



Troubleshooting object and storage issues

StorageGRID

NetApp

October 03, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-115/troubleshoot/verifying-object-integrity.html> on October 03, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Troubleshooting object and storage issues	1
Confirming object data locations	1
Object store (storage volume) failures	3
Verifying object integrity	4
What background verification is	4
Changing the background verification rate	5
What foreground verification is	7
Running foreground verification	8
Troubleshooting lost and missing object data	10
Investigating lost objects	11
Searching for and restoring potentially lost objects	16
Resetting lost and missing object counts	22
Troubleshooting the Low object data storage alert	23
Troubleshooting the Storage Status (SSTS) alarm	25
Troubleshooting delivery of platform services messages (SMTT alarm)	29

Troubleshooting object and storage issues

There are several tasks you can perform to help determine the source of object and storage issues.

Confirming object data locations

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

What you'll need

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

Steps

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

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

Object Metadata Lookup

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

Identifier	<input type="text" value="source/testobject"/>	Look Up
------------	--	----------------

3. Click **Look Up**.

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

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

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

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

System Metadata

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

Replicated Copies

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

Raw Metadata

```
{
  "TYPE": "CTNT",
  "CHND": "A12E96FF-B13F-4905-9E9E-45373F6E7DA8",
  "NAME": "testobject",
  "CBID": "0x8823DE7EC7C10416",
  "PHND": "FEA0AE51-534A-11EA-9FCD-31FF00C36056",
  "PPTH": "source",
  "META": {
    "BASE": {
      "PAWS": "2",
      "PAWS2": "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. These object stores are physical partitions that act as mount points for StorageGRID system's storage. Object stores are also known as storage volumes.

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

Disk Devices						
Name	World Wide Name	I/O Load	Read Rate	Write Rate		
croot(8:1,sda1)	N/A	1.62%	0 bytes/s	177 KB/s		
cvloc(8:2,sda2)	N/A	17.28%	0 bytes/s	2 MB/s		
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	11 KB/s		
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	0 bytes/s		
sds(8:48,sdd)	N/A	0.00%	0 bytes/s	0 bytes/s		

Volumes						
Mount Point	Device	Status	Size	Available	Write Cache Status	
/	croot	Online	21.00 GB	14.25 GB		Unknown
/var/local	cvloc	Online	85.86 GB	84.39 GB		Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.18 GB		Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB		Enabled
/var/local/rangedb/2	sds	Online	107.32 GB	107.18 GB		Enabled

Object Stores						
ID	Size	Available	Replicated Data	EC Data	Object Data (%)	Health
0000	107.32 GB	96.45 GB	 994.37 KB	 0 bytes	 0.00%	No Errors
0001	107.32 GB	107.18 GB	 0 bytes	 0 bytes	 0.00%	No Errors
0002	107.32 GB	107.18 GB	 0 bytes	 0 bytes	 0.00%	No Errors

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

1. Select **Support > Tools > Grid Topology**.
2. Select **site > Storage Node > LDR > Storage > Overview > Main**.



Storage State - Desired:
Storage State - Current:
Storage Status:

Online

Online

No Errors



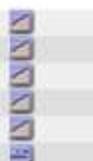
Utilization

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



Replication

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



Object Store Volumes

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

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

Related information

[Maintain & recover](#)

Verifying object integrity

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

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

What background verification is

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

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

- **Replicated objects:** If the background verification process finds a replicated object that is corrupt, the corrupt copy is removed from its location and quarantined elsewhere on the Storage Node. Then, a new uncorrupted copy is generated and placed to satisfy the active ILM policy. The new copy might not be placed on the Storage Node that was used for the original copy.
- **Erasure-coded objects:** If the background verification process detects that a fragment of an erasure-coded object is corrupt, StorageGRID automatically attempts to rebuild the missing fragment in place on the same Storage Node, using the remaining data and parity fragments. If the corrupted fragment cannot be rebuilt, the Corrupt Copies Detected (ECOR) attribute is incremented by one, and an attempt is made to retrieve another copy of the object. If retrieval is successful, an ILM evaluation is performed to create a replacement copy of the erasure-coded object.

