILM make replicated or erasure coded objects? Does ILM ensure site redundancy? Do ILM rules use the Strict, Balanced, or Dual Commit ingest behaviors?
	Log files and system data	<p>Collect log files and system data for your system. Select SUPPORT > Tools > Logs.</p> <p>You can collect logs for the entire grid, or for selected nodes.</p> <p>If you are collecting logs only for selected nodes, be sure to include at least one Storage Node that has the ADC service. (The first three Storage Nodes at a site include the ADC service.)</p> <p>Collect log files and system data</p>
	Baseline information	<p>Collect baseline information regarding ingest operations, retrieval operations, and storage consumption.</p> <p>Establish baselines</p>
	Timeline of recent changes	<p>Create a timeline that summarizes any recent changes to the system or its environment.</p> <p>Create a timeline of recent changes</p>
	History of efforts to diagnose the issue	If you have taken steps to diagnose or troubleshoot the issue yourself, make sure to record the steps you took and the outcome.

Troubleshoot object and storage issues

Confirm object data locations

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

object data is being stored where intended.

What you'll need

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

Steps

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

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

3. If you want to look up a specific version of the object, enter the version ID (optional).

Object Metadata Lookup

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

Identifier	<input type="text" value="source/testobject"/>
Version ID (optional)	<input type="text" value="MEJGMkMyQzgtNEY5OC0xMUU3LTkzMEYtRDkyNTAwQkY5I"/>

Look Up

4. Select **Look Up**.

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

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

- For segmented objects and multipart objects, a list of object segments including segment identifiers and data sizes. For objects with more than 100 segments, only the first 100 segments are shown.
- All object metadata in the unprocessed, internal storage format. This raw metadata includes internal system metadata that is not guaranteed to persist from release to release.

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

System Metadata

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

Replicated Copies

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

Raw Metadata

```
{
  "TYPE": "CTNT",
  "CHND": "A12E96FF-B13F-4905-9E9E-45373F6E7DA8",
  "NAME": "testobject",
  "CBID": "0x8823DE7EC7C10416",
  "PHND": "FEA0AE51-534A-11EA-9FCD-31FF00C36D56",
  "PPTH": "source",
  "META": {
    "BASE": {
      "PAW5": "2",

```

Related information

[Manage objects with ILM](#)

[Use S3](#)

[Use Swift](#)

Object store (storage volume) failures








The underlying storage on a Storage Node is divided into object stores. Object stores are also known as storage volumes.

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
















Disk devices

Name ? ⇅	World Wide Name ? ⇅	I/O load ? ⇅	Read rate ? ⇅	Write rate ? ⇅
sdC(8:16,sdb)	N/A	0.05%	0 bytes/s	4 KB/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s
sdf(8:64,sde)	N/A	0.00%	0 bytes/s	82 bytes/s
sdg(8:80,sdf)	N/A	0.00%	0 bytes/s	82 bytes/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s
croot(8:1,sda1)	N/A	0.04%	0 bytes/s	4 KB/s
cvloc(8:2,sda2)	N/A	0.95%	0 bytes/s	52 KB/s

Volumes


Mount point ? ⇅	Device ? ⇅	Status ? ⇅	Size ? ⇅	Available ? ⇅	Write cache status ? ⇅
/	croot	Online	21.00 GB	14.73 GB 	Unknown
/var/local	cvloc	Online	85.86 GB	80.94 GB 	Unknown
/var/local/rangedb/0	sdC	Online	107.32 GB	107.17 GB 	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB 	Enabled
/var/local/rangedb/2	sde	Online	107.32 GB	107.18 GB 	Enabled
/var/local/rangedb/3	sdf	Online	107.32 GB	107.18 GB 	Enabled
/var/local/rangedb/4	sdg	Online	107.32 GB	107.18 GB 	Enabled

Object stores

ID ? ⇅	Size ? ⇅	Available ? ⇅	Replicated data ? ⇅	EC data ? ⇅	Object data (%) ? ⇅	Health ? ⇅
0000	107.32 GB	96.44 GB 	1.55 MB 	0 bytes 	0.00%	No Errors
0001	107.32 GB	107.18 GB 	0 bytes 	0 bytes 	0.00%	No Errors
0002	107.32 GB	107.18 GB 	0 bytes 	0 bytes 	0.00%	No Errors
0003	107.32 GB	107.18 GB 	0 bytes 	0 bytes 	0.00%	No Errors
0004	107.32 GB	107.18 GB 	0 bytes 	0 bytes 	0.00%	No Errors




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

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **site > Storage Node > LDR > Storage > Overview > Main**.





Overview: LDR (DC1-S1) - Storage







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

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













Utilization

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

Replication

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

Object Store Volumes

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

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

Related information

[Recover and maintain](#)

Verify object integrity

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

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

What is background verification?

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

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

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



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

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

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

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

Alerts and alarms (legacy) related to background verification

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

If background verification cannot replace a corrupted object because it cannot locate another copy, the **Objects lost** alert is triggered.

Change the background verification rate

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

What you'll need

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

About this task

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

- **Adaptive:** Default setting. The task is designed to verify at a maximum of 4 MB/s or 10 objects/s (whichever is exceeded first).
- **High:** Storage verification proceeds quickly, at a rate that can slow ordinary system activities.

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

Steps

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

Overview Alarms Reports Configuration

Main

Configuration: LDR (Storage Node) - Verification
Updated: 2021-11-11 07:13:00 MST

Reset Missing Objects Count ☐

Background Verification

Verification Rate Adaptive

Reset Corrupt Objects Count ☐

Quarantined Objects

Delete Quarantined Objects ☐

Apply Changes



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

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

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

- If the object identifier can be extracted, StorageGRID automatically creates a new copy of the object data. The new copy can be made anywhere in the StorageGRID system that satisfies the active ILM policy.
- If the object identifier cannot be extracted (because it has been corrupted), the **Corrupt Objects Unidentified** metric is incremented, and the **Unidentified corrupt object detected** alert is

triggered.

- c. If corrupt replicated object data is found, contact technical support to determine the root cause of the corruption.

8. Monitor the results of background verification for erasure-coded objects.

If background verification finds corrupt fragments of erasure-coded object data, the Corrupt Fragments Detected attribute is incremented. StorageGRID recovers by rebuilding the corrupt fragment in place on the same Storage Node.

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

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

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

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



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

- a. Select **SUPPORT > Tools > Grid topology**.
- b. Select **Storage Node > LDR > Verification > Configuration**.
- c. Select **Delete Quarantined Objects**.
- d. Select **Apply Changes**.

What is object existence check?

Object existence check verifies whether all expected replicated copies of objects and erasure-coded fragments exist on a Storage Node. Object existence check does not verify the object data itself (background verification does that); instead, it provides a way to verify the integrity of storage devices, especially if a recent hardware issue could have affected data integrity.

Unlike background verification, which occurs automatically, you must manually start an object existence check job.

Object existence check reads the metadata for every object stored in StorageGRID and verifies the existence of both replicated object copies and erasure-coded object fragments. Any missing data is handled as follows:

- **Replicated copies:** If a copy of replicated object data is missing, StorageGRID automatically attempts to replace the copy from a copy stored elsewhere in the system. The Storage Node runs an existing copy through an ILM evaluation, which will determine that the current ILM policy is no longer being met for this object because another copy is missing. A new copy is generated and placed to satisfy the system's active ILM policy. This new copy might not be placed in the same location where the missing copy was stored.
- **Erasure-coded fragments:** If a fragment of an erasure-coded object is missing, StorageGRID

automatically attempts to rebuild the missing fragment in place on the same Storage Node using the remaining fragments. If the missing fragment cannot be rebuilt (because too many fragments have been lost), ILM attempts to find another copy of the object, which it can use to generate a new erasure-coded fragment.

Run object existence check

You create and run one object existence check job at a time. When you create a job, you select the Storage Nodes and volumes you want to verify. You also select the consistency control for the job.

What you'll need

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Maintenance or Root Access permission.
- You have ensured that the Storage Nodes you want to check are online. Select **NODES** to view the table of nodes. Ensure that no alert icons appear next to the node name for the nodes you want to check.
- You have ensured that the following procedures are **not** running on the nodes you want to check:
 - Grid expansion to add a Storage Node
 - Storage Node decommission
 - Recovery of a failed storage volume
 - Recovery of a Storage Node with a failed system drive
 - EC rebalance
 - Appliance node clone

Object existence check does not provide useful information while these procedures are in progress.

