



StorageGRID 11.9 software documentation

StorageGRID software

NetApp
October 01, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-119/index.html> on October 01, 2025. Always check docs.netapp.com for the latest.

Table of Contents

StorageGRID 11.9 software documentation	1
StorageGRID appliances	2
Release notes	3
Get started with a StorageGRID system	4
Learn about StorageGRID	4
What is StorageGRID?	4
Hybrid clouds with StorageGRID	6
StorageGRID architecture and network topology	7
Grid nodes and services	9
How StorageGRID manages data	21
Explore StorageGRID	32
Networking guidelines	40
Networking guidelines	40
StorageGRID network types	41
Network topology examples	45
Networking requirements	51
Network-specific requirements	53
Deployment-specific networking considerations	55
Network installation and provisioning	58
Post-installation guidelines	58
Network port reference	59
Quick start for StorageGRID	67
Install, upgrade, and hotfix StorageGRID	70
StorageGRID appliances	70
Install StorageGRID on Red Hat Enterprise Linux	70
Quick start for installing StorageGRID on Red Hat Enterprise Linux	70
Plan and prepare for installation on Red Hat	71
Automate StorageGRID installation on Red Hat Enterprise Linux	96
Deploy virtual grid nodes (Red Hat)	99
Configure the grid and complete installation (Red Hat)	121
Installation REST API	133
Where to go next	134
Troubleshoot installation issues	135
Example /etc/sysconfig/network-scripts	135
Install StorageGRID on Ubuntu or Debian	138
Quick start for installing StorageGRID on Ubuntu or Debian	138
Plan and prepare for installation on Ubuntu or Debian	139
Automate the installation (Ubuntu or Debian)	165
Deploy virtual grid nodes (Ubuntu or Debian)	167
Configure grid and complete installation (Ubuntu or Debian)	189
Installation REST API	202
Where to go next	202
Troubleshoot installation issues	203

Example /etc/network/interfaces	204
Install StorageGRID on VMware	205
Quick start for installing StorageGRID on VMware	205
Plan and prepare for installation on VMware	206
Automate the installation (VMware)	215
Deploy virtual machine grid nodes (VMware)	228
Configure the grid and complete installation (VMware)	237
Installation REST API	250
Where to go next	250
Troubleshoot installation issues	251
Upgrade StorageGRID software	253
Upgrade StorageGRID software	253
What's new in StorageGRID 11.9	253
Removed or deprecated features and capabilities	255
Changes to the Grid Management API	257
Changes to the Tenant Management API	258
Plan and prepare for upgrade	258
Upgrade software	265
Troubleshoot upgrade issues	281
Apply StorageGRID hotfix	283
StorageGRID hotfix procedure	283
How your system is affected when you apply a hotfix	284
Obtain required materials for hotfix	285
Download hotfix file	286
Check system's condition before applying hotfix	287
Apply hotfix	287
Configure and manage a StorageGRID system	291
Administer StorageGRID	291
Administer StorageGRID	291
Get started with Grid Manager	291
Control access to StorageGRID	321
Use grid federation	367
Manage security	403
Manage tenants	468
Configure client connections	487
Manage networks and connections	526
Use AutoSupport	544
Manage Storage Nodes	559
Manage Admin Nodes	573
Manage objects with ILM	575
Manage objects with ILM	575
ILM and object lifecycle	576
Create and assign storage grades	597
Use storage pools	599
Use Cloud Storage Pools	607

Manage erasure-coding profiles	626
Configure regions (optional and S3 only)	630
Create ILM rule	631
Manage ILM policies	647
Work with ILM policies and ILM rules	662
Use S3 Object Lock	666
Example ILM rules and policies	674
System hardening	694
General considerations for system hardening	694
Hardening guidelines for software upgrades	695
Hardening guidelines for StorageGRID networks	696
Hardening guidelines for StorageGRID nodes	696
Hardening guidelines for TLS and SSH	699
Other hardening guidelines	700
Configure StorageGRID for FabricPool	702
Configure StorageGRID for FabricPool	702
Information needed to attach StorageGRID as a cloud tier	704
Use FabricPool setup wizard	705
Configure StorageGRID manually	718
Configure ONTAP System Manager	727
Configure the DNS server	729
StorageGRID best practices for FabricPool	730
Remove FabricPool data from StorageGRID	733
Use StorageGRID tenants and clients	735
Use a tenant account	735
Use a tenant account	735
How to sign in and sign out	736
Understand Tenant Manager dashboard	741
Tenant Management API	744
Use grid federation connections	749
Manage groups and users	762
Manage S3 access keys	780
Manage S3 buckets	785
Manage S3 platform services	807
Use S3 REST API	837
S3 REST API supported versions and updates	837
Quick reference: Supported S3 API requests	840
Test S3 REST API configuration	859
How StorageGRID implements S3 REST API	860
Support for Amazon S3 REST API	875
StorageGRID custom operations	922
Bucket and group access policies	941
S3 operations tracked in the audit logs	966
Use Swift REST API (end of life)	967
Use Swift REST API	967

Monitor and troubleshoot a StorageGRID system	968
Monitor StorageGRID system	968
Monitor a StorageGRID system	968
View and manage the dashboard	968
View the Nodes page	971
Information to monitor regularly	1003
Manage alerts	1032
Log files reference	1067
Configure audit message and log destinations	1085
Use SNMP monitoring	1098
Collect additional StorageGRID data	1110
Troubleshoot StorageGRID system	1145
Troubleshoot a StorageGRID system	1145
Troubleshoot object and storage issues	1151
Troubleshoot metadata issues	1178
Troubleshoot certificate errors	1180
Troubleshoot Admin Node and user interface issues	1182
Troubleshoot network, hardware, and platform issues	1185
Troubleshoot an external syslog server	1193
Review audit logs	1196
Audit messages and logs	1196
Audit message flow and retention	1196
Access audit log file	1199
Audit log file rotation	1200
Audit log file format	1200
Audit message format	1212
Audit messages and the object lifecycle	1216
Audit messages	1223
Expand a grid	1269
Expansion types	1269
Plan StorageGRID expansion	1270
Add storage capacity	1270
Add metadata capacity	1277
Add grid nodes to add capabilities to your system	1279
Add a new site	1279
Gather required materials	1280
Download and extract the StorageGRID installation files	1281
Verify hardware and networking	1287
Add storage volumes	1287
Add storage volumes to Storage Nodes	1287
VMware: Add storage volumes to Storage Node	1290
Linux: Add direct-attached or SAN volumes to Storage Node	1291
Add grid nodes or site	1295
Add grid nodes to existing site or add new site	1295
Update subnets for Grid Network	1295

Deploy new grid nodes	1296
Perform expansion	1302
Configure expanded system	1309
Configuration steps after expansion	1309
Verify that Storage Node is active	1311
Copy Admin Node database	1311
Copy Prometheus metrics	1313
Copy audit logs	1314
Rebalance erasure-coded data after adding Storage Nodes	1315
Troubleshoot expansion	1318
Maintain a StorageGRID system	1320
Grid maintenance	1320
Before you begin	1320
Maintenance procedures for appliances	1320
Download Recovery Package	1320
Decommission nodes or site	1321
Decommission node or site	1321
Decommission nodes	1321
Decommission site	1338
Rename grid, site, or node	1360
Use the rename procedure	1360
Add or update display names	1364
Node procedures	1370
Node maintenance procedures	1370
Server Manager procedures	1370
Reboot, shutdown, and power procedures	1379
Port remap procedures	1391
Network procedures	1394
Update subnets for Grid Network	1394
Configure IP addresses	1396
Add interfaces to existing node	1413
Configure DNS servers	1417
Modify DNS configuration for single grid node	1418
Manage NTP servers	1420
Restore network connectivity for isolated nodes	1421
Host and middleware procedures	1422
Linux: Migrate grid node to new host	1422
VMware: Configure virtual machine for automatic restart	1425
Recover or replace nodes	1426
Warnings and considerations for grid node recovery	1426
Preconditions for recovering grid nodes	1426
Order of node recovery if a server hosting more than one grid node fails	1426
IP addresses for recovered nodes	1427
Gather required materials for grid node recovery	1427
Download and extract StorageGRID installation files	1428

Select node recovery procedure	1433
Recover from Storage Node failures	1434
Recover from Storage Node failures	1434
Recover appliance Storage Node	1435
Recover from storage volume failure where system drive is intact	1455
Recover from system drive failure	1469
Restore object data using Grid Manager	1487
Monitor repair-data jobs	1491
Recover from Admin Node failures	1493
Primary or non-primary Admin Node recovery	1493
Recover from primary Admin Node failures	1493
Recover from non-primary Admin Node failures	1501
Recover from Gateway Node failures	1509
Replace Gateway Node	1509
Select Start Recovery to configure Gateway Node	1509
Recover from Archive Node failures	1511
Recover from Archive Node failures	1511
Replace Linux node	1511
Replace Linux node	1511
Deploy new Linux hosts	1511
Restore grid nodes to the host	1512
What's next: Perform additional recovery steps, if required	1516
Replace VMware node	1517
Replace failed node with services appliance	1518
Replace failed node with services appliance	1518
Install services appliance (platform change only)	1519
Prepare appliance for reinstallation (platform replacement only)	1519
Start software installation on services appliance	1520
Monitor services appliance installation	1524
How technical support recovers a site	1527
How to enable StorageGRID in your environment	1529
How to manage StorageGRID using BlueXP	1530
Legal notices	1531
Copyright	1531
Trademarks	1531
Patents	1531
Privacy policy	1531
Open source	1531

StorageGRID 11.9 software documentation

StorageGRID appliances

Go to [StorageGRID Appliance Documentation](#) to learn how to install, configure, and maintain StorageGRID storage and services appliances.

Release notes

Obtain release-specific information about fixed issues and known issues.

Log in to the NetApp Support Site to [view or download a PDF file](#) containing the StorageGRID 11.9 release notes.

Get started with a StorageGRID system


Learn about StorageGRID

What is StorageGRID?

NetApp® StorageGRID® is a software-defined object storage suite that supports a wide range of use cases across public, private, and hybrid multicloud environments. StorageGRID offers native support for the Amazon S3 API and delivers industry-leading innovations such as automated lifecycle management to store, secure, protect, and preserve unstructured data cost effectively over long periods.

StorageGRID provides secure, durable storage for unstructured data at scale. Integrated, metadata-driven lifecycle management policies optimize where your data lives throughout its life. Content is placed in the right location, at the right time, and on the right storage tier to reduce cost.

StorageGRID is composed of geographically distributed, redundant, heterogeneous nodes, which can be integrated with both existing and next-generation client applications.



Support for Archive Nodes has been removed. Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by [ILM Cloud Storage Pools](#), which offer more functionality.

StorageGRID benefits

Advantages of the StorageGRID system include the following:

- Massively scalable and easy-to-use a geographically distributed data repository for unstructured data.
- Standard object storage protocols:
 - Amazon Web Services Simple Storage Service (S3)
 - OpenStack Swift



Support for Swift client applications has been deprecated and will be removed in a future release.

- Hybrid cloud enabled. Policy-based information lifecycle management (ILM) stores objects to public clouds, including Amazon Web Services (AWS) and Microsoft Azure. StorageGRID platform services enable content replication, event notification, and metadata searching of objects stored to public clouds.
- Flexible data protection to ensure durability and availability. Data can be protected using replication and layered erasure coding. At-rest and in-flight data verification ensures integrity for long-term retention.
- Dynamic data lifecycle management to help manage storage costs. You can create ILM rules that manage data lifecycle at the object level, customizing data locality, durability, performance, cost, and retention time.
- High availability of data storage and some management functions, with integrated load balancing to optimize the data load across StorageGRID resources.
- Support for multiple storage tenant accounts to segregate the objects stored on your system by different entities.
- Numerous tools for monitoring the health of your StorageGRID system, including a comprehensive alert system, a graphical dashboard, and detailed statuses for all nodes and sites.
- Support for software or hardware-based deployment. You can deploy StorageGRID on any of the following:
 - Virtual machines running in VMware.
 - Container engines on Linux hosts.
 - StorageGRID engineered appliances.
 - Storage appliances provide object storage.
 - Services appliances provide grid administration and load balancing services.
- Compliant with the relevant storage requirements of these regulations:
 - Securities and Exchange Commission (SEC) in 17 CFR § 240.17a-4(f), which regulates exchange members, brokers or dealers.
 - Financial Industry Regulatory Authority (FINRA) Rule 4511(c), which defers to the format and media requirements of SEC Rule 17a-4(f).
 - Commodity Futures Trading Commission (CFTC) in regulation 17 CFR § 1.31(c)-(d), which regulates commodity futures trading.
- Non-disruptive upgrade and maintenance operations. Maintain access to content during upgrade, expansion, decommission, and maintenance procedures.
- Federated identity management. Integrates with Active Directory, OpenLDAP, or Oracle Directory Service for user authentication. Supports single sign-on (SSO) using the Security Assertion Markup Language 2.0 (SAML 2.0) standard to exchange authentication and authorization data between StorageGRID and Active Directory Federation Services (AD FS).

Hybrid clouds with StorageGRID

Use StorageGRID in a hybrid cloud configuration by implementing policy-driven data management to store objects in Cloud Storage Pools, leveraging StorageGRID platform services, and tiering data from ONTAP to StorageGRID with NetApp FabricPool.

Cloud Storage Pools

Cloud Storage Pools allow you to store objects outside of the StorageGRID system. For example, you might want to move infrequently accessed objects to lower-cost cloud storage, such as Amazon S3 Glacier, S3 Glacier Deep Archive, Google Cloud, or the Archive access tier in Microsoft Azure Blob storage. Or, you might want to maintain a cloud backup of StorageGRID objects, which can be used to recover data lost because of a storage volume or Storage Node failure.

Third-party partner storage is also supported, including disk and tape storage.



Using Cloud Storage Pools with FabricPool is not supported because of the added latency to retrieve an object from the Cloud Storage Pool target.

S3 platform services

S3 platform services give you the ability to use remote services as endpoints for object replication, event notifications, or search integration. Platform services operate independently of the grid's ILM rules, and are enabled for individual S3 buckets. The following services are supported:

- The CloudMirror replication service automatically mirrors specified objects to a target S3 bucket, which can be on Amazon S3 or a second StorageGRID system.
- The Event notification service sends messages about specified actions to an external endpoint that supports receiving Simple Notification Service (Amazon SNS) events.
- The search integration service sends object metadata to an external Elasticsearch service, allowing metadata to be searched, visualized, and analyzed using third party tools.

For example, you might use CloudMirror replication to mirror specific customer records into Amazon S3 and then leverage AWS services to perform analytics on your data.

ONTAP data tiering using FabricPool

You can reduce the cost of ONTAP storage by tiering data to StorageGRID using FabricPool. FabricPool enables automated tiering of data to low-cost object storage tiers, either on or off premises.

Unlike manual tiering solutions, FabricPool reduces total cost of ownership by automating the tiering of data to lower the cost of storage. It delivers the benefits of cloud economics by tiering to public and private clouds including StorageGRID.

Related information

- [What is Cloud Storage Pool?](#)
- [Manage platform services](#)
- [Configure StorageGRID for FabricPool](#)

StorageGRID architecture and network topology

A StorageGRID system consists of multiple types of grid nodes at one or more data center sites.

See the [descriptions of grid node types](#).


For additional information about StorageGRID network topology, requirements, and grid communications, see the [Networking guidelines](#).

Deployment topologies

The StorageGRID system can be deployed to a single data center site or to multiple data center sites.

Single site

In a deployment with a single site, the infrastructure and operations of the StorageGRID system are centralized.




Multiple sites

In a deployment with multiple sites, different types and numbers of StorageGRID resources can be installed at each site. For example, more storage might be required at one data center than at another.

Different sites are often located in geographically different locations across different failure domains, such as

an earthquake fault line or flood plain. Data sharing and disaster recovery are achieved by automated distribution of data to other sites.




Multiple logical sites can also exist within a single data center to allow the use of distributed replication and erasure coding for increase availability and resiliency.

Grid node redundancy

In a single-site or multi-site deployment, you can optionally include more than one Admin Node or Gateway Node for redundancy. For example, you can install more than one Admin Node at a single site or across several sites. However, each StorageGRID system can only have one primary Admin Node.

System architecture

This diagram shows how grid nodes are arranged within a StorageGRID system.



S3 clients store and retrieve objects in StorageGRID. Other clients are used to send email notifications, to access the StorageGRID management interface, and optionally to access the audit share.

S3 clients can connect to a Gateway Node or an Admin Node to use the load-balancing interface to Storage Nodes. Alternatively, S3 clients can connect directly to Storage Nodes using HTTPS.

Objects can be stored within StorageGRID on software or hardware-based Storage Nodes, or in Cloud Storage Pools, which consist of external S3 buckets or Azure Blob storage containers.

Grid nodes and services

Grid nodes and services

The basic building block of a StorageGRID system is the grid node. Nodes contain services, which are software modules that provide a set of capabilities to a grid node.

Types of grid nodes

The StorageGRID system uses four types of grid nodes:

Admin Nodes

Provide management services such as system configuration, monitoring, and logging. When you sign in to the Grid Manager, you are connecting to an Admin Node. Each grid must have one primary Admin Node and might have additional non-primary Admin Nodes for redundancy. You can connect to any Admin Node, and each Admin Node displays a similar view of the StorageGRID system. However, maintenance procedures must be performed using the primary Admin Node.

Admin Nodes can also be used to load balance S3 client traffic.

See [What is an Admin Node?](#)

Storage Nodes

Manage and store object data and metadata. Each site in your StorageGRID system must have at least three Storage Nodes.

See [What is a Storage Node?](#)

Gateway Nodes (optional)

Provide a load-balancing interface that client applications can use to connect to StorageGRID. A load balancer seamlessly directs clients to an optimal Storage Node, so that the failure of nodes or even an entire site is transparent.

See [What is a Gateway Node?](#)

Hardware and software nodes

StorageGRID nodes can be deployed as StorageGRID appliance nodes or as software-based nodes.

StorageGRID appliance nodes

StorageGRID hardware appliances are specially designed for use in a StorageGRID system. Some appliances can be used as Storage Nodes. Other appliances can be used as Admin Nodes or Gateway Nodes. You can combine appliance nodes with software-based nodes or deploy fully engineered, all-appliance grids that have no dependencies on external hypervisors, storage, or compute hardware.

See the following to learn about the available appliances:

- [StorageGRID Appliance Documentation](#)
- [NetApp Hardware Universe](#)

Software-based nodes

Software-based grid nodes can be deployed as VMware virtual machines or within container engines on a Linux host.

- Virtual machine (VM) in VMware vSphere: See [Install StorageGRID on VMware](#).
- Within a container engine on Red Hat Enterprise Linux: See [Install StorageGRID on Red Hat Enterprise Linux](#).
- Within a container engine on Ubuntu or Debian: See [Install StorageGRID on Ubuntu or Debian](#).

Use the [NetApp Interoperability Matrix Tool \(IMT\)](#) to determine the supported versions.

During initial installation of a new software-based Storage Node you can specify that it only be used to [store metadata](#).

StorageGRID services

The following is a complete list of StorageGRID services.

Service	Description	Location
Account Service Forwarder	Provides an interface for the Load Balancer service to query the Account Service on remote hosts and provides notifications of Load Balancer Endpoint configuration changes to the Load Balancer service.	Load Balancer service on Admin Nodes and Gateway Nodes
ADC (Administrative Domain Controller)	Maintains topology information, provides authentication services, and responds to queries from the LDR and CMN services.	At least three Storage Nodes containing the ADC service at each site
AMS (Audit Management System)	Monitors and logs all audited system events and transactions to a text log file.	Admin Nodes
Cassandra Reaper	Performs automatic repairs of object metadata.	Storage Nodes
Chunk service	Manages erasure-coded data and parity fragments.	Storage Nodes
CMN (Configuration Management Node)	Manages system-wide configurations and grid tasks. Each grid has one CMN service.	Primary Admin Node
DDS (Distributed Data Store)	Interfaces with the Cassandra database to manage object metadata.	Storage Nodes
DMV (Data Mover)	Moves data to cloud endpoints.	Storage Nodes
Dynamic IP (dynip)	Monitors the grid for dynamic IP changes and updates local configurations.	All nodes
Grafana	Used for metrics visualization in the Grid Manager.	Admin Nodes

Service	Description	Location
High Availability	Manages high availability virtual IPs on nodes configured on the High Availability Groups page. This service is also known as the keepalived service.	Admin and Gateway Nodes
Identity (idnt)	Federates user identities from LDAP and Active Directory.	Storage Nodes that use the ADC service
Lambda Arbitrator	Manages S3 Select SelectObjectContent requests.	All nodes
Load Balancer (nginx-gw)	Provides load balancing of S3 traffic from clients to Storage Nodes. The Load Balancer service can be configured through the Load Balancer Endpoints configuration page. This service is also known as the nginx-gw service.	Admin and Gateway Nodes
LDR (Local Distribution Router)	Manages the storage and transfer of content within the grid.	Storage Nodes
MISCd Information Service Control Daemon	Provides an interface for querying and managing services on other nodes and for managing environmental configurations on the node such as querying the state of services running on other nodes.	All nodes
nginx	Acts as an authentication and secure communication mechanism for various grid services (such as Prometheus and Dynamic IP) to be able to talk to services on other nodes over HTTPS APIs.	All nodes
nginx-gw	Powers the Load Balancer service.	Admin and Gateway Nodes
NMS (Network Management System)	Powers the monitoring, reporting, and configuration options that are displayed through the Grid Manager.	Admin Nodes
Persistence	Manages files on the root disk that need to persist across a reboot.	All nodes
Prometheus	Collects time series metrics from services on all nodes.	Admin Nodes
RSM (Replicated State Machine)	Ensures platform service requests are sent to their respective endpoints.	Storage Nodes that use the ADC service

Service	Description	Location
SSM (Server Status Monitor)	Monitors hardware conditions and reports to the NMS service.	An instance is present on every grid node
Trace collector	Performs trace collection to gather information for use by technical support. The trace collector service uses open source Jaeger software.	Admin Nodes

What is an Admin Node?

Admin Nodes provide management services such as system configuration, monitoring, and logging. Admin Nodes can also be used to load balance S3 client traffic. Each grid must have one primary Admin Node and might have any number of non-primary Admin Nodes for redundancy.

Differences between primary and non-primary Admin Nodes

When you sign in to the Grid Manager or the Tenant Manager, you are connecting to an Admin Node. You can connect to any Admin Node, and each Admin Node displays a similar view of the StorageGRID system. However, the primary Admin Node provides more functionality than non-primary Admin Nodes. For example, most maintenance procedures must be performed from the primary Admin Nodes.

The table summarizes the capabilities of primary and non-primary Admin Nodes.

Capabilities	Primary Admin Node	Non-primary Admin Node
Includes the AMS service	Yes	Yes
Includes the CMN service	Yes	No
Includes the NMS service	Yes	Yes
Includes the Prometheus service	Yes	Yes
Includes the SSM service	Yes	Yes
Includes the Load Balancer and High Availability services	Yes	Yes
Supports the Management Application Program Interface (mgmt-api)	Yes	Yes
Can be used for all network-related maintenance tasks, for example IP address change and updating NTP servers	Yes	No

Capabilities	Primary Admin Node	Non-primary Admin Node
Can perform EC rebalance after Storage Node expansion	Yes	No
Can be used for the volume restoration procedure	Yes	Yes
Can collect log files and system data from one or more nodes	Yes	No
Sends alert notifications, AutoSupport packages, and SNMP traps and informs	Yes. Acts as the preferred sender .	Yes. Acts as a standby sender.

Preferred sender Admin Node

If your StorageGRID deployment includes multiple Admin Nodes, the primary Admin Node is the preferred sender for alert notifications, AutoSupport packages, and SNMP traps and informs.

Under normal system operations, only the preferred sender sends notifications. However, all other Admin Nodes monitor the preferred sender. If a problem is detected, other Admin Nodes act as *standby senders*.

Multiple notifications might sent in these cases:

- If Admin Nodes become "islanded" from each other, both the preferred sender and the standby senders will attempt to send notifications, and multiple copies of notifications might be received.
- If standby sender detects problems with the preferred sender and starts sending notifications, the preferred sender might regain its ability to send notifications. If this occurs, duplicate notifications might be sent. The standby sender will stop sending notifications when it no longer detects errors on the preferred sender.



When you test AutoSupport packages, all Admin Nodes send the test. When you test alert notifications, you must sign in to every Admin Node to verify connectivity.

Primary services for Admin Nodes

The following table shows the primary services for Admin Nodes; however, this table does not list all node services.

Service	Key function
Audit Management System (AMS)	Tracks system activity and events.
Configuration Management Node (CMN)	Manages system-wide configuration.

Service	Key function
High Availability	Manages high availability virtual IP addresses for groups of Admin Nodes and Gateway Nodes. Note: This service is also found on Gateway Nodes.
Load Balancer	Provides load balancing of S3 traffic from clients to Storage Nodes. Note: This service is also found on Gateway Nodes.
Management Application Program Interface (mgmt-api)	Processes requests from the Grid Management API and the Tenant Management API.
Network Management System (NMS)	Provides functionality for the Grid Manager.
Prometheus	Collects and stores time-series metrics from the services on all nodes.
Server Status Monitor (SSM)	Monitors the operating system and underlying hardware.

What is a Storage Node?

Storage Nodes manage and store object data and metadata. Storage Nodes include the services and processes required to store, move, verify, and retrieve object data and metadata on disk.

Each site in your StorageGRID system must have at least three Storage Nodes.

Types of Storage Nodes

During installation, you can select the type of Storage Node you want to install. These types are available for software-based Storage Nodes and for appliance-based Storage Nodes that support the feature:

- Combined data and metadata Storage Node
- Metadata-only Storage Node
- Data-only Storage Node

You can select the Storage Node type in these situations:

- When initially installing a Storage Node
- When you add a Storage Node during StorageGRID system expansion



You can't change the type after the Storage Node installation is complete.

Data and metadata Storage Node (combined)

By default, all new Storage Nodes will store both object data and metadata. This type of Storage Node is called a *combined* Storage Node.

Metadata-only Storage Node

Using a Storage Node exclusively for metadata can make sense if your grid stores a very large number of small objects. Installing dedicated metadata capacity provides a better balance between the space needed for a very large number of small objects and the space needed for the metadata for those objects. Additionally, metadata-only Storage Nodes hosted on high-performance appliances can increase performance.

Metadata-only Storage Nodes have specific hardware requirements:

- When using StorageGRID appliances, metadata-only nodes can be configured only on SGF6112 appliances with twelve 1.9 TB or twelve 3.8 TB drives.
- When using software-based nodes, metadata-only node resources must match the existing Storage Nodes resources. For example:
 - If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)
 - If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0)).
- When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.

When installing metadata-only nodes, the grid must also contain a minimum number of nodes for data storage:

- For a single-site grid, configure at least two combined or data-only Storage Nodes.
- For a multi-site grid, configure at least one combined or data-only Storage Node *per site*.



Although metadata-only Storage Nodes contain the [LDR service](#) and can process S3 client requests, StorageGRID performance might not increase.

Data-only Storage Node

Using a Storage Node exclusively for data can make sense if your Storage Nodes have differing performance characteristics. For example, to potentially increase performance, you could have data-only, high-capacity spinning-disk Storage Nodes accompanied by metadata-only high-performance Storage Nodes.

When installing data-only nodes, the grid must contain the following:

- A minimum of two combined or data-only Storage Nodes *per grid*
- At least one combined or data-only Storage Node *per site*
- A minimum of three combined or metadata-only Storage Nodes *per site*

Primary services for Storage Nodes

The following table shows the primary services for Storage Nodes; however, this table does not list all node services.



Some services, such as the ADC service and the RSM service, typically exist only on three Storage Nodes at each site.

Service	Key function
Account (acct)	Manages tenant accounts.
Administrative Domain Controller (ADC)	Maintains topology and grid-wide configuration. Note: Data-only Storage Nodes don't host the ADC service. Details The Administrative Domain Controller (ADC) service authenticates grid nodes and their connections with each other. The ADC service is hosted on a minimum of 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. 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.
Cassandra	Stores and protects object metadata. Note: Data-only Storage Nodes don't host the Cassandra service.
Cassandra Reaper	Performs automatic repairs of object metadata. Note: Data-only Storage Nodes don't host the Cassandra Reaper service.

Service	Key function
Chunk	Manages erasure-coded data and parity fragments.
Data Mover (dmv)	Moves data to Cloud Storage Pools.
Distributed Data Store (DDS)	Monitors object metadata storage. Details Each Storage Node includes the Distributed Data Store (DDS) service. This service interfaces with the Cassandra database to perform background tasks on the object metadata stored in the StorageGRID system. 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).
Identity (idnt)	Federates user identities from LDAP and Active Directory.

Service	Key function
Local Distribution Router (LDR)	Processes object storage protocol requests and manages object data on disk. Details Each <i>combined</i>, <i>data-only</i>, and <i>metadata-only</i> Storage Node includes the Local Distribution Router (LDR) service. This service handles content transport functions, 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 • S3 protocol interface The LDR service also maps each S3 object to its unique UUID. 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. 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 an object store fills to capacity, the remaining object stores continue to store objects until there is no more room on the Storage Node. Metadata protection 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. For details, see Manage object metadata storage.

Service	Key function
Replicated State Machine (RSM)	Ensures that S3 platform services requests are sent to their respective endpoints.
Server Status Monitor (SSM)	Monitors the operating system and underlying hardware.

What is a Gateway Node?

Gateway Nodes provide a dedicated load-balancing interface that S3 client applications can use to connect to StorageGRID. Load balancing maximizes speed and connection capacity by distributing the workload across multiple Storage Nodes. Gateway Nodes are optional.

The StorageGRID Load Balancer service is provided on all Admin Nodes and all Gateway Nodes. It performs Transport Layer Security (TLS) termination of client requests, inspects the requests, and establishes new secure connections to the Storage Nodes. The Load Balancer service seamlessly directs clients to an optimal Storage Node, so that the failure of nodes or even an entire site is transparent.

You configure one or more load balancer endpoints to define the port and network protocol (HTTPS or HTTP) that incoming and outgoing client requests will use to access the Load Balancer services on Gateway and Admin Nodes. The load balancer endpoint also defines the client type (S3), the binding mode, and optionally a list of allowed or blocked tenants. See [Considerations for load balancing](#).

As required, you can group the network interfaces of multiple Gateway Nodes and Admin Nodes into a high availability (HA) group. If the active interface in the HA group fails, a backup interface can manage the client application workload. See [Manage high availability \(HA\) groups](#).

Primary services for Gateway Nodes

The following table shows the primary services for Gateway Nodes; however, this table does not