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

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

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

Alerts and alarms (legacy) related to background verification

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

If background verification cannot replace a corrupted object because it cannot locate another copy, the **Objects lost** alert and the **LOST (Lost Objects)** legacy alarm are triggered.

Changing the background verification rate

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

What you'll need

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

About this task

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

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

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

Steps

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

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

Reset Missing Objects Count

Foreground Verification

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

Background Verification

Verification Rate: Adaptive

Reset Corrupt Objects Count

Quarantined Objects

Delete Quarantined Objects

Apply Changes 



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

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

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

- If the object identifier can be extracted, StorageGRID automatically creates a new copy of the object data. The new copy can be made anywhere in the StorageGRID system that satisfies the active ILM policy.

- If the object identifier cannot be extracted (because it has been corrupted), the **Corrupt Objects Unidentified** metric is incremented, and the **Unidentified corrupt object detected** alert is triggered.
- 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.

- Select **Support > Tools > Grid Topology**.
- Select **Storage Node > LDR > Erasure Coding**.
- 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.

- Select **Support > Tools > Grid Topology**.
- Select **Storage Node > LDR > Verification > Configuration**.
- Select **Reset Corrupt Object Count**.
- 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.

- Select **Support > Tools > Grid Topology**.
- Select **Storage Node > LDR > Verification > Configuration**.
- Select **Delete Quarantined Objects**.
- Click **Apply Changes**.

What foreground verification is

Foreground verification is a user-initiated process that checks if all expected object data exists on a Storage Node. Foreground verification is used to verify the integrity of a storage device.

Foreground verification is a faster alternative to background verification that checks the existence, but not the integrity, of object data on a Storage Node. If foreground verification finds that many items are missing, there might be an issue with all or part of a storage device associated with the Storage Node.

Foreground verification checks both replicated object data and erasure-coded object data, as follows:

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

- **Erasure-coded objects:** If a fragment of an erasure-coded object is found to be missing, StorageGRID automatically attempts to rebuild the missing fragment in place on the same Storage Node using the remaining fragments. If the missing fragment cannot be rebuilt (because too many fragments have been lost), the Corrupt Copies Detected (ECOR) attribute is incremented by one. ILM then attempts to find another copy of the object, which it can use to generate a new erasure-coded copy.

If foreground verification identifies an issue with erasure coding on a storage volume, the foreground verification task pauses with an error message that identifies the affected volume. You must perform a recovery procedure for any affected storage volumes.

If no other copies of a missing replicated object or a corrupted erasure-coded object can be found in the grid, the **Objects lost** alert and the LOST (Lost Objects) legacy alarm are triggered.

Running foreground verification

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

What you'll need

- You have ensured that the following grid tasks are not running:
 - Grid Expansion: Add Server (GEXP), when adding a Storage Node
 - Storage Node Decommissioning (LDCM) on the same Storage Node If these grid tasks are running, wait for them to complete or release their lock.
- You have ensured that the storage is online. (Select **Support > Tools > Grid Topology**. Then, select **Storage Node > LDR > Storage > Overview > Main**. Ensure that **Storage State - Current** is **Online**.)
- You have ensured that the following recovery procedures are not running on the same Storage Node:
 - Recovery of a failed storage volume
 - Recovery of a Storage Node with a failed system drive Foreground verification does not provide useful information while recovery procedures are in progress.

About this task

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

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

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

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

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

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

Steps

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

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

Background Verification

Verification Rate: Adaptive

Reset Corrupt Objects Count:

Apply Changes 

4. Click **Apply Changes**.

Wait until the page auto-refreshes and reloads before you leave the page. Once refreshed, object stores become unavailable for selection on that Storage Node.

An LDR Foreground Verification grid task is generated and runs until it completes, pauses, or is aborted.

5. Monitor missing objects or missing fragments:

- a. Select **Storage Node > LDR > Verification**.
- b. On the Overview tab under **Verification Results**, note the value of **Missing Objects Detected**.

Note: The same value is reported as **Lost Objects** on the Nodes page. Go to **Nodes > Storage Node**, and select the **Objects** tab.

If the number of **Missing Objects Detected** is large (if there are a hundreds of missing objects), there

is likely an issue with the Storage Node's storage. Contact technical support.

c. Select **Storage Node > LDR > Erasure Coding**.

d. On the Overview tab under **Verification Results**, note the value of **Missing Fragments Detected**.

If the number of **Missing Fragments Detected** is large (if there are a hundreds of missing fragments), there is likely an issue with the Storage Node's storage. Contact technical support.

If foreground verification does not detect a significant number of missing replicated object copies or a significant number of missing fragments, then the storage is operating normally.

6. Monitor the completion of the foreground verification grid task:

a. Select **Support > Tools > Grid Topology**. Then select **site > Admin Node > CMN > Grid Task > Overview > Main**.

b. Verify that the foreground verification grid task is progressing without errors.

Note: A notice-level alarm is triggered on grid task status (SCAS) if the foreground verification grid task pauses.

c. If the grid task pauses with a critical storage error, recover the affected volume and then run foreground verification on the remaining volumes to check for additional errors.

Attention: If the foreground verification grid task pauses with the message Encountered a critical storage error in volume *volID*, you must perform the procedure for recovering a failed storage volume. See the recovery and maintenance instructions.

After you finish

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

Related information

[Maintain & recover](#)

Troubleshooting lost and missing object data

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

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

If this retrieval is successful, the StorageGRID system replaces the missing copy of the object. Otherwise, the **Objects lost** alert and the legacy LOST (Lost Objects) alarm are triggered, as follows:

- For replicated copies, if another copy cannot be retrieved, the object is considered lost, and the alert and alarm are triggered.

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

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

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

Related information

[Investigating lost objects](#)

[Resetting lost and missing object counts](#)

Investigating lost objects

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

What you'll need

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

About this task

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

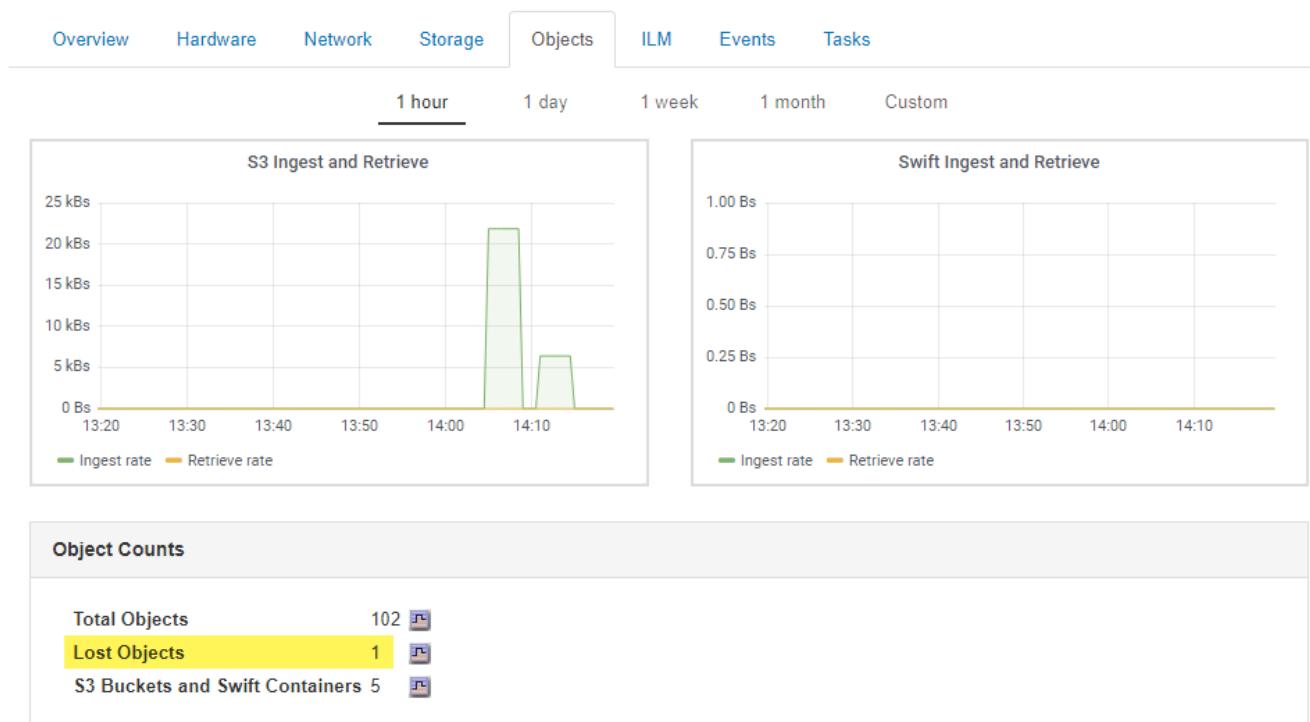
Investigate lost object alarms or alerts immediately. You might need to take action to prevent further data loss. In some cases, you might be able to restore a lost object if you take prompt action.

The number of Lost Objects can be seen in the Grid Manager.

Steps

1. Select **Nodes**.
2. Select **Storage Node > Objects**.
3. Review the number of Lost Objects shown in the Object Counts table.

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



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

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

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

at risk.

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

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

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

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