About this task

An object existence check job can take days or weeks to complete, depending on the number of objects in the grid, the selected storage nodes and volumes, and the selected consistency control. You can run only one job at a time, but you can select multiple Storage Nodes and volumes at the same time.

Steps

1. Select **MAINTENANCE > Tasks > Object existence check**.
2. Select **Create job**. The Create an object existence check job wizard appears.
3. Select the nodes containing the volumes you want to verify. To select all online nodes, select the **Node name** check box in the column header.

You can search by node name or site.

You cannot select nodes that are not connected to the grid.

4. Select **Continue**.
5. Select one or more volumes for each node in the list. You can search for volumes using the storage volume number or node name.

To select all volumes for each node you selected, select the **Storage volume** check box in the column header.

6. Select **Continue**.

7. Select the consistency control for the job.

The consistency control determines how many copies of object metadata are used for the object existence check.

- **Strong-site**: Two copies of metadata at a single site.
- **Strong-global**: Two copies of metadata at each site.
- **All** (default): All three copies of metadata at each site.

For more information about consistency control, see the descriptions in the wizard.

8. Select **Continue**.

9. Review and verify your selections. You can select **Previous** to go to a previous step in the wizard to update your selections.

An Object existence check job is generated and runs until one of the following occurs:

- The job completes.
- You pause or cancel the job. You can resume a job that you have paused, but you cannot resume a job that you have canceled.
- The job stalls. The **Object existence check has stalled** alert is triggered. Follow the corrective actions specified for the alert.
- The job fails. The **Object existence check has failed** alert is triggered. Follow the corrective actions specified for the alert.
- A “Service unavailable” or an “Internal server error” message appears. After one minute, refresh the page to continue monitoring the job.



As needed, you can navigate away from the Object existence check page and return to continue monitoring the job.

10. As the job runs, view the **Active job** tab and note the value of Missing object copies detected.

This value represents the total number of missing copies of replicated objects and erasure-coded objects with one or more missing fragments.

If the number of Missing object copies detected is greater than 100, there might be an issue with the Storage Node's storage.

Object existence check

Perform an object existence check if you suspect some storage volumes have been damaged or are corrupt and you want to verify that objects still exist on these volumes.

If you have questions about running object existence check, contact technical support.

Active job

Job history

Status: Accepted

Consistency control: All

Job ID: 2334602652907829302

Start time: 2021-11-10 14:43:02 MST

Missing object copies detected: 0

Elapsed time: —

Progress: 0%

Estimated time to completion: —

Pause

Cancel

Volumes

Details

Selected node	Selected storage volumes	Site
DC1-S1	0, 1, 2	Data Center 1
DC1-S2	0, 1, 2	Data Center 1
DC1-S3	0, 1, 2	Data Center 1

11. After the job has completed, take any additional required actions:

- If Missing object copies detected is zero, then no issues were found. No action is required.
- If Missing object copies detected is greater than zero and the **Objects lost** alert has not been triggered, then all missing copies were repaired by the system. Verify that any hardware issues have been corrected to prevent future damage to object copies.
- If Missing object copies detected is greater than zero and the **Objects lost** alert has been triggered, then data integrity could be affected. Contact technical support.
- You can investigate lost object copies by using grep to extract the LLST audit messages: `grep LLST audit_file_name`.

This procedure is similar to the one for [investigating lost objects](#), although for object copies you search for LLST instead of OLSST.

12. If you selected the strong-site or strong-global consistency control for the job, wait approximately three weeks for metadata consistency and then rerun the job on the same volumes again.

When StorageGRID has had time to achieve metadata consistency for the nodes and volumes included in the job, rerunning the job could clear erroneously reported missing object copies or cause additional object copies to be checked if they were missed.

a. Select **MAINTENANCE > Object existence check > Job history**.

b. Determine which jobs are ready to be rerun:

- Look at the **End time** column to determine which jobs were run more than three weeks ago.

- ii. For those jobs, scan the Consistency control column for strong-site or strong-global.
- c. Select the check box for each job you want to rerun, then select **Rerun**.

Object existence check

Perform an object existence check if you suspect some storage volumes have been damaged or are corrupt and you want to verify that objects still exist on these volumes.

If you have questions about running object existence check, contact technical support.

Active job
Job history

Delete
Rerun

Displaying 4 results

<input type="checkbox"/>	Job ID ?	Status ?	Nodes (volumes) ?	Missing object copies detected ?	Consistency control ?	Start time ?	End time ?
<input checked="" type="checkbox"/>	2334602652907829302	Completed	DC1-S1 (3 volumes) DC1-S2 (3 volumes) DC1-S3 (3 volumes) and 7 more	0	All	2021-11-10 14:43:02 MST	2021-11-10 14:43:06 MST (3 weeks ago)
<input type="checkbox"/>	11725651898848823235 (Rerun job)	Completed	DC1-S2 (2 volumes) DC1-S3 (2 volumes) DC1-S4 (2 volumes) and 4 more	0	Strong-site	2021-11-10 14:42:10 MST	2021-11-10 14:42:11 MST (17 minutes ago)

- d. In the Rerun jobs wizard, review the selected nodes and volumes and the consistency control.
- e. When you are ready to rerun the jobs, select **Rerun**.

The Active job tab appears. All the jobs you selected are rerun as one job at a consistency control of strong-site. A **Related jobs** field in the Details section lists the job IDs for the original jobs.

After you finish

If you still have concerns about data integrity, go to **SUPPORT > Tools > Grid topology > site > Storage Node > LDR > Verification > Configuration > Main** and increase the Background Verification Rate. Background verification checks the correctness of all stored object data and repairs any issues that it finds. Finding and repairing potential issues as quickly as possible reduces the risk of data loss.

Troubleshoot lost and missing object data

Objects can be retrieved for several reasons, including read requests from a client application, background verifications of replicated object data, ILM re-evaluations, and the restoration of object data during the recovery of a Storage Node.

The StorageGRID system uses location information in an object's metadata to determine from which location to retrieve the object. If a copy of the object is not found in the expected location, the system attempts to retrieve another copy of the object from elsewhere in the system, assuming that the ILM policy contains a rule to make two or more copies of the object.

If this retrieval is successful, the StorageGRID system replaces the missing copy of the object. Otherwise, the

Objects lost alert is triggered, as follows:

- For replicated copies, if another copy cannot be retrieved, the object is considered lost, and the alert is triggered.
- For erasure coded copies, if a copy cannot be retrieved from the expected location, the Corrupt Copies Detected (ECOR) attribute is incremented by one before an attempt is made to retrieve a copy from another location. If no other copy is found, the alert is triggered.

You should investigate all **Objects lost** alerts immediately to determine the root cause of the loss and to determine if the object might still exist in an offline, or otherwise currently unavailable, Storage Node or Archive Node.

In the case where object data without copies is lost, there is no recovery solution. However, you must reset the Lost objects counter to prevent known lost objects from masking any new lost objects.

Related information

[Investigate lost objects](#)

[Reset lost and missing object counts](#)

Investigate lost objects

When the **Objects lost** alert is triggered, you must investigate immediately. Collect information about the affected objects and contact technical support.

What you'll need

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

About this task


The **Objects lost** alert indicates that StorageGRID believes that there are no copies of an object in the grid. Data might have been permanently lost.

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

Steps

1. Select **NODES**.
2. Select **Storage Node > Objects**.
3. Review the number of Lost objects shown in the Object counts table.

This number indicates the total number of objects this grid node detects as missing from the entire StorageGRID system. The value is the sum of the Lost objects counters of the Data store component within the LDR and DDS services.



