



Troubleshoot object and storage issues

StorageGRID software

NetApp

December 03, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-119/troubleshoot/confirming-object-data-locations.html> on December 03, 2025. Always check docs.netapp.com for the latest.

Table of Contents

- Troubleshoot object and storage issues 1
 - Confirm object data locations 1
 - Object store (storage volume) failures 2
 - Verify object integrity 5
 - What is background verification? 5
 - What is object existence check? 8
 - Troubleshoot S3 PUT Object size too large alert 12
- Troubleshoot lost and missing object data 14
 - Troubleshoot lost and missing object data 14
 - Investigate lost objects 15
 - Search for and restore potentially lost objects 18
 - Reset lost and missing object counts 22
- Troubleshoot the Low object data storage alert 23
- Troubleshoot Low read-only watermark override alerts 25
 - Understand the alert 26
 - Resolve the alert 26

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.

Before you begin

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

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).



The screenshot shows a web form titled "Object Metadata Lookup". Below the title is a subtitle: "Enter the identifier for any object stored in the grid to view its metadata." There are two input fields. The first is labeled "Identifier" and contains the text "source/testobject". The second is labeled "Version ID (optional)" and contains a long alphanumeric string: "MEJGMkMyQzgtNEYSOC0xMUU3LTkzMEYtRDkyNTAwQkY5N0Mx". Below these fields is a blue button labeled "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": {
      "PAWS": "2",
```

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 [storage volume alerts](#). 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.

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 policies. 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 about 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 can't 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 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 can't 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 related to background verification

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

If background verification can't replace a corrupted object because it can't 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.

Before you begin

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

About this task

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 

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 policies.
- If the object identifier can't 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 aren't required, you can delete them.



If the **Objects lost** alert 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 policies. 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 can't 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 for the job.

Before you begin

- 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. 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** checkbox in the column header.

You can search by node name or site.

You can't select nodes that aren't 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** checkbox in the column header.

6. Select **Continue**.
7. Select the consistency for the job.

The consistency 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, 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 can't 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 OLST.

12. If you selected the strong-site or strong-global consistency 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:

- i. 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 checkbox 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

Search by Job ID/ node name/ consistency control/ start time

Displaying 4 results

☐	Job ID	Status	Nodes (volumes)	Missing object copies detected	Consistency control	Start time	End time
☒	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)
☐	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.
- 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 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 S3 PUT Object size too large alert

The S3 PUT Object size too large alert is triggered if a tenant attempts a non-multipart PutObject operation that exceeds the S3 size limit of 5 GiB.

Before you begin

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

Determine which tenants use objects that are larger than 5 GiB, so you can notify them.

Steps

1. Go to **CONFIGURATION > Monitoring > Audit and syslog server**.
2. If Client Writes are Normal, access the audit log:
 - a. Enter `ssh admin@primary_Admin_Node_IP`
 - b. Enter the password listed in the `Passwords.txt` file.
 - c. Enter the following command to switch to root: `su -`
 - d. Enter the password listed in the `Passwords.txt` file.

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

- e. Change to the directory where the audit logs are located.

The audit log directory and applicable nodes depend on your audit destination settings.

Option	Destination
Local nodes (default)	<code>/var/local/log/localaudit.log</code>
Admin Nodes/local nodes	• Admin Nodes (Primary and Non-Primary): <code>/var/local/audit/export/audit.log</code> • All nodes: The <code>/var/local/log/localaudit.log</code> file is typically empty or missing in this mode.
External syslog server	<code>/var/local/log/localaudit.log</code>

Depending on your audit destination settings, enter: `cd /var/local/log` or `/var/local/audit/export/`

To learn more, refer to [Select audit information destinations](#).

- f. Identify which tenants are using objects larger than 5 GiB.
 - i. Enter `zgrep SPUT * | egrep "CSIZ\(UI64\):([5-9]|[1-9][0-9]+)[0-9]{9}"`
 - ii. For each audit message in the results, look at `S3AI` field to determine the tenant account ID. Use the other fields in the message to determine which IP address was used by the client, the bucket, and the object:

Code	Description
SAIP	Source IP
S3AI	Tenant ID
S3BK	Bucket
S3KY	Object

Code	Description
CSIZ	Size (bytes)

Example audit log results