```

```
    "Location Type": "CLDI\ (Location online\ )",
    "NOID\ (Node ID\ )": "12448208",
    "VOLI\ (Volume ID\ )": "3222345473",
    "Object File Path":
    "/var/local/rangedb/1/p/17/11/00rH0%DkRt78Il1a\#3udu",
    "LTIM\ (Location timestamp\ )": "2020-02-
12T19:36:17.880569"
    \},
    \{
        "Location Type": "CLDI\ (Location online\ )",
        "NOID\ (Node ID\ )": "12288733",
        "VOLI\ (Volume ID\ )": "3222345984",
        "Object File Path":
        "/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb\#3s;L",
        "LTIM\ (Location timestamp\ )": "2020-02-
12T19:36:17.934425"
    }
]
```

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

```

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

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

```

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

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

Related information

[Maintain & recover](#)

[Review audit logs](#)

Searching for and restoring potentially lost objects

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

What you'll need

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

About this task

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



Contact technical support for assistance with this procedure.

Steps

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

For example:

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

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

For example:

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

An LLST audit message looks like this sample message.

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

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

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

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

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

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

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

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

a. Log in to the grid node:

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

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

b. Determine if the file path for the object exists.

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

For example, enter:

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

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

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

3. If the object path was found, attempt to restore the object to StorageGRID:

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

For example, enter:

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

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

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

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

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

ade 12448208: /proc/STOR > Object found succeeded.
First packet of file was valid. Extracted key: 38186FE53E3C49A5
Renamed '/var/local/rangedb/1/p/17/11/00rH0%DkRs&LgA%#3tN6' to
'/var/local/rangedb/1/p/17/11/00rH0%DkRt78I1a#3udu'
```

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

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

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

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

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

```

    "*ctp(HTTP content MIME type)": "binary/octet-stream"
},
"BYCB(System metadata)": {
    "CSIZ(Plaintext object size)": "5242880",
    "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
    "BSIZ(Content block size)": "5252084",
    "CVER(Content block version)": "196612",
    "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
    "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
    "ITME": "1581534970983000"
},
"CMSM": {
    "LATM(Object last access time)": "2020-02-
12T19:16:10.983000"
},
"AWS3": {
    "LOCC": "us-east-1"
},
"CLCO\ (Locations\)": \[
    \{
        "Location Type": "CLDI\ (Location online\)",
        "NOID\ (Node ID)": "12448208",
        "VOLI\ (Volume ID)": "3222345473",
        "Object File Path":
        "/var/local/rangedb/1/p/17/11/00rH0%DkRt78Il1a\#3udu",
        "LTIM\ (Location timestamp)": "2020-02-
12T19:36:17.880569"
    },
    \{
        "Location Type": "CLDI\ (Location online\)",
        "NOID\ (Node ID)": "12288733",
        "VOLI\ (Volume ID)": "3222345984",
        "Object File Path":
        "/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb\#3s;L",
        "LTIM\ (Location timestamp)": "2020-02-
12T19:36:17.934425"
    }
]
}

```

c. Sign out of the LDR console. Enter: exit

5. From an Admin Node, search the audit logs for the ORLM audit message for this object to confirm that information lifecycle management (ILM) has placed copies as required.
 - a. Log in to the grid node:
 - i. Enter the following command: `ssh admin@grid_node_IP`
 - ii. Enter the password listed in the `Passwords.txt` file.
 - iii. Enter the following command to switch to root: `su -`
 - iv. Enter the password listed in the `Passwords.txt` file. When you are logged in as root, the prompt changes from `$` to `#`.
 - b. Change to the directory where the audit logs are located: `cd /var/local/audit/export/`
 - c. Use grep to extract the audit messages associated with the object to an output file. Enter: `grep uuid-valueaudit_file_name > output_file_name`

For example:

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

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

For example:

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

An ORLM audit message looks like this sample message.