4. From an Admin Node, access the audit log to determine the unique identifier (UUID) of the object that triggered the **Objects lost** alert:
 - a. Log in to the grid node:
 - i. Enter the following command: `ssh admin@grid_node_IP`
 - ii. Enter the password listed in the `Passwords.txt` file.
 - iii. Enter the following command to switch to root: `su -`
 - iv. Enter the password listed in the `Passwords.txt` file. When you are logged in as root, the prompt changes from `$` to `#`.
 - b. Change to the directory where the audit logs are located. Enter: `cd /var/local/audit/export/`
 - c. Use `grep` to extract the Object Lost (OLST) audit messages. Enter: `grep OLST audit_file_name`
 - d. Note the UUID value included in the message.

```
>Admin: # grep OLSST audit.log
2020-02-12T19:18:54.780426
[AUDT:[CBID(UI64):0x38186FE53E3C49A5][UUID(CSTR):926026C4-00A4-449B-
AC72-BCCA72DD1311]
[PATH(CSTR):"source/cats"][NOID(UI32):12288733][VOLI(UI64):3222345986
][RSLT(FC32):NONE][AVER(UI32):10]
[ATIM(UI64):1581535134780426][ATYP(FC32):OLST][ANID(UI32):12448208][A
MID(FC32):ILMX][ATID(UI64):7729403978647354233]]
```

5. Use the `ObjectByUUID` command to find the object by its identifier (UUID), and then determine if data is at risk.
 - a. Telnet to localhost 1402 to access the LDR console.
 - b. Enter: `/proc/OBRP/ObjectByUUID UUID_value`

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

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

{
  "TYPE(Object Type)": "Data object",
  "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
  "NAME": "cats",
  "CBID": "0x38186FE53E3C49A5",
  "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-
ACBB00BB82DD",
  "PPTH(Parent path)": "source",
  "META": {
    "BASE(Protocol metadata)": {
      "PAWS(S3 protocol version)": "2",
      "ACCT(S3 account ID)": "44084621669730638018",
      "*ctp(HTTP content MIME type)": "binary/octet-stream"
    },
    "BYCB(System metadata)": {
      "CSIZ(Plaintext object size)": "5242880",
      "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
      "BSIZ(Content block size)": "5252084",
      "CVER(Content block version)": "196612",
      "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
      "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
```

```

        "ITME": "1581534970983000"
    },
    "CMSM": {
        "LATM(Object last access time)": "2020-02-
12T19:16:10.983000"
    },
    "AWS3": {
        "LOCC": "us-east-1"
    }
},
"CLCO\ (Locations\)": \[
    \{
        "Location Type": "CLDI\ (Location online\)",
        "NOID\ (Node ID\)": "12448208",
        "VOLI\ (Volume ID\)": "3222345473",
        "Object File Path":
"/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila\#3udu",
        "LTIM\ (Location timestamp\)": "2020-02-
12T19:36:17.880569"
    },
    \{
        "Location Type": "CLDI\ (Location online\)",
        "NOID\ (Node ID\)": "12288733",
        "VOLI\ (Volume ID\)": "3222345984",
        "Object File Path":
"/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb\#3s;L",
        "LTIM\ (Location timestamp\)": "2020-02-
12T19:36:17.934425"
    }
]
}

```

In the second example, the object with UUID 926026C4-00A4-449B-AC72-BCCA72DD1311 has no locations listed.

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

```
{
  "TYPE(Object Type)": "Data object",
  "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
  "NAME": "cats",
  "CBID": "0x38186FE53E3C49A5",
  "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-
ACBB00BB82DD",
  "PPTH(Parent path)": "source",
  "META": {
    "BASE(Protocol metadata)": {
      "PAWS(S3 protocol version)": "2",
      "ACCT(S3 account ID)": "44084621669730638018",
      "*ctp(HTTP content MIME type)": "binary/octet-stream"
    },
    "BYCB(System metadata)": {
      "CSIZ(Plaintext object size)": "5242880",
      "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
      "BSIZ(Content block size)": "5252084",
      "CVER(Content block version)": "196612",
      "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
      "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
      "ITME": "1581534970983000"
    },
    "CMSM": {
      "LATM(Object last access time)": "2020-02-
12T19:16:10.983000"
    },
    "AWS3": {
      "LOCC": "us-east-1"
    }
  }
}
```

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

Metadata	Conclusion
No object found ("ERROR": "")	<p>If the object is not found, the message "ERROR": "" is returned.</p> <p>If the object is not found, you can reset the count of Objects lost to clear the alert. The lack of an object indicates that the object was intentionally deleted.</p>
Locations > 0	<p>If there are locations listed in the output, the Objects lost alert might be a false positive.</p> <p>Confirm that the objects exist. Use the Node ID and filepath listed in the output to confirm that the object file is in the listed location.</p> <p>(The procedure for searching for potentially lost objects explains how to use the Node ID to find the correct Storage Node.)</p> <p>If the objects exist, you can reset the count of Objects lost to clear the alert.</p>
Locations = 0	<p>If there are no locations listed in the output, the object is potentially missing. You can try to search for and restore the object yourself, or you can contact technical support.</p> <p>Technical support might ask you to determine if there is a storage recovery procedure in progress. That is, has a <i>repair-data</i> command been issued on any Storage Node, and is the recovery still in progress? See the information about restoring object data to a storage volume.</p>

Related information

[Review audit logs](#)

Search for and restore potentially lost objects

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

What you'll need

- You must have the UUID of any lost object, as identified in "Investigating lost objects."
- You must have the `Passwords.txt` file.

About this task

You can follow this procedure to look for replicated copies of the lost object elsewhere in the grid. In most cases, the lost object will not be found. However, in some cases, you might be able to find and restore a lost replicated object if you take prompt action.



Contact technical support for assistance with this procedure.

Steps

1. From an Admin Node, search the audit logs for possible object locations:

a. Log in to the grid node:

- i. Enter the following command: `ssh admin@grid_node_IP`
- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file. When you are logged in as root, the prompt changes from `$` to `#`.

b. Change to the directory where the audit logs are located: `cd /var/local/audit/export/`

c. Use `grep` to extract the audit messages associated with the potentially lost object and send them to an output file. Enter: `grep uuid-valueaudit_file_name > output_file_name`

For example:

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

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

For example:

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

An LLST audit message looks like this sample message.

```
[AUDT:\[NOID\ (UI32\):12448208\] [CBIL(UI64):0x38186FE53E3C49A5]
[UUID(CSTR):"926026C4-00A4-449B-AC72-BCCA72DD1311"] [LTYP(FC32):CLDI]
[PCLD\ (CSTR\):"/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%\#3tN6"\]
[TSRC(FC32):SYST] [RSLT(FC32):NONE] [AVER(UI32):10] [ATIM(UI64):
1581535134379225] [ATYP(FC32):LLST] [ANID(UI32):12448208] [AMID(FC32):CL
SM]
[ATID(UI64):7086871083190743409]]
```

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

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

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

f. Find the Storage Node for this LDR node ID.

There are two ways to use the node ID to find the Storage Node:

- In the Grid Manager, select **SUPPORT > Tools > Grid topology**. Then select **Data Center > Storage Node > LDR**. The LDR node ID is in the Node Information table. Review the information for each Storage Node until you find the one that hosts this LDR.
- Download and unzip the Recovery Package for the grid. There is a `ldocs` directory in the SAID package. If you open the `index.html` file, the Servers Summary shows all node IDs for all grid nodes.

2. Determine if the object exists on the Storage Node indicated in the audit message:

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

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

- Determine if the file path for the object exists.

For the file path of the object, use the value of `PCLD` from the LLST audit message.

For example, enter:

```
ls '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6'
```

Note: Always enclose the object file path in single quotes in commands to escape any special characters.

- If the object path is not found, the object is lost and cannot be restored using this procedure. Contact technical support.
- If the object path is found, continue with step [Restore the object to StorageGRID](#). You can attempt to restore the found object back to StorageGRID.

- If the object path was found, attempt to restore the object to StorageGRID:
 - From the same Storage Node, change the ownership of the object file so that it can be managed by StorageGRID. Enter: `chown ldr-user:bcast 'file_path_of_object'`
 - Telnet to localhost 1402 to access the LDR console. Enter: `telnet 0 1402`
 - Enter: `cd /proc/STOR`
 - Enter: `Object_Found 'file_path_of_object'`

For example, enter:

```
Object_Found '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6'
```

Issuing the `Object_Found` command notifies the grid of the object's location. It also triggers the active ILM policy, which makes additional copies as specified in the policy.

Note: If the Storage Node where you found the object is offline, you can copy the object to any Storage Node that is online. Place the object in any `/var/local/rangedb` directory of the online Storage Node. Then, issue the `Object_Found` command using that file path to the object.

- If the object cannot be restored, the `Object_Found` command fails. Contact technical support.
- If the object was successfully restored to StorageGRID, a success message appears. For example:

```
ade 12448208: /proc/STOR > Object_Found
'/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6'

ade 12448208: /proc/STOR > Object found succeeded.
First packet of file was valid. Extracted key: 38186FE53E3C49A5
Renamed '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6' to
'/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila#3udu'
```

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

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

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

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

{
  "TYPE(Object Type)": "Data object",
  "CHND(Content handle)": "926026C4-00A4-449B-AC72-BCCA72DD1311",
  "NAME": "cats",
  "CBID": "0x38186FE53E3C49A5",
  "PHND(Parent handle, UUID)": "221CABD0-4D9D-11EA-89C3-ACBB00BB82DD",
  "PPTH(Parent path)": "source",
  "META": {
    "BASE(Protocol metadata)": {
      "PAWS(S3 protocol version)": "2",
      "ACCT(S3 account ID)": "44084621669730638018",
      "*ctp(HTTP content MIME type)": "binary/octet-stream"
    },
    "BYCB(System metadata)": {
      "CSIZ(Plaintext object size)": "5242880",
      "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
      "BSIZ(Content block size)": "5252084",
```

```

        "CVER(Content block version)": "196612",
        "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
        "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
        "ITME": "1581534970983000"
    },
    "CMSM": {
        "LATM(Object last access time)": "2020-02-12T19:16:10.983000"
    },
    "AWS3": {
        "LOCC": "us-east-1"
    }
},
"CLCO\ (Locations\)": \[
    \{
        "Location Type": "CLDI\ (Location online\)",
        "NOID\ (Node ID\)": "12448208",
        "VOLI\ (Volume ID\)": "3222345473",
        "Object File Path":
"/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila\#3udu",
        "LTIM\ (Location timestamp\)": "2020-02-12T19:36:17.880569"
    },
    \{
        "Location Type": "CLDI\ (Location online\)",
        "NOID\ (Node ID\)": "12288733",
        "VOLI\ (Volume ID\)": "3222345984",
        "Object File Path":
"/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb\#3s;L",
        "LTIM\ (Location timestamp\)": "2020-02-12T19:36:17.934425"
    }
]
}

```

a. Sign out of the LDR console. Enter: `exit`

1. From an Admin Node, search the audit logs for the ORLM audit message for this object to confirm that information lifecycle management (ILM) has placed copies as required.

b. Log in to the grid node:

- i. Enter the following command: `ssh admin@grid_node_IP`
- ii. Enter the password listed in the `Passwords.txt` file.
- iii. Enter the following command to switch to root: `su -`
- iv. Enter the password listed in the `Passwords.txt` file. When you are logged in as root, the prompt changes from `$` to `#`.

- c. Change to the directory where the audit logs are located: `cd /var/local/audit/export/`
- d. Use `grep` to extract the audit messages associated with the object to an output file. Enter: `grep uuid-valueaudit_file_name > output_file_name`

For example:

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

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

For example:

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

An ORLM audit message looks like this sample message.

```
[AUDT:[CBID(UI64):0x38186FE53E3C49A5][RULE(CSTR):"Make 2 Copies"]
[STAT(FC32):DONE][CSIZ(UI64):0][UUID(CSTR):"926026C4-00A4-449B-AC72-
BCCA72DD1311"]
[LOCS(CSTR):"***CLDI 12828634 2148730112**, CLDI 12745543 2147552014"]
[RSLT(FC32):SUCS][AVER(UI32):10][ATYP(FC32):ORLM][ATIM(UI64):15633982306
69]
[ATID(UI64):15494889725796157557][ANID(UI32):13100453][AMID(FC32):BCMS]]
```

- f. Find the `LOCS` field in the audit message.

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

Related information

[Investigate lost objects](#)

[Reset lost and missing object counts](#)

[Review audit logs](#)

Reset lost and missing object counts

After investigating the StorageGRID system and verifying that all recorded lost objects are permanently lost or that it is a false alarm, you can reset the value of the Lost Objects attribute to zero.

What you'll need

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

About this task


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

- **SUPPORT > Tools > Grid topology > Site > Storage Node > LDR > Data Store > Overview > Main**
- **SUPPORT > Tools > Grid topology > Site > Storage Node > DDS > Data Store > Overview > Main**

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

Steps

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



4. Click **Apply Changes**.

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

5. Optionally, reset other related attribute values that might have been incremented in the process of identifying the lost object.
 - a. Select **Site > Storage Node > LDR > Erasure Coding > Configuration**.
 - b. Select **Reset Reads Failure Count** and **Reset Corrupt Copies Detected Count**.
 - c. Click **Apply Changes**.
 - d. Select **Site > Storage Node > LDR > Verification > Configuration**.
 - e. Select **Reset Missing Objects Count** and **Reset Corrupt Objects Count**.
 - f. If you are confident that quarantined objects are not required, you can select **Delete Quarantined Objects**.

Quarantined objects are created when background verification identifies a corrupt replicated object copy. In most cases StorageGRID automatically replaces the corrupt object, and it is safe to delete the quarantined objects. However, if the **Objects lost** alert or the LOST alarm is triggered, technical support might want to access the quarantined objects.

g. Click **Apply Changes**.

It can take a few moments for the attributes to reset after you click **Apply Changes**.

Troubleshoot the Low object data storage alert

The **Low object data storage** alert monitors how much space is available for storing object data on each Storage Node.

What you'll need

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

About this task

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

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

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

In this condition:

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

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

Steps

1. Select **ALERTS > Current**.

The Alerts page appears.

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




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

3. Review the details in the dialog box, and note the following:
 - Time triggered
 - The name of the site and node
 - The current values of the metrics for this alert
4. Select **NODES > Storage Node or Site > Storage**.

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

The following values are shown:

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



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

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

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

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



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

Related information

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

[Expand your grid](#)

[Administer StorageGRID](#)

Troubleshoot Low read-only watermark override alerts

If you use custom values for storage volume watermarks, you might need to resolve the **Low read-only watermark override** alert. If possible, you should update your system to

start using the optimized values.

In previous releases, the three [storage volume watermarks](#) were global settings — the same values applied to every storage volume on every Storage Node. Starting in StorageGRID 11.6, the software can optimize these watermarks for each storage volume, based on the size of the Storage Node and the relative capacity of the volume.

When you upgrade to StorageGRID 11.6, optimized read-only and read-write watermarks are automatically applied to all storage volumes, unless either of the following is true:

- Your system is close to capacity and would not be able to accept new data if optimized watermarks were applied. StorageGRID will not change watermark settings in this case.
- You previously set any of the storage volume watermarks to a custom value. StorageGRID will not override custom watermark settings with optimized values. However, StorageGRID might trigger the **Low read-only watermark override** alert if your custom value for the Storage Volume Soft Read-Only Watermark is too small.

Understand the alert

If you use custom values for storage volume watermarks, the **Low read-only watermark override** alert might be triggered for one or more Storage Nodes.

Each instance of the alert indicates that the custom value of the **Storage Volume Soft Read-Only Watermark** is smaller than the minimum optimized value for that Storage Node. If you continue to use the custom setting, the Storage Node might run critically low on space before it can safely transition to the read-only state. Some storage volumes might become inaccessible (automatically unmounted) when the node reaches capacity.

For example, suppose you previously set the **Storage Volume Soft Read-Only Watermark** to 5 GB. Now suppose that StorageGRID has calculated the following optimized values for the four storage volumes in Storage Node A:

Volume 0	12 GB
Volume 1	12 GB
Volume 2	11 GB
Volume 3	15 GB

The **Low read-only watermark override** alert is triggered for Storage Node A because your custom watermark (5 GB) is smaller than the minimum optimized value for all volumes in that node (11 GB). If you continue using the custom setting, the node might run critically low on space before it can safely transition to the read-only state.

Resolve the alert

Follow these steps if one or more **Low read-only watermark override** alerts have been triggered. You can also use these instructions if you currently use custom watermark settings and want to start using optimized settings even if no alerts have been triggered.

What you'll need

- You have completed the upgrade to StorageGRID 11.6.
- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the Root access permission.

About this task

You can resolve the **Low read-only watermark override** alert by updating custom watermark settings to the new watermark overrides. However, if one or more Storage Nodes are close to full or you have special ILM requirements, you should first view the optimized storage watermarks and determine if it is safe to use them.

