



Manage Storage Nodes

StorageGRID 11.7

NetApp
April 12, 2024

Table of Contents

- Manage Storage Nodes 1
 - Manage Storage Nodes: Overview 1
 - What is a Storage Node? 1
 - Use Storage options 4
 - Manage object metadata storage 9
 - Compress stored objects 16
 - Storage Node configuration settings 17
 - Manage full Storage Nodes 20

Manage Storage Nodes

Manage Storage Nodes: Overview

Storage Nodes provide disk storage capacity and services. Managing Storage Nodes entails the following:

- Managing storage options
- Understanding what storage volume watermarks are and how you can use watermark overrides to control when Storage Nodes become read-only
- Monitoring and managing the space used for object metadata
- Configuring global settings for stored objects
- Applying Storage Node configuration settings
- Managing full Storage Nodes

What is a Storage Node?

Storage Nodes manage and store object data and metadata. Each StorageGRID system must have at least three Storage Nodes. If you have multiple sites, each site within your StorageGRID system must also have three Storage Nodes.

A Storage Node includes the services and processes required to store, move, verify, and retrieve object data and metadata on disk. You can view detailed information about the Storage Nodes on the **NODES** page.

What is the ADC service?

The Administrative Domain Controller (ADC) service authenticates grid nodes and their connections with each other. The ADC service is hosted on each of the first three Storage Nodes at a site.

The ADC service maintains topology information including the location and availability of services. When a grid node requires information from another grid node or an action to be performed by another grid node, it contacts an ADC service to find the best grid node to process its request. In addition, the ADC service retains a copy of the StorageGRID deployment's configuration bundles, allowing any grid node to retrieve current configuration information. You can view ADC information for a Storage Node on the Grid Topology page (**SUPPORT > Grid topology**).

To facilitate distributed and islanded operations, each ADC service synchronizes certificates, configuration bundles, and information about services and topology with the other ADC services in the StorageGRID system.

In general, all grid nodes maintain a connection to at least one ADC service. This ensures that grid nodes are always accessing the latest information. When grid nodes connect, they cache other grid nodes' certificates, enabling systems to continue functioning with known grid nodes even when an ADC service is unavailable. New grid nodes can only establish connections by using an ADC service.

The connection of each grid node lets the ADC service gather topology information. This grid node information includes the CPU load, available disk space (if it has storage), supported services, and the grid node's site ID. Other services ask the ADC service for topology information through topology queries. The ADC service responds to each query with the latest information received from the StorageGRID system.


What is the DDS service?

Hosted by a Storage Node, the Distributed Data Store (DDS) service interfaces with the Cassandra database to perform background tasks on the object metadata stored in the StorageGRID system.

Object counts

The DDS service tracks the total number of objects ingested into the StorageGRID system as well as the total number of objects ingested through each of the system’s supported interfaces (S3 or Swift).

You can see the Total Objects count on the Nodes page > Objects tab for any Storage Node.



Queries

You can identify the average time that it takes to run a query against the metadata store through the specific DDS service, the total number of successful queries, and the total number of queries that failed because of a timeout issue.

You might want to review query information to monitor the health of the metadata store, Cassandra, which impacts the system’s ingest and retrieval performance. For example, if the latency for an average query is slow and the number of failed queries due to timeouts is high, the metadata store might be encountering a higher load or performing another operation.

You can also view the total number of queries that failed because of consistency failures. Consistency level failures result from an insufficient number of available metadata stores at the time a query is performed through the specific DDS service.

You can use the Diagnostics page to obtain additional information about the current state of your grid. See [Run diagnostics](#).

Consistency guarantees and controls

StorageGRID guarantees read-after-write consistency for newly created objects. Any GET operation following a successfully completed PUT operation will be able to read the newly written data. Overwrites of existing objects, metadata updates, and deletes remain eventually consistent.

What is the LDR service?

Hosted by each Storage Node, the Local Distribution Router (LDR) service handles content transport for the StorageGRID system. Content transport encompasses many tasks including data storage, routing, and request handling. The LDR service does most of the StorageGRID system's hard work by handling data transfer loads and data traffic functions.