```
[AUDT:[CBID(UI64):0x38186FE53E3C49A5][RULE(CSTR):"Make 2 Copies"]
[STAT(FC32):DONE][CSIZ(UI64):0][UUID(CSTR):"926026C4-00A4-449B-AC72-
BCCA72DD1311"]
[LOCS(CSTR):"**CLDI 12828634 2148730112**, CLDI 12745543 2147552014"]
[RSLT(FC32):SUCS][AVER(UI32):10][ATYP(FC32):ORLM][ATIM(UI64):15633982
30669]
[ATID(UI64):15494889725796157557][ANID(UI32):13100453][AMID(FC32):BCM
S]]
```

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

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

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

Related information

[Investigating lost objects](#)

[Confirming object data locations](#)

[Resetting lost and missing object counts](#)

[Review audit logs](#)

Resetting lost and missing object counts

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

What you'll need

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

About this task

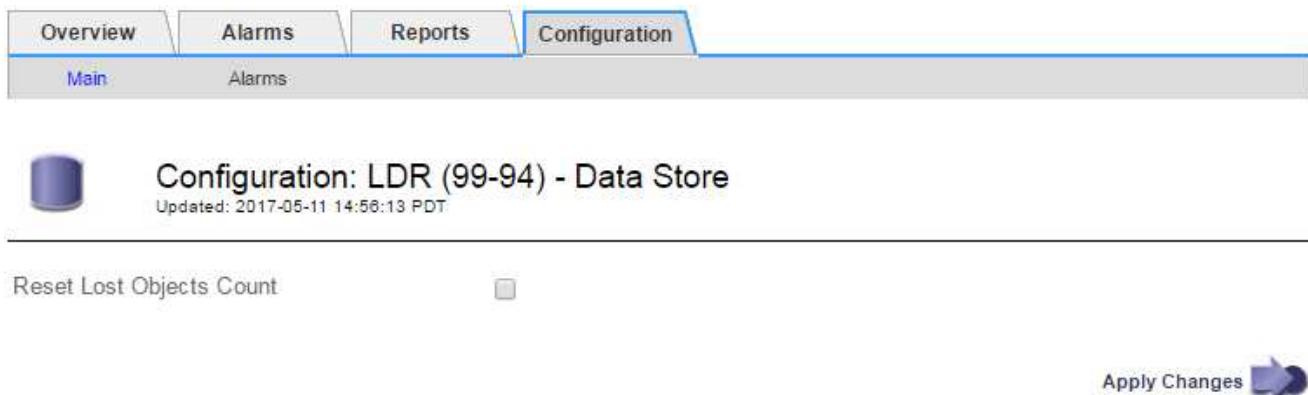
You can reset the Lost Objects counter from either of the following pages:

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

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

Steps

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



The screenshot shows a navigation bar with tabs: Overview, Alarms, Reports, and Configuration. The Configuration tab is highlighted with a blue border. Below the navigation bar, there is a sub-navigation with Main and Alarms. The main content area has a title 'Configuration: LDR (99-94) - Data Store' and a subtitle 'Updated: 2017-05-11 14:56:13 PDT'. At the bottom of the main content area, there is a button labeled 'Reset Lost Objects Count' with a blue border. To the right of this button is a small square icon. At the bottom right of the page, there is a button labeled 'Apply Changes' with a blue arrow icon.

4. Click **Apply Changes**.

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

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

- a. Select **Site > Storage Node > LDR > Erasure Coding > Configuration**.
- b. Select **Reset Reads Failure Count** and **Reset Corrupt Copies Detected Count**.
- c. Click **Apply Changes**.
- d. Select **Site > Storage Node > LDR > Verification > Configuration**.
- e. Select **Reset Missing Objects Count** and **Reset Corrupt Objects Count**.