Assess object data usage for entire grid

1. Select **NODES**.
2. For each site in the grid, expand the list of nodes.
3. Review the percentage values shown in the **Object data used** column for each Storage Node at every site.

Nodes

View the list and status of sites and grid nodes.

🔍

Total node count: 13

Name	Type	Object data used	Object metadata used	CPU usage
StorageGRID	Grid	61%	4%	—
⬆ Data Center 1	Site	56%	3%	—
DC1-ADM	Primary Admin Node	—	—	6%
DC1-GW	Gateway Node	—	—	1%
⚠ DC1-SN1	Storage Node	71%	3%	30%
⚠ DC1-SN2	Storage Node	25%	3%	42%
⚠ DC1-SN3	Storage Node	63%	3%	42%
⚠ DC1-SN4	Storage Node	65%	3%	41%

4. If none of the Storage Nodes are close to full (for example, all **Object data used** values are less than 80%), you can start using the override settings. Go to [Use optimized watermarks](#).



There are some exceptions to this general rule. For example, if ILM rules use Strict ingest behavior or if specific storage pools are close to full, you should first perform the steps in [View optimized storage watermarks](#) and [Determine if you can use optimized watermarks](#).

5. If one more more Storage Nodes are close to full, perform the steps in [View optimized storage watermarks](#)

and [Determine if you can use optimized watermarks](#).

View optimized storage watermarks

StorageGRID uses two Prometheus metrics to show the optimized values it has calculated for the **Storage Volume Soft Read-Only Watermark**. You can view the minimum and maximum optimized values for each Storage Node in your grid.

1. Select **SUPPORT > Tools > Metrics**.
2. In the Prometheus section, select the link to access the Prometheus user interface.
3. To see the recommended minimum soft read-only watermark, enter the following Prometheus metric, and select **Execute**:

```
storagegrid_storage_volume_minimum_optimized_soft_readonly_watermark
```

The last column shows the minimum optimized value of the Soft Read-Only Watermark for all storage volumes on each Storage Node. If this value is greater than the custom setting for the **Storage Volume Soft Read-Only Watermark**, the **Low read-only watermark override** alert is triggered for the Storage Node.

4. To see the recommended maximum soft read-only watermark, enter the following Prometheus metric, and select **Execute**:

```
storagegrid_storage_volume_maximum_optimized_soft_readonly_watermark
```

The last column shows the maximum optimized value of the Soft Read-Only Watermark for all storage volumes on each Storage Node.

5. Note the maximum optimized value for each Storage Node.

Determine if you can use optimized watermarks

1. Select **NODES**.
2. Repeat these steps for each online Storage Node:
 - a. Select **Storage Node > Storage**.
 - b. Scroll down to the Object Stores table.
 - c. Compare the **Available** value for each object store (volume) to the maximum optimized watermark you noted for that Storage Node.
3. If at least one volume on every online Storage Node has more space available than maximum optimized watermark for that node, go to [Use optimized watermarks](#) to start using the optimized watermarks.

Otherwise, [expand your grid](#) as soon as possible. Either add storage volumes to an existing node or add new Storage Nodes. Then, go to [Use optimized watermarks](#) to update watermark settings.

4. If you need to continue using custom values for the storage volume watermarks, [silence](#) or [disable](#) the **Low read-only watermark override** alert.



The same custom watermark values are applied to every storage volume on every Storage Node. Using smaller-than-recommended values for storage volume watermarks might cause some storage volumes to become inaccessible (automatically unmounted) when the node reaches capacity.

Use optimized watermarks


1. Go to **CONFIGURATION > System > Storage options**.
2. Select **Configuration** from the Storage Options menu.
3. Change all three Watermark Overrides to 0.
4. Select **Apply Changes**.

Optimized storage volume watermark settings are now in effect for each storage volume, based on the size of the Storage Node and the relative capacity of the volume.

Storage Options

Overview

Configuration



Storage Options Overview

Updated: 2021-11-22 13:57:51 MST

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark Override	0 B
Storage Volume Soft Read-Only Watermark Override	0 B
Storage Volume Hard Read-Only Watermark Override	0 B
Metadata Reserved Space	3,000 GB

Ports

Description	Settings
CLB S3 Port	8082
CLB Swift Port	8083
LDR S3 Port	18082
LDR Swift Port	18083

Troubleshoot the Storage Status (SSTS) alarm

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

What you'll need

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

About this task

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



Storage Options Overview

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

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

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

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

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


Steps

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

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




OverviewAlarmsReportsConfiguration


MainHistory



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

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

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

Apply Changes 

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



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

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

Overview

Alarms

Reports

Configuration

Main

Overview: LDR (.DC1-S1-101-193) - Storage

Updated: 2019-10-09 12:51:07 MDT

Storage State - Desired:

Online

Storage State - Current:

Read-only

Storage Status:

Insufficient Free Space

Utilization

Total Space:

164 GB

Total Usable Space:

19.6 GB

Total Usable Space (Percent):

11.937 %

Total Data:

139 GB

Total Data (Percent):

84.567 %

Replication

Block Reads:

0

Block Writes:

2,279,881

Objects Retrieved:

0

Objects Committed:

88,882

Objects Deleted:

16

Delete Service State:

Enabled

Object Store Volumes

ID	Total	Available	Replicated Data	EC Data	Stored (%)	Health	
0000	54.7 GB	2.93 GB	46.2 GB	0 B	84.486 %	No Errors	
0001	54.7 GB	8.32 GB	46.3 GB	0 B	84.644 %	No Errors	
0002	54.7 GB	8.36 GB	46.3 GB	0 B	84.57 %	No Errors	

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

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

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

Overview


Alarms

Reports

Configuration

Charts

Text



Reports (Charts): LDR (DC1-S1-101-193) - Storage

Attribute:

Total Usable Space

Quick Query:

Custom Query

Update

Vertical Scaling:
☒

Raw Data:
☐


YYYY/MM/DD HH:MM:SS

Start Date:

2019/10/09 12:00:00

End Date:

2019/10/09 13:10:33



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

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

Overview

Alarms

Reports

Configuration

Charts

Text

Reports (Charts): LDR (DC1-S1-101-193) - Storage

Attribute:

Total Usable Space (Percent)

Quick Query:

Custom Query

Update

Vertical Scaling:
☒

Raw Data:
☐


YYYY/MM/DD HH:MM:SS

Start Date:

2019/10/09 12:00:00

End Date:

2019/10/09 13:10:33



6. As required, add storage capacity by [expanding the StorageGRID system](#).

For procedures on how to manage a full Storage Node, see the [instructions for administering StorageGRID](#).

Troubleshoot delivery of platform services messages (SMTT alarm)

The Total Events (SMTT) alarm is triggered in the Grid Manager if a platform service message is delivered to an destination that cannot accept the data.

About this task

For example, an S3 multipart upload can succeed even though the associated replication or notification message cannot be delivered to the configured endpoint. Or, a message for CloudMirror replication can fail to be delivered if the metadata is too long.

The SMTT alarm contains a Last Event message that says, Failed to publish notifications for *bucket-name object key* for the last object whose notification failed.

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

For additional information about troubleshooting platform services, see the [instructions for administering StorageGRID](#). You might need to [access the tenant from the Tenant Manager](#) to debug a platform service error.

Steps

1. To view the alarm, select **NODES** > *site* > *grid node* > **Events**.
2. View Last Event at the top of the table.

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

3. Follow the guidance provided in the SMTT alarm contents to correct the issue.
4. Select **Reset event counts**.
5. Notify the tenant of the objects whose platform services messages have not been delivered.
6. Instruct the tenant to trigger the failed replication or notification by updating the object's metadata or tags.

Troubleshoot metadata issues

You can perform several tasks to help determine the source of metadata problems.

Troubleshoot the Low metadata storage alert


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

What you'll need

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

About this task

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



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

You can [monitor object metadata capacity for each Storage Node](#) to help you anticipate errors and correct them before they occur.

StorageGRID uses the following Prometheus metric to measure how full the allowed metadata space is:

```
storagegrid_storage_utilization_metadata_bytes/storagegrid_storage_utilization_metadata_allowed_bytes
```

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

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



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

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

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

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

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

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



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

Steps

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



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

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

Troubleshoot the Services: Status - Cassandra (SVST) alarm

The Services: Status - Cassandra (SVST) alarm indicates that you might need to rebuild the Cassandra database for a Storage Node. Cassandra is used as the metadata store for StorageGRID.