The LDR service handles the following tasks:

- Queries
- Information lifecycle management (ILM) activity
- Object deletion
- Object data storage
- Object data transfers from another LDR service (Storage Node)
- Data storage management
- Protocol interfaces (S3 and Swift)

The LDR service also manages the mapping of S3 and Swift objects to the unique "content handles" (UUIDs) that the StorageGRID system assigns to each ingested object.

Queries

LDR queries include queries for object location during retrieve and archive operations. You can identify the average time that it takes to run a query, the total number of successful queries, and the total number of queries that failed because of a timeout issue.

You can review query information to monitor the health of the metadata store, which impacts the system's ingest and retrieval performance. For example, if the latency for an average query is slow and the number of failed queries due to timeouts is high, the metadata store might be encountering a higher load or performing another operation.

You can also view the total number of queries that failed because of consistency failures. Consistency level failures result from an insufficient number of available metadata stores at the time a query is performed through the specific LDR service.

You can use the Diagnostics page to obtain additional information about the current state of your grid. See [Run diagnostics](#).










ILM activity

Information lifecycle management (ILM) metrics allow you to monitor the rate at which objects are evaluated for ILM implementation. You can view these metrics on the dashboard or at **NODES > Storage Node > ILM**.

Object stores

The underlying data storage of an LDR service is divided into a fixed number of object stores (also known as storage volumes). Each object store is a separate mount point.

You can see the object stores for a Storage Node on the Nodes page > Storage tab.

Object stores						
ID	Size	Available	Replicated data	EC data	Object data (%)	Health
0000	107.32 GB	96.44 GB 	124.60 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

The object stores in a Storage Node are identified by a hexadecimal number from 0000 to 002F, which is known as the volume ID. Space is reserved in the first object store (volume 0) for object metadata in a Cassandra database; any remaining space on that volume is used for object data. All other object stores are used exclusively for object data, which includes replicated copies and erasure-coded fragments.

To ensure even space usage for replicated copies, object data for a given object is stored to one object store based on available storage space. When one or more object stores fill to capacity, the remaining object stores continue to store objects until there is no more room on the Storage Node.

Metadata protection

Object metadata is information related to or a description of an object; for example, object modification time, or storage location. StorageGRID stores object metadata in a Cassandra database, which interfaces with the LDR service.

To ensure redundancy and thus protection against loss, three copies of object metadata are maintained at each site. This replication is non-configurable and performed automatically.

[Manage object metadata storage](#)


Use Storage options

What is object segmentation?

Object segmentation is the process of splitting up an object into a collection of smaller fixed-size objects to optimize storage and resources usage for large objects. S3 multi-part upload also creates segmented objects, with an object representing each part.

When an object is ingested into the StorageGRID system, the LDR service splits the object into segments, and

creates a segment container that lists the header information of all segments as content.



On retrieval of a segment container, the LDR service assembles the original object from its segments and returns the object to the client.

The container and segments aren't necessarily stored on the same Storage Node. Container and segments can be stored on any Storage Node within the storage pool specified in the ILM rule.

Each segment is treated by the StorageGRID system independently and contributes to the count of attributes such as Managed Objects and Stored Objects. For example, if an object stored to the StorageGRID system is split into two segments, the value of Managed Objects increases by three after the ingest is complete, as follows:


segment container + segment 1 + segment 2 = three stored objects

You can improve performance when handling large objects by ensuring that:

- Each Gateway and Storage Node has sufficient network bandwidth for the throughput required. For example, configure separate Grid and Client Networks on 10 Gbps Ethernet interfaces.
- Enough Gateway and Storage Nodes are deployed for the throughput required.
- Each Storage Node has sufficient disk I/O performance for the throughput required.

What are storage volume watermarks?

StorageGRID uses three storage volume watermarks to ensure that Storage Nodes are safely transitioned to a read-only state before they run critically low on space and to allow Storage Nodes that have been transitioned to a read-only state to become read-write again.