- f. If you are confident that quarantined objects are not required, you can select **Delete Quarantined Objects**.

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

- g. Click **Apply Changes**.

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

Related information

[Administer StorageGRID](#)

Troubleshooting the Low object data storage alert

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

What you'll need

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

About this task

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

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

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

In this condition:

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

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

Steps

1. Select **Alerts > Current**.

The Alerts page appears.

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



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

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

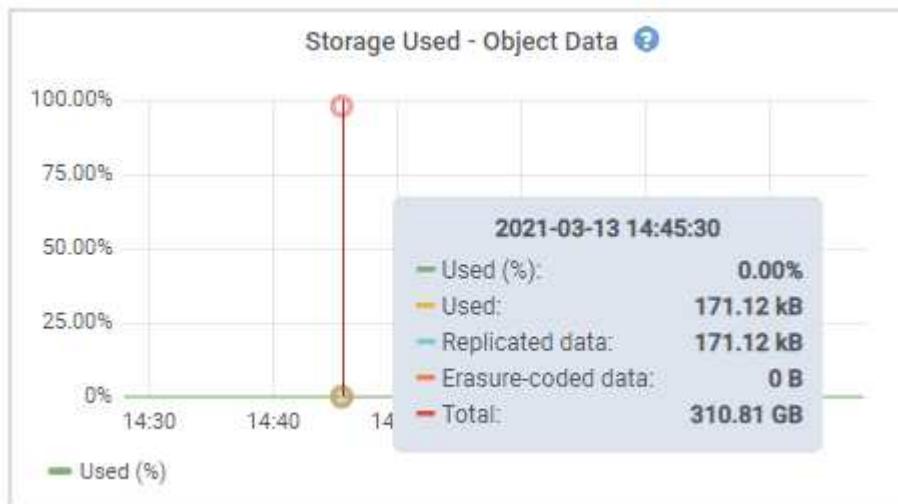
- Time triggered
- The name of the site and node
- The current values of the metrics for this alert

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

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

The following values are shown:

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



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

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

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

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



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

Related information

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

[Expand your grid](#)

[Administer StorageGRID](#)

Troubleshooting the Storage Status (SSTS) alarm

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

What you'll need

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

About this task

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



Storage Options Overview

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

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

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

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

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

Steps