What you'll need

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

About this task

If Cassandra is stopped for more than 15 days (for example, the Storage Node is powered off), Cassandra will

not start when the node is brought back online. You must rebuild the Cassandra database for the affected DDS service.

You can [run diagnostics](#) to obtain additional information on the current state of your grid.




If two or more of the Cassandra database services are down for more than 15 days, contact technical support, and do not proceed with the steps below.

Steps


1. Select **SUPPORT > Tools > Grid topology**.
2. Select **Site > Storage Node > SSM > Services > Alarms > Main** to display alarms.

This example shows that the SVST alarm was triggered.



Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
Minor	SVST (Services: Status - Cassandra)	Not Running	2014-08-14 14:56:28 PDT	Not Running	Not Running		<input type="checkbox"/>

The SSM Services Main page also indicates that Cassandra is not running.



Service	Version	Status	Threads	Load	Memory
Account Service	10.4.0-20161224.0333.803cd91	Running	7	0.002 %	12 MB
Administrative Domain Controller (ADC)	10.4.0-20170329.0039.8800cae	Running	52	0.14 %	63.1 MB
Cassandra	4.6.12-1.byc.0-20170308.0109.ba3598a	Not Running	0	0 %	0 B
Content Management System (CMS)	10.4.0-20170220.1846.1a76aed	Running	18	0.055 %	20.6 MB
Distributed Data Store (DDS)	10.4.0-20170329.0039.8800cae	Running	104	1.301 %	76 MB
Identity Service	10.4.0-20170203.2038.a457d45	Running	6	0 %	8.75 MB
Keystone Service	10.4.0-20170104.1815.6e52138	Running	5	0 %	7.77 MB
Local Distribution Router (LDR)	10.4.0-20170329.0039.8800cae	Running	109	0.218 %	96.6 MB
Server Manager	10.4.0-20170306.2303.9649faf	Running	4	3.58 %	19.1 MB

3. Try restarting Cassandra from the Storage Node:
 - a. Log in to the grid node:
 - i. Enter the following command: `ssh admin@grid_node_IP`

- ii. Enter the password listed in the `Passwords.txt` file.
 - iii. Enter the following command to switch to root: `su -`
 - iv. Enter the password listed in the `Passwords.txt` file. When you are logged in as root, the prompt changes from `$` to `#`.
- b. Enter: `/etc/init.d/cassandra status`
 - c. If Cassandra is not running, restart it: `/etc/init.d/cassandra restart`
4. If Cassandra does not restart, determine how long Cassandra has been down. If Cassandra has been down for longer than 15 days, you must rebuild the Cassandra database.



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


You can determine how long Cassandra has been down by charting it or by reviewing the `servermanager.log` file.

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

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

Overview
Alarms
Reports
Configuration


Charts
Text


Reports (Charts): SSM (DC1-S3) - Services

Attribute: Services: Status - Cassandra
Quick Query: Last Month Update

Vertical Scalling: ☒
Raw Data: ☐

Start Date: 2014/07/14 16:49:38
End Date: 2014/08/14 16:49:38



6. To review the servermanager.log file on the Storage Node:
 - a. Log in to the grid node:
 - i. Enter the following command: `ssh admin@grid_node_IP`
 - ii. Enter the password listed in the `Passwords.txt` file.
 - iii. Enter the following command to switch to root: `su -`
 - iv. Enter the password listed in the `Passwords.txt` file. When you are logged in as root, the prompt changes from `$` to `#`.

b. Enter: `cat /var/local/log/servermanager.log`

The contents of the servermanager.log file are displayed.

If Cassandra has been down for longer than 15 days, the following message is displayed in the servermanager.log file:

```
"2014-08-14 21:01:35 +0000 | cassandra | cassandra not
started because it has been offline for longer than
its 15 day grace period - rebuild cassandra"
```

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

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

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

For instructions, see [Recover Storage Node down more than 15 days](#).

- e. Contact technical support if alarms do not clear after Cassandra is rebuilt.

Troubleshoot Cassandra Out of Memory errors (SMTT alarm)

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

About this task

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

Steps

1. To view the event, select **SUPPORT > Tools > Grid topology > Configuration**.
2. Verify that the Cassandra Heap Out Of Memory Errors count is 1 or greater.

You can [run diagnostics](#) to obtain additional information on the current state of your grid.

3. Go to `/var/local/core/`, compress the `Cassandra.hprof` file, and send it to technical support.
4. Make a backup of the `Cassandra.hprof` file, and delete it from the `/var/local/core/` directory.

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

5. After the issue is resolved, select the **Reset** check box for the Cassandra Heap Out Of Memory Errors count. Then select **Apply Changes**.



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

Troubleshoot certificate errors

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


About this task

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

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

- Your custom management interface certificate expires.
- You revert from a custom management interface certificate to the default server certificate.

The following example shows a certificate error when the custom management interface certificate expired:



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

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

Steps

If you received an alert notification about an expired certificate, access the certificate details: . Select **CONFIGURATION > Security > Certificates** and then [select the appropriate certificate tab](#).

1. Check the validity period of the certificate.
Some web browsers and S3 or Swift clients do not accept certificates with a validity period greater than 398 days.
2. If the certificate has expired or will expire soon, upload or generate a new certificate.
 - For a server certificate, see the steps for [configuring a custom server certificate for the Grid Manager and the Tenant Manager](#).
 - For a client certificate, see the steps for [configuring a client certificate](#).
3. For server certificate errors, try either or both of the following options:

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

Troubleshoot Admin Node and user interface issues

There are several tasks you can perform to help determine the source of issues related to Admin Nodes and the StorageGRID user interface.

Troubleshoot sign-on errors

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

What you'll need

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

About this task

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

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

Steps

1. Wait 10 minutes, and try signing in again.

If the error is not resolved automatically, go to the next step.

2. If your StorageGRID system has more than one Admin Node, try signing in to the Grid Manager from another Admin Node.
 - If you are able to sign in, you can use the **Dashboard**, **NODES**, **Alerts**, and **SUPPORT** options to help determine the cause of the error.
 - If you have only one Admin Node or you still cannot sign in, go to the next step.

3. Determine if the node's hardware is offline.
4. If single sign-on (SSO) is enabled for your StorageGRID system, refer to the steps for configuring single sign-on, in the instructions for administering StorageGRID.

You might need to temporarily disable and re-enable SSO for a single Admin Node to resolve any issues.



If SSO is enabled, you cannot sign on using a restricted port. You must use port 443.

5. Determine if the account you are using belongs to a federated user.

If the federated user account is not working, try signing in to the Grid Manager as a local user, such as root.

- If the local user can sign in:
 - i. Review any displayed alarms.
 - ii. Select **CONFIGURATION > Access Control > Identity federation**.
 - iii. Click **Test Connection** to validate your connection settings for the LDAP server.
 - iv. If the test fails, resolve any configuration errors.
- If the local user cannot sign in and you are confident that the credentials are correct, go to the next step.

6. Use Secure Shell (ssh) to log in to the Admin Node:

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

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

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

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

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

```

$ storagegrid-status
Host Name                99-211
IP Address                10.96.99.211
Operating System Kernel  4.19.0                Verified
Operating System Environment Debian 10.1            Verified
StorageGRID Webscale Release 11.4.0                Verified
Networking                Verified
Storage Subsystem        Verified
Database Engine          5.5.9999+default Running
Network Monitoring       11.4.0                Running
Time Synchronization     1:4.2.8p10+dfsg Running
ams                      11.4.0                Running
cmn                      11.4.0                Running
nms                      11.4.0                Running
ssm                      11.4.0                Running
mi                      11.4.0                Running
dynip                   11.4.0                Running
nginx                   1.10.3                Running
tomcat                  9.0.27                Running
grafana                 6.4.3                Running
mgmt api                11.4.0                Running
prometheus              11.4.0                Running
persistence             11.4.0                Running
ade exporter            11.4.0                Running
alertmanager            11.4.0                Running
attrDownPurge           11.4.0                Running
attrDownSamp1           11.4.0                Running
attrDownSamp2           11.4.0                Running
node exporter           0.17.0+ds             Running
sg snmp agent           11.4.0                Running

```

8. Confirm that the nginx-gw service is running # `service nginx-gw status`

9. Use Lumberjack to collect logs: # `/usr/local/sbin/lumberjack.rb`

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

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

10. Review the following logs:

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

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

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

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

Admin Nodes use the ADC service during the authentication process.