Storage volume watermarks only apply to the space used for replicated and erasure-coded object data. To learn about the space reserved for object metadata on volume 0, go to [Manage object metadata storage](#).

What is the Soft Read-Only Watermark?

The **Storage Volume Soft Read-Only Watermark** is the first watermark to indicate that a Storage Node's usable space for object data is becoming full.

If each volume in a Storage Node has less free space than that volume's Soft Read-Only Watermark, the Storage Node transitions into *read-only mode*. Read-only mode means that the Storage Node advertises read-only services to the rest of the StorageGRID system, but fulfills all pending write requests.

For example, suppose each volume in a Storage Node has a Soft Read-Only Watermark of 10 GB. As soon as each volume has less than 10 GB of free space, the Storage Node transitions to soft read-only mode.

What is the Hard Read-Only Watermark?

The **Storage Volume Hard Read-Only Watermark** is the next watermark to indicate that a node's usable space for object data is becoming full.

If the free space on a volume is less than that volume's Hard Read-Only Watermark, writes to the volume will fail. Writes to other volumes can continue, however, until the free space on those volumes is less than their Hard Read-Only Watermarks.

For example, suppose each volume in a Storage Node has a Hard Read-Only Watermark of 5 GB. As soon as each volume has less than 5 GB of free space, the Storage Node no longer accepts any write requests.

The Hard Read-Only Watermark is always less than the Soft Read-Only Watermark.

What is the Read-Write Watermark?

The **Storage Volume Read-Write Watermark** only applies to Storage Nodes that have transitioned to read-only mode. It determines when the node can become read-write again. When the free space on any one storage volume in a Storage Node is greater than that volume's Read-Write Watermark, the node automatically transitions back to the read-write state.

For example, suppose the Storage Node has transitioned to read-only mode. Also suppose that each volume has a Read-Write Watermark of 30 GB. As soon as the free space for any volume increases to 30 GB, the node becomes read-write again.

The Read-Write Watermark is always larger than both the Soft Read-Only Watermark and the Hard Read-Only Watermark.

View storage volume watermarks

You can view the current watermark settings and the system-optimized values. If optimized watermarks aren't being used, you can determine if you can or should adjust the settings.

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.

View current watermark settings

You can view the current storage watermark settings in the Grid Manager.


Steps

1. Select **CONFIGURATION > System > Storage options**.
2. In the Storage Watermarks section, look at the settings for the three storage volume watermark overrides.

Storage Options

Overview

Configuration



Storage Options Overview

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

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

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

Ports

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

- If the watermark overrides are **0**, all three watermarks are optimized for every storage volume on every Storage Node, based on the size of the Storage Node and the relative capacity of the volume.

This is the default and recommended setting. You should not update these values. As required, you can optionally [View optimized storage watermarks](#).

- If the watermark overrides are non-0 values, custom (non-optimized) watermarks are being used. Using custom watermark settings is not recommended. Use the instructions for [troubleshooting Low read-only watermark override alerts](#) to determine if you can or should adjust the settings.

View optimized storage watermarks

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

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

```
storagegrid_storage_volume_minimum_optimized_soft_readonly_watermark
```

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

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

```
storagegrid_storage_volume_maximum_optimized_soft_readonly_watermark
```

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

Manage object metadata storage

The object metadata capacity of a StorageGRID system controls the maximum number of objects that can be stored on that system. To ensure that your StorageGRID system has adequate space to store new objects, you must understand where and how StorageGRID stores object metadata.

What is object metadata?

Object metadata is any information that describes an object. StorageGRID uses object metadata to track the locations of all objects across the grid and to manage each object's lifecycle over time.


For an object in StorageGRID, object metadata includes the following types of information:

- System metadata, including a unique ID for each object (UUID), the object name, the name of the S3 bucket or Swift 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, segment identifiers and data sizes.

How is object metadata stored?