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

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

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

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



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

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

Overview Alarms Reports Configuration

Main



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

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

Storage State - Desired:	Online	
Storage State - Current:	Read-only	
Storage Status:	Insufficient Free Space	

Utilization

Total Space:	164 GB	
Total Usable Space:	19.6 GB	
Total Usable Space (Percent):	11.937 %	
Total Data:	139 GB	
Total Data (Percent):	84.567 %	

Replication

Block Reads:	0	
Block Writes:	2,279,881	
Objects Retrieved:	0	
Objects Committed:	88,882	
Objects Deleted:	16	
Delete Service State:	Enabled	

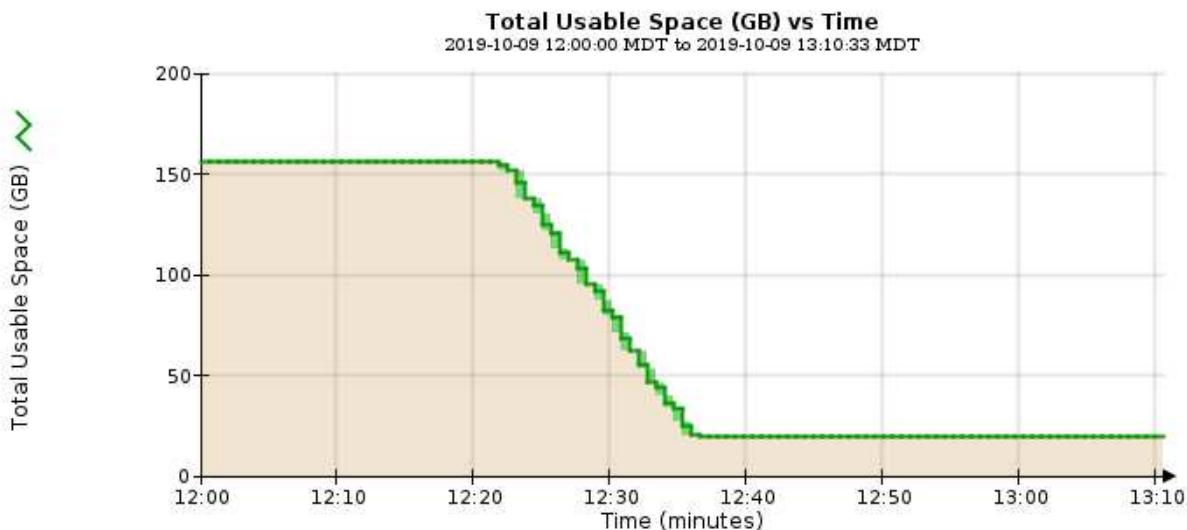
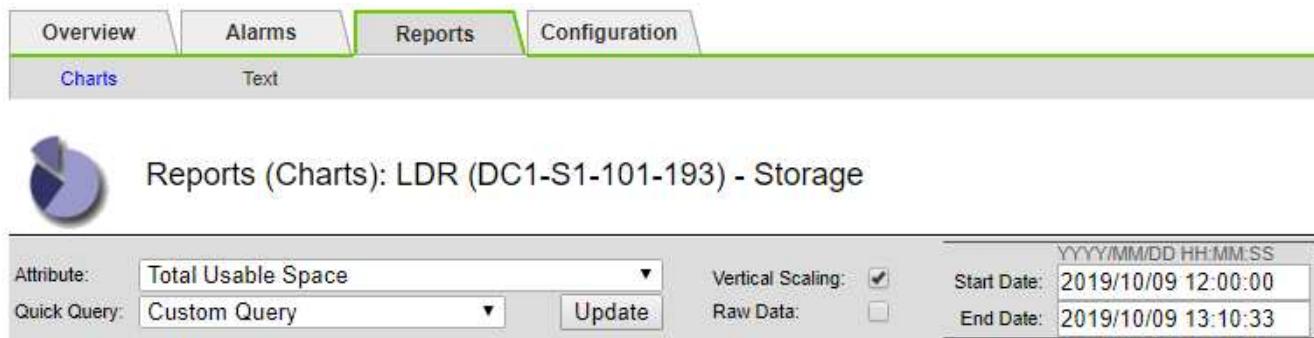
Object Store Volumes

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

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

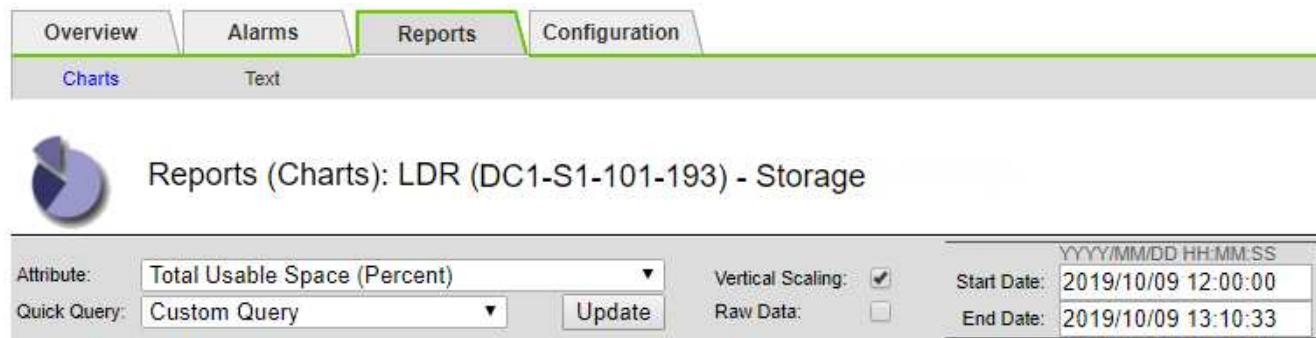
4. To understand how storage has been used over time, select the **Reports** tab, and plot Total Usable Space over the last few hours.

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



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

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



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

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

Related information

[Expand your grid](#)

[Administer StorageGRID](#)

Troubleshooting delivery of platform services messages (SMTT alarm)

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

About this task

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

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

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

Steps

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

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

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

Related information

[Administer StorageGRID](#)

[Use a tenant account](#)

[Log files reference](#)

[Resetting event counts](#)

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.