12. From the Admin Node, log in to each of the ADC Storage Nodes, using the IP addresses you identified.
- Enter the following command: `ssh admin@grid_node_IP`
 - Enter the password listed in the `Passwords.txt` file.
 - Enter the following command to switch to root: `su -`
 - Enter the password listed in the `Passwords.txt` file.

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

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

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

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

Provide the logs you collected to technical support. See also [Log files reference](#).

Troubleshoot user interface issues

You might see issues with the Grid Manager or the Tenant Manager after upgrading to a new version of StorageGRID software.

Web interface does not respond as expected

The Grid Manager or the Tenant Manager might not respond as expected after StorageGRID software is upgraded.

If you experience issues with the web interface:

- Make sure you are using a [supported web browser](#).



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

- Clear your web browser cache.

Clearing the cache removes outdated resources used by the previous version of StorageGRID software,

and permits the user interface to operate correctly again. For instructions, see the documentation for your web browser.

Check the status of an unavailable Admin Node

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

What you'll need


You must have specific access permissions.

Steps

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

OverviewAlarmsReportsConfiguration

Main



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

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

Operating System:

Linux
3.16.0-4-amd64

Services

Service	Version	Status	Threads	Load	Memory
Audit Management System (AMS)	10.4.0-20170113.2207.3ec2cd0	Running	52	0.043 %	35.7 MB
CIFS Filesharing (nmbd)	2.4.2.14+dfsg-0+deb8u2	Running	1	0 %	5.5 MB
CIFS Filesharing (smbd)	2.4.2.14+dfsg-0+deb8u2	Running	1	0 %	14.5 MB
CIFS Filesharing (winbindd)	2.4.2.14+dfsg-0+deb8u2	Not Running	0	0 %	0 B
Configuration Management Node (CMN)	10.4.0-20170113.2207.3ec2cd0	Running	52	0.055 %	41.3 MB
Database Engine	5.5.53-0+deb8u1	Running	47	0.354 %	1.33 GB
Grid Deployment Utility Server	10.4.0-20170112.2125.c4253bb	Running	3	0 %	32.8 MB
Management Application Program Interface (mgmt-api)	10.4.0-20170113.2136.07c4997	Not Running	0	0 %	0 B
NFS Filesharing	10.4.0-20161224.0333.803cd91	Not Running	0	0 %	0 B
NMS Data Cleanup	10.4.0-20161224.0333.803cd91	Running	22	0.008 %	52.4 MB
NMS Data Downsampler 1	10.4.0-20161224.0333.803cd91	Running	22	0.049 %	195 MB
NMS Data Downsampler 2	10.4.0-20161224.0333.803cd91	Running	22	0.009 %	157 MB
NMS Processing Engine	10.4.0-20161224.0333.803cd91	Running	40	0.132 %	200 MB

5. Determine if alarms have been triggered.

6. Take the appropriate actions to resolve the issue.

Related information

[Administer StorageGRID](#)

Troubleshoot network, hardware, and platform issues

There are several tasks you can perform to help determine the source of issues related to StorageGRID network, hardware, and platform issues.

Troubleshoot “422: Unprocessable Entity” errors

The error 422: Unprocessable Entity can occur in a number of circumstances. Check the error message to determine what caused your issue.

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

Error message	Root cause and corrective action
<pre>422: Unprocessable Entity Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration. Unable to authenticate, please verify your username and password: LDAP Result Code 8 "Strong Auth Required": 00002028: LdapErr: DSID-0C090256, comment: The server requires binds to turn on integrity checking if SSL\TLS are not already active on the connection, data 0, v3839</pre>	<p>This message might occur if you select the Do not use TLS option for Transport Layer Security (TLS) when configuring identity federation using Windows Active Directory (AD).</p> <p>Using the Do not use TLS option is not supported for use with AD servers that enforce LDAP signing. You must select either the Use STARTTLS option or the Use LDAPS option for TLS.</p>

Error message	Root cause and corrective action
<pre>422: Unprocessable Entity Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration.Unable to begin TLS, verify your certificate and TLS configuration: LDAP Result Code 200 "Network Error": TLS handshake failed (EOF)</pre>	<p>This message appears if you try to use an unsupported cipher to make a Transport Layer Security (TLS) connection from StorageGRID to an external system used for identify federation or Cloud Storage Pools.</p> <p>Check the ciphers that are offered by the external system. The system must use one of the ciphers supported by StorageGRID for outgoing TLS connections, as shown in the instructions for administering StorageGRID.</p>

Related information

[Administer StorageGRID](#)

Troubleshoot the Grid Network MTU mismatch alert

The **Grid Network MTU mismatch** alert is triggered when the maximum transmission unit (MTU) setting for the Grid Network interface (eth0) differs significantly across nodes in the grid.

About this task

The differences in MTU settings could indicate that some, but not all, eth0 networks are configured for jumbo frames. An MTU size mismatch of greater than 1000 might cause network performance problems.

Steps

- List the MTU settings for eth0 on all nodes.
 - Use the query provided in the Grid Manager.
 - Navigate to *primary Admin Node IP address/metrics/graph* and enter the following query: `node_network_mtu_bytes{interface='eth0'}`
- Modify the MTU settings as necessary to ensure they are the same for the Grid Network interface (eth0) on all nodes.
 - For appliance nodes, see the installation and maintenance instructions for your appliance.
 - For Linux- and VMware-based nodes, use the following command: `/usr/sbin/change-ip.py [-h] [-n node] mtu network [network...]`

Example: `change-ip.py -n node 1500 grid admin`

Note: On Linux-based nodes, if the desired MTU value for the network in the container exceeds the value already configured on the host interface, you must first configure the host interface to have the desired MTU value, and then use the `change-ip.py` script to change the MTU value of the network in the container.

Use the following arguments for modifying the MTU on Linux- or VMware-based nodes.

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

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

Related information

[SG100 and SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

Troubleshoot the Network Receive Error (NRER) alarm

Network Receive Error (NRER) alarms can be caused by connectivity issues between StorageGRID and your network hardware. In some cases, NRER errors can clear without manual intervention. If the errors do not clear, take the recommended actions.

About this task

NRER alarms can be caused by the following issues with networking hardware that connects to StorageGRID:

- Forward error correction (FEC) is required and not in use
- Switch port and NIC MTU mismatch
- High link error rates
- NIC ring buffer overrun

Steps

1. Follow the troubleshooting steps for all potential causes of the NRER alarm given your network configuration.
 - If the error is caused by FEC mismatch, perform the following steps:

Note: These steps are applicable only for NRER errors caused by FEC mismatch on StorageGRID appliances.

- i. Check the FEC status of the port in the switch attached to your StorageGRID appliance.
- ii. Check the physical integrity of the cables from the appliance to the switch.
- iii. If you want to change FEC settings to try to resolve the NRER alarm, first ensure that the appliance is configured for **Auto** mode on the Link Configuration page of the StorageGRID Appliance Installer (see the installation and maintenance instructions for your appliance). Then, change the FEC settings on the switch ports. The StorageGRID appliance ports will adjust their FEC settings to match, if possible.

(You cannot configure FEC settings on StorageGRID appliances. Instead, the appliances attempt to discover and mirror the FEC settings on the switch ports they are connected to. If the links are forced to 25-GbE or 100-GbE network speeds, the switch and NIC might fail to negotiate a common FEC setting. Without a common FEC setting, the network will fall back to “no-FEC” mode. When FEC is not enabled, the connections are more susceptible to errors caused by electrical noise.)

Note: StorageGRID appliances support Firecode (FC) and Reed Solomon (RS) FEC, as well as no FEC.

- If the error is caused by a switch port and NIC MTU mismatch, check that the MTU size configured on the node is the same as the MTU setting for the switch port.

The MTU size configured on the node might be smaller than the setting on the switch port the node is connected to. If a StorageGRID node receives an Ethernet frame larger than its MTU, which is possible with this configuration, the NRER alarm might be reported. If you believe this is what is happening, either change the MTU of the switch port to match the StorageGRID network interface MTU, or change the MTU of the StorageGRID network interface to match the switch port, depending on your end-to-end MTU goals or requirements.