```
audit.log:2023-01-05T18:47:05.525999
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1672943621106262][TIME(UI64):80
4317333][SAIP(IPAD):"10.96.99.127"][S3AI(CSTR):"933908492661540043
43"][SACC(CSTR):"bhavna"][S3AK(CSTR):"06OX85M40Q90Y280B7YT"][SUSR(
CSTR):"urn:sgws:identity::93390849266154004343:root"][SBAI(CSTR):"
93390849266154004343"][SBAC(CSTR):"bhavna"][S3BK(CSTR):"test"][S3K
Y(CSTR):"large-
object"][CBID(UI64):0x077EA25F3B36C69A][UUID(CSTR):"A80219A2-CD1E-
466F-9094-
B9C0FDE2FFA3"]][CSIZ(UI64):6040000000][MTME(UI64):1672943621338958]
[AVER(UI32):10][ATIM(UI64):1672944425525999][ATYP(FC32):SPUT][ANID
(UI32):12220829][AMID(FC32):S3RQ][ATID(UI64):4333283179807659119]]
```

3. If Client Writes aren't Normal, use the tenant ID from the alert to identify the tenant:

- Go to **SUPPORT > Tools > Logs**. Collect application logs for the Storage Node in the alert. Specify 15 minutes before and after the alert.
- Extract the file and go to `broadcast.log`:

```
/GID<grid_id>_<time_stamp>/<site_node>/<time_stamp>/grid/broadcast.log
```

- Search the log for `method=PUT` and identify the client in the `clientIP` field.

Example broadcast.log

```
Jan  5 18:33:41 BHAVNAJ-DC1-S1-2-65 ADE: |12220829 1870864574 S3RQ
%CEA 2023-01-05T18:33:41.208790| NOTICE 1404 af23cb66b7e3efa5 S3RQ:
EVENT_PROCESS_CREATE - connection=1672943621106262 method=PUT
name=</test/4MiB-0> auth=<V4> clientIP=<10.96.99.127>
```

- Inform tenants that the maximum PutObject size is 5 GiB and to use multipart uploads for objects greater than 5 GiB.
- Ignore the alert for one week if the application has been changed.

Troubleshoot lost and missing object data

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 can't be retrieved, the object is considered lost, and the alert is triggered.
- For erasure-coded copies, if a copy can't 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 Nodes. See [Investigate lost objects](#).

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. See [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.

Before you begin

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

The audit log directory and applicable nodes depend on your audit destination settings.

Option	Destination
Local nodes (default)	<code>/var/local/log/localaudit.log</code>
Admin Nodes/local nodes	- Admin Nodes (Primary and Non-Primary): <code>/var/local/audit/export/audit.log</code> - All nodes: The <code>/var/local/log/localaudit.log</code> file is typically empty or missing in this mode.

Option	Destination
External syslog server	/var/local/log/localaudit.log

Depending on your audit destination settings, enter: `cd /var/local/log` or `/var/local/audit/export/`

To learn more, refer to [Select audit information destinations](#).

- 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. Look up the metadata for the lost object by using the UUID:
 - a. Select **ILM > Object metadata lookup**.
 - b. Enter the UUID, and select **Look Up**.
 - c. Review the locations in the metadata, and take the appropriate action:

Metadata	Conclusion
Object <object_identifier> not found	If the object is not found, the message "ERROR":"" is returned. 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.
Locations > 0	If there are locations listed in the output, the Objects lost alert might be a false positive. 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. (The procedure for searching for potentially lost objects explains how to use the Node ID to find the correct Storage Node.) If the objects exist, you can reset the count of Objects lost to clear the alert.

Metadata	Conclusion
Locations = 0	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. Technical support might ask you to determine if there is a storage recovery procedure in progress. See the information about restoring object data using Grid Manager and restoring object data to a storage volume.

Search for and restore potentially lost objects

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

Before you begin

- You have the UUID of any lost object, as identified in [Investigate lost objects](#).
- You 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.

The audit log directory and applicable nodes depend on your audit destination settings.

Option	Destination
Local nodes (default)	<code>/var/local/log/localaudit.log</code>

Option	Destination
Admin Nodes/local nodes	• Admin Nodes (Primary and Non-Primary): /var/local/audit/export/audit.log • All nodes: The /var/local/log/localaudit.log file is typically empty or missing in this mode.
External syslog server	/var/local/log/localaudit.log

Depending on your audit destination settings, enter: `cd /var/local/log` or `/var/local/audit/export/`

To learn more, refer to [Select audit information destinations](#).

- c. Use `grep` to extract the [audit messages associated with the potentially lost object](#) and send them to an output file. Enter: `grep uuid-value audit_file_name > output_file_name`

For example:

```
Admin: # grep 926026C4-00A4-449B-AC72-BCCA72DD1311 audit.log >
/var/local/tmp/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 /var/local/tmp/messages_about_lost_objects.txt
```

An LLST audit message looks like this example 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):15815351
34379225]
[ATYP(FC32):LLST][ANID(UI32):12448208][AMID(FC32):CLSM][ATID(UI64):70
86871083190743409]]
```

- 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 associated with this LDR node ID. In the Grid Manager, select **SUPPORT > Tools > Grid topology**. Then select **Data Center > Storage Node > LDR**.

The Node ID for the LDR service is in the Node Information table. Review the information for each Storage Node until you find the one that hosts this LDR.

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

- 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. 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'
```



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 can't be restored using this procedure. Contact technical support.
- If the object path is found, continue with the next step. 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:

- a. 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'`
- b. Telnet to localhost 1402 to access the LDR console. Enter: `telnet 0 1402`
- c. Enter: `cd /proc/STOR`
- d. 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 policies, which make additional copies as specified in each policy.



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 can't 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 the next step.

4. If the object was successfully restored to StorageGRID, verify that the new locations were created:
 - a. Sign in to the Grid Manager using a [supported web browser](#).
 - b. Select **ILM > Object metadata lookup**.
 - c. Enter the UUID, and select **Look Up**.
 - d. Review the metadata, and verify the new locations.
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. Refer to [substep 1. b](#).
 - c. Use `grep` to extract the audit messages associated with the object to an output file. Enter: `grep uuid-value audit_file_name > output_file_name`

For example:

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

An ORLM audit message looks like this example 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.

6. [Reset the lost and missing object counts](#) in the Grid Manager.

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.

Before you begin

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

About this task


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.
- Select **Site > Storage Node > LDR > Erasure Coding > Configuration**.
 - Select **Reset Reads Failure Count** and **Reset Corrupt Copies Detected Count**.
 - Click **Apply Changes**.
 - Select **Site > Storage Node > LDR > Verification > Configuration**.
 - Select **Reset Missing Objects Count** and **Reset Corrupt Objects Count**.
 - If you are confident that quarantined objects aren't 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.

- 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.

Before you begin

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

About this task

The **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. Position 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, [add storage capacity](#) to your grid.

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



For more information, see [Manage full Storage Nodes](#).

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 or higher, 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.

Before you begin

- You have completed the upgrade to StorageGRID 11.6 or higher.
- 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

Steps

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.

Search... 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. Follow the appropriate step:

- 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](#).
- If ILM rules use Strict ingest behavior or if specific storage pools 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.

Steps

- Select **SUPPORT > Tools > Metrics**.
- In the Prometheus section, select the link to access the Prometheus user interface.
- 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

Steps

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 the 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

Steps

1. Go to **SUPPORT > Other > Storage watermarks**.
2. Select the **Use optimized values** checkbox.
3. Select **Save**.

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.

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.