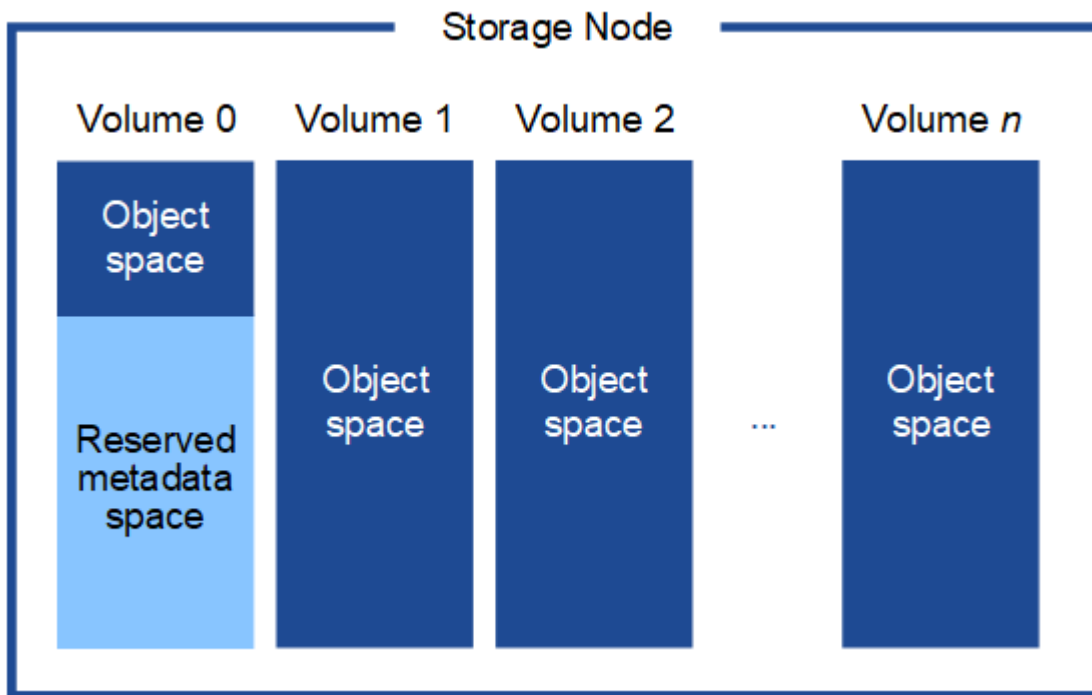
StorageGRID maintains object metadata in a Cassandra database, which is stored independently of object data. To provide redundancy and to protect object metadata from loss, StorageGRID stores three copies of the metadata for all objects in the system at each site.

This figure represents the Storage Nodes at two sites. Each site has the same amount of object metadata, and each site's metadata is subdivided among all Storage Nodes at that site.



Where is object metadata stored?

This figure represents the storage volumes for a single Storage Node.



As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. It uses the reserved space to store object metadata and to perform essential database operations. Any remaining space on storage volume 0 and all other storage volumes in the Storage Node are used exclusively for object data (replicated copies and erasure-coded fragments).

The amount of space that is reserved for object metadata on a particular Storage Node depends on several factors, which are described below.

Metadata Reserved Space setting

The *Metadata Reserved Space* is a system-wide setting that represents the amount of space that will be reserved for metadata on volume 0 of every Storage Node. As shown in the table, the default value of this setting is based on:

- The software version you were using when you initially installed StorageGRID.
- The amount of RAM on each Storage Node.


Version used for initial StorageGRID installation	Amount of RAM on Storage Nodes	Default Metadata Reserved Space setting
11.5 to 11.7	128 GB or more on each Storage Node in the grid	8 TB (8,000 GB)
	Less than 128 GB on any Storage Node in the grid	3 TB (3,000 GB)
11.1 to 11.4	128 GB or more on each Storage Node at any one site	4 TB (4,000 GB)
	Less than 128 GB on any Storage Node at each site	3 TB (3,000 GB)
11.0 or earlier	Any amount	2 TB (2,000 GB)

View Metadata Reserved Space setting

Follow these steps to view the Metadata Reserved Space setting for your StorageGRID system.

Steps

1. Select **CONFIGURATION > System > Storage options**.
2. In the Storage Watermarks table, locate **Metadata Reserved Space**.



Storage Options Overview
 Updated: 2021-12-10 13:53:01 MST

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark Override	0 B
Storage Volume Soft Read-Only Watermark Override	0 B
Storage Volume Hard Read-Only Watermark Override	0 B
Metadata Reserved Space	8,000 GB

In the screenshot, the **Metadata Reserved Space** value is 8,000 GB (8 TB). This is the default setting for a

new StorageGRID 11.6 or higher installation in which each Storage Node has 128 GB or more of RAM.

Actual reserved space for metadata

In contrast to the system-wide Metadata Reserved Space setting, the *actual reserved space* for object metadata is determined for each Storage Node. For any given Storage Node, the actual reserved space for metadata depends on the size of volume 0 for the node and the system-wide **Metadata Reserved Space** setting.


Size of volume 0 for the node	Actual reserved space for metadata
Less than 500 GB (non production use)	10% of volume 0
500 GB or more	The smaller of these values: <ul style="list-style-type: none">• Volume 0• Metadata Reserved Space setting

View actual reserved space for metadata

Follow these steps to view the actual reserved space for metadata on a particular Storage Node.

Steps

1. From the Grid Manager, select **NODES > Storage Node**.
2. Select the **Storage** tab.
3. Position your cursor over the Storage Used - Object Metadata chart and locate the **Actual reserved** value.



In the screenshot, the **Actual reserved** value is 8 TB. This screenshot is for a large Storage Node in a new StorageGRID 11.6 installation. Because the system-wide Metadata Reserved Space setting is smaller than volume 0 for this Storage Node, the actual reserved space for this node equals the Metadata Reserved Space setting.


Example for actual reserved metadata space

Suppose you install a new StorageGRID system using version 11.7. For this example, assume that each Storage Node has more than 128 GB of RAM and that volume 0 of Storage Node 1 (SN1) is 6 TB. Based on these values:

- The system-wide **Metadata Reserved Space** is set to 8 TB. (This is the default value for a new StorageGRID 11.6 or higher installation if each Storage Node has more than 128 GB RAM.)
- The actual reserved space for metadata for SN1 is 6 TB. (The entire volume is reserved because volume 0 is smaller than the **Metadata Reserved Space** setting.)

Allowed metadata space

Each Storage Node’s actual reserved space for metadata is subdivided into the space available for object metadata (the *allowed metadata space*) and the space required for essential database operations (such as compaction and repair) and future hardware and software upgrades. The allowed metadata space governs overall object capacity.



The following table shows how StorageGRID calculates the **allowed metadata space** for different Storage Nodes, based on the amount of memory for the node and the actual reserved space for metadata.

Amount of memory on Storage Node

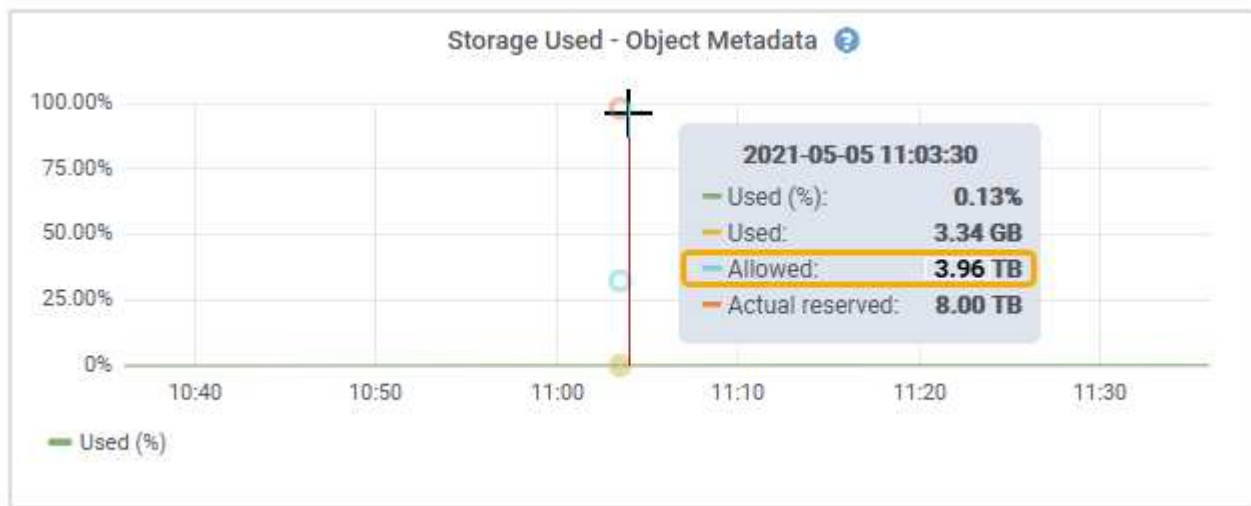
	< 128 GB	>= 128 GB
Actual reserved space for metadata <= 4 TB	60% of actual reserved space for metadata, up to a maximum of 1.32 TB	60% of actual reserved space for metadata, up to a maximum of 1.98 TB
> 4 TB	(Actual reserved space for metadata – 1 TB) × 60%, up to a maximum of 1.32 TB	(Actual reserved space for metadata – 1 TB) × 60%, up to a maximum of 3.96 TB

View allowed metadata space

Follow these steps to view the allowed metadata space for a Storage Node.

Steps

1. From the Grid Manager, select **NODES**.
2. Select the Storage Node.
3. Select the **Storage** tab.
4. Position your cursor over the Storage used - object metadata chart and locate the **Allowed** value.



In the screenshot, the **Allowed** value is 3.96 TB, which is the maximum value for a Storage Node whose actual reserved space for metadata is more than 4 TB.

The **Allowed** value corresponds to this Prometheus metric:

```
storagegrid_storage_utilization_metadata_allowed_bytes
```

Example for allowed metadata space

Suppose you install a StorageGRID system using version 11.6. For this example, assume that each Storage Node has more than 128 GB of RAM and that volume 0 of Storage Node 1 (SN1) is 6 TB. Based on these values:

- The system-wide **Metadata Reserved Space** is set to 8 TB. (This is the default value for StorageGRID 11.6 or higher when each Storage Node has more than 128 GB RAM.)
- The actual reserved space for metadata for SN1 is 6 TB. (The entire volume is reserved because volume 0 is smaller than the **Metadata Reserved Space** setting.)
- The allowed space for metadata on SN1 is 3 TB, based on the calculation shown in the [table for allowed space for metadata](#): (Actual reserved space for metadata – 1 TB) × 60%, up to a maximum of 3.96 TB.

How Storage Nodes of different sizes affect object capacity


As described above, StorageGRID evenly distributes object metadata across the Storage Nodes at each site. For this reason, if a site contains Storage Nodes of different sizes, the smallest node at the site determines the site's metadata capacity.

Consider the following example:

- You have a single-site grid containing three Storage Nodes of different sizes.
- The **Metadata Reserved Space** setting is 4 TB.
- The Storage Nodes have the following values for the actual reserved metadata space and the allowed metadata space.

Storage Node	Size of volume 0	Actual reserved metadata space	Allowed metadata space
SN1	2.2 TB	2.2 TB	1.32 TB
SN2	5 TB	4 TB	1.98 TB
SN3	6 TB	4 TB	1.98 TB

Because object metadata is evenly distributed across the Storage Nodes at a site, each node in this example can only hold 1.32 TB of metadata. The additional 0.66 TB of allowed metadata space for SN2 and SN3 can't be used.



Similarly, because StorageGRID maintains all object metadata for a StorageGRID system at each site, the overall metadata capacity of a StorageGRID system is determined by the object metadata capacity of the smallest site.

And because object metadata capacity controls the maximum object count, when one node runs out of metadata capacity, the grid is effectively full.

Related information

- To learn how to monitor the object metadata capacity for each Storage Node, see the instructions for [Monitoring StorageGRID](#).
- To increase the object metadata capacity for your system, [expand your grid](#) by adding new Storage Nodes.

Compress stored objects

You can enable object compression to reduce the size of objects stored in StorageGRID, so that objects consume less storage.

Before you begin

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

About this task

By default, object compression is disabled. If you enable compression, StorageGRID attempts to compress each object when saving it, using lossless compression.



If you change this setting, it will take about one minute for the new setting to be applied. The configured value is cached for performance and scaling.

Before enabling object compression, be aware of the following:

- You should not select **Compress stored objects** unless you know that the data being stored is compressible.
- Applications that save objects to StorageGRID might compress objects before saving them. If a client application has already compressed an object before saving it to StorageGRID, selecting this option will not further reduce an object's size.
- Don't select **Compress stored objects** if you are using NetApp FabricPool with StorageGRID.
- If **Compress stored objects** is selected, S3 and Swift client applications should avoid performing GET Object operations that specify a range of bytes be returned. These "range read" operations are inefficient because StorageGRID must effectively uncompress the objects to access the requested bytes. GET Object operations that request a small range of bytes from a very large object are especially inefficient; for example, it is inefficient to read a 10 MB range from a 50 GB compressed object.

If ranges are read from compressed objects, client requests can time out.



If you need to compress objects and your client application must use range reads, increase the read timeout for the application.

Steps

1. Select **CONFIGURATION > System > Object compression**.
2. Select the **Compress stored objects** checkbox.
3. Select **Save**.

Storage Node configuration settings

Each Storage Node uses several configuration settings and counters. You might need to view the current settings or reset counters to clear alarms (legacy system).



Except when specifically instructed in documentation, you should consult with technical support before modifying any Storage Node configuration settings. As required, you can reset event counters to clear legacy alarms.

Follow these steps to access a Storage Node's configuration settings and counters.

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **site > Storage Node**.
3. Expand the Storage Node and select the service or component.
4. Select the **Configuration** tab.

The following tables summarize Storage Node configuration settings.

LDR

Attribute Name	Code	Description
HTTP State	HSTE	The current state of HTTP for S3, Swift, and other internal StorageGRID traffic: <ul style="list-style-type: none">• Offline: No operations are allowed, and any client application that attempts to open an HTTP session to the LDR service receives an error message. Active sessions are gracefully closed.• Online: Operation continues normally
Auto-Start HTTP	HTAS	<ul style="list-style-type: none">• If selected, the state of the system on restart depends on the state of the LDR > Storage component. If the LDR > Storage component is Read-only on restart, the HTTP interface is also Read-only. If the LDR > Storage component is Online, then HTTP is also Online. Otherwise, the HTTP interface remains in the Offline state.• If not selected, the HTTP interface remains Offline until explicitly enabled.

LDR > Data Store

Attribute Name	Code	Description
Reset Lost Objects Count	RCOR	Reset the counter for the number of lost objects on this service.

LDR > Storage

Attribute Name	Code	Description
Storage State — Desired	SSDS	<p>A user-configurable setting for the desired state of the storage component. The LDR service reads this value and attempts to match the status indicated by this attribute. The value is persistent across restarts.</p> <p>For example, you can use this setting to force storage to become read-only even when there is ample available storage space. This can be useful for troubleshooting.</p> <p>The attribute can take one of the following values:</p> <ul style="list-style-type: none">• Offline: When the desired state is Offline, the LDR service takes the LDR > Storage component offline.• Read-only: When the desired state is Read-only, the LDR service moves the storage state to read-only and stops accepting new content. Note that content might continue to be saved to the Storage Node for a short time until open sessions are closed.• Online: Leave the value at Online during normal system operations. The Storage State — Current of the storage component will be dynamically set by the service based on the condition of the LDR service, such as the amount of available object storage space. If space is low, the component becomes Read-only.
Health Check Timeout	SHCT	<p>The time limit in seconds within which a health check test must complete in order for a storage volume to be considered healthy. Only change this value when directed to do so by Support.</p>

LDR > Verification

Attribute Name	Code	Description
Reset Missing Objects Count	VCMI	<p>Resets the count of Missing Objects Detected (OMIS). Use only after object existence check completes. Missing replicated object data is restored automatically by the StorageGRID system.</p>
Verification Rate	VPRI	<p>Set the rate at which background verification takes place. See information about configuring the background verification rate.</p>

Attribute Name	Code	Description
Reset Corrupt Objects Count	VCCR	Reset the counter for corrupt replicated object data found during background verification. This option can be used to clear the Corrupt Objects Detected (OCOR) alarm condition.
Delete Quarantined Objects	OQRT	<p>Delete corrupt objects from the quarantine directory, reset the count of quarantined objects to zero, and clear the Quarantined Objects Detected (OQRT) alarm. This option is used after corrupt objects have been automatically restored by the StorageGRID system.</p> <p>If a Lost Objects alarm is triggered, technical support might want to access the quarantined objects. In some cases, quarantined objects might be useful for data recovery or for debugging the underlying issues that caused the corrupt object copies.</p>

LDR > Erasure Coding

Attribute Name	Code	Description
Reset Writes Failure Count	RSWF	Reset the counter for write failures of erasure-coded object data to the Storage Node.
Reset Reads Failure Count	RSRF	Reset the counter for read failures of erasure-coded object data from the Storage Node.
Reset Deletes Failure Count	RSDF	Reset the counter for delete failures of erasure-coded object data from the Storage Node.
Reset Corrupt Copies Detected Count	RSCC	Reset the counter for the number of corrupt copies of erasure-coded object data on the Storage Node.
Reset Corrupt Fragments Detected Count	RSCD	Reset the counter for corrupt fragments of erasure-coded object data on the Storage Node.
Reset Missing Fragments Detected Count	RSMD	Reset the counter for missing fragments of erasure-coded object data on the Storage Node. Use only after object existence check completes.

LDR > Replication

Attribute Name	Code	Description
Reset Inbound Replication Failure Count	RICR	Reset the counter for inbound replication failures. This can be used to clear the RIRF (Inbound Replication — Failed) alarm.
Reset Outbound Replication Failure Count	ROCR	Reset the counter for outbound replication failures. This can be used to clear the RORF (Outbound Replications — Failed) alarm.
Disable Inbound Replication	DSIR	<p>Select to disable inbound replication as part of a maintenance or testing procedure. Leave unchecked during normal operation.</p> <p>When inbound replication is disabled, objects can be retrieved from the Storage Node for copying to other locations in the StorageGRID system, but objects can't be copied to this Storage Node from other locations: the LDR service is read-only.</p>
Disable Outbound Replication	DSOR	<p>Select to disable outbound replication (including content requests for HTTP retrievals) as part of a maintenance or testing procedure. Leave unchecked during normal operation.</p> <p>When outbound replication is disabled, objects can be copied to this Storage Node, but objects can't be retrieved from the Storage Node to be copied to other locations in the StorageGRID system. The LDR service is write-only.</p>

Manage full Storage Nodes

As Storage Nodes reach capacity, you must expand the StorageGRID system through the addition of new storage. There are three options available: adding storage volumes, adding storage expansion shelves, and adding Storage Nodes.

Add storage volumes

Each Storage Node supports a maximum number of storage volumes. The defined maximum varies by platform. If a Storage Node contains fewer than the maximum number of storage volumes, you can add volumes to increase its capacity. See the instructions for [expanding a StorageGRID system](#).

Add storage expansion shelves

Some StorageGRID appliance Storage Nodes, such as the SG6060, can support additional storage shelves. If you have StorageGRID appliances with expansion capabilities that have not already been expanded to maximum capacity, you can add storage shelves to increase capacity. See the instructions for [expanding a StorageGRID system](#).

Add Storage Nodes

You can increase storage capacity by adding Storage Nodes. Careful consideration of currently active ILM rules and capacity requirements must be taken when adding storage. See the instructions for [expanding a StorageGRID system](#).

Copyright information

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

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

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

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

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

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

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

Trademark information

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