For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The **Grid Network MTU mismatch** alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values do not have to be the same for all network types.



To change the MTU setting, see the installation and maintenance guide for your appliance.

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



You might notice high error rates in an environment with high electrical noise.

- If the error is a NIC ring buffer overrun, contact technical support.

The ring buffer can be overrun when the StorageGRID system is overloaded and unable to process network events in a timely manner.

2. After you resolve the underlying problem, reset the error counter.

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

Related information

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

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

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[SG100 and SG1000 services appliances](#)

Troubleshoot time synchronization errors

You might see issues with time synchronization in your grid.

If you encounter time synchronization problems, verify that you have specified at least four external NTP sources, each providing a Stratum 3 or better reference, and that all external NTP sources are operating normally and are accessible by your StorageGRID nodes.



When specifying the external NTP source for a production-level StorageGRID installation, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

Related information

[Recover and maintain](#)

Linux: Network connectivity issues

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

MAC address cloning

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



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

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

Promiscuous mode

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

Related information

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

[Install Ubuntu or Debian](#)

Linux: Node status is “orphaned”

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

About this task

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

- Check logs for errors and messages.
- Attempt to start the node again.
- If necessary, use container engine commands to stop the existing node container.
- Restart the node.

Steps

1. Check logs for both the service daemon and the orphaned node for obvious errors or messages about exiting unexpectedly.
2. Log in to the host as root or using an account with sudo permission.
3. Attempt to start the node again by running the following command: `$ sudo storagegrid node start node-name`

```
$ sudo storagegrid node start DC1-S1-172-16-1-172
```

If the node is orphaned, the response is

```
Not starting ORPHANED node DC1-S1-172-16-1-172
```

4. From Linux, stop the container engine and any controlling storagegrid-node processes. For example: `sudo docker stop --time seconds container-name`

For `seconds`, enter the number of seconds you want to wait for the container to stop (typically 15 minutes or less). For example:

```
sudo docker stop --time 900 storagegrid-DC1-S1-172-16-1-172
```

5. Restart the node: `storagegrid node start node-name`

```
storagegrid node start DC1-S1-172-16-1-172
```

Linux: Troubleshoot IPv6 support

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

About this task

You can see the IPv6 address that has been assigned to a grid node in the following locations in the Grid Manager:

- Select **NODES**, and select the node. Then, select **Show more** next to **IP Addresses** on the Overview tab.

DC1-S2 (Storage Node)

Overview Hardware Network Storage Objects ILM Tasks

Node information

Name: DC1-S2

Type: Storage Node

ID: 352bd978-ff3e-45c5-aac1-24c7278206fa

Connection state: ✓ Connected

Storage used:

Object data	0%
Object metadata	0%

Software version: 11.6.0 (build 20210924.1557.00a5eb9)

IP addresses:

172.16.1.227 - eth0 (Grid Network)

10.224.1.227 - eth1 (Admin Network)

[Hide additional IP addresses](#)

Interface	IP address
eth0 (Grid Network)	172.16.1.227
eth0 (Grid Network)	fd20:328:328:0:250:56ff:fe87:b532

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

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

Steps

1. Log in to the host as root or using an account with sudo permission.

2. Run the following command: `sysctl net.ipv6.conf.all.disable_ipv6`

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

The result should be 0.

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



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

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

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

The result should be 1.

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



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

5. Exit the container: `exit`

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

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

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

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

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

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

8. Run these commands to restart the StorageGRID container:

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

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

Troubleshoot an external syslog server

The following table describes external syslog server error messages and lists corrective actions.

Error message	Description and recommended actions
Cannot resolve hostname	<p>The FQDN you entered for the syslog server could not be resolved to an IP address.</p> <ol style="list-style-type: none">1. Check the hostname you entered. If you entered an IP address, make sure it is a valid IP address in W.X.Y.Z ("dotted decimal") notation.2. Check that the DNS servers are configured correctly.3. Confirm that each node can access the IP addresses for the DNS server.
Connection refused	<p>A TCP or TLS connection to the syslog server was refused. There might be no service listening on the TCP or TLS port for the host, or a firewall might be blocking access.</p> <ol style="list-style-type: none">1. Check that you entered the correct FQDN or IP address, port, and protocol for the syslog server.2. Confirm that the host for the syslog service is running a syslog daemon that is listening on the specified port.3. Confirm that a firewall is not blocking access to TCP/TLS connections from the nodes to the IP and port of the syslog server.
Network unreachable	<p>The syslog server is not on a directly attached subnet. A router returned an ICMP failure message to indicate it could not forward the test messages from the listed nodes to the syslog server.</p> <ol style="list-style-type: none">1. Check that you entered the correct FQDN or IP address for the syslog server.2. For each node listed, check the Grid Network Subnet List, the Admin Networks Subnet Lists, and the Client Network gateways. Confirm these are configured to route traffic to the syslog server over the expected network interface and gateway (Grid, Admin, or Client).

Error message	Description and recommended actions
Host unreachable	<p>The syslog server is on a directly attached subnet (subnet used by the listed nodes for their Grid, Admin, or Client IP addresses). The nodes attempted to send test messages, but did not receive responses to ARP requests for the syslog server's MAC address.</p> <ol style="list-style-type: none"> 1. Check that you entered the correct FQDN or IP address for the syslog server. 2. Check that the host running the syslog service is up.
Connection timed out	<p>A TCP/TLS connection attempt was made, but no response was received from the syslog server for a long time. There might be a routing misconfiguration or a firewall might be dropping traffic without sending any response (a common configuration).</p> <ol style="list-style-type: none"> 1. Check that you entered the correct FQDN or IP address for the syslog server. 2. For each node listed, check the Grid Network Subnet List, the Admin Networks Subnet Lists, and the Client Network gateways. Confirm these are configured to route traffic to the syslog server via the network interface and gateway (Grid, Admin, or Client) over which you expect the syslog server to be reached. 3. Confirm that a firewall is not blocking access to TCP/TLS connections from the nodes listed to the IP and port of the syslog server.
Connection closed by partner	<p>A TCP connection to the syslog server was successfully established but was later closed. Reasons for this might include:</p> <ul style="list-style-type: none"> • The syslog server might have been restarted or rebooted. • The node and the syslog server might have different TCP/TLS settings. • An intermediate firewall might be closing idle TCP connections. • A non-syslog server listening on the syslog server port might have closed the connection. <ol style="list-style-type: none"> 1. Check that you entered the correct FQDN or IP address, port, and protocol for the syslog server. 2. If you are using TLS, confirm the syslog server is also using TLS. If you are using TCP, confirm the syslog server is also using TCP. 3. Check that an intermediate firewall is not configured to close idle TCP connections.

Error message	Description and recommended actions
TLS certificate error	<p>The server certificate received from the syslog server was not compatible with the CA certificate bundle and client certificate you provided.</p> <ol style="list-style-type: none"> 1. Confirm that the CA certificate bundle and client certificate (if any) are compatible with the server certificate on the syslog server. 2. Confirm that the identities in the server certificate from the syslog server include the expected IP or FQDN values.
Forwarding suspended	<p>Syslog records are no longer being forwarded to the syslog server and StorageGRID is unable to detect the reason.</p> <p>Review the debugging logs provided with this error to attempt to determine the root cause.</p>
TLS session terminated	<p>The syslog server terminated the TLS session and StorageGRID is unable to detect the reason.</p> <ol style="list-style-type: none"> 1. Review the debugging logs provided with this error to attempt to determine the root cause. 2. Check that you entered the correct FQDN or IP address, port, and protocol for the syslog server. 3. If you are using TLS, confirm the syslog server is also using TLS. If you are using TCP, confirm the syslog server is also using TCP. 4. Confirm that the CA certificate bundle and client certificate (if any) are compatible with the server certificate from the syslog server. 5. Confirm that the identities in the server certificate from the syslog server include the expected IP or FQDN values.
Results query failed	<p>The Admin Node used for syslog server configuration and testing is unable to request test results from the nodes listed. One or more nodes might be down.</p> <ol style="list-style-type: none"> 1. Follow standard troubleshooting steps to ensure that the nodes are online and all expected services are running. 2. Restart the miscd service on the nodes listed.

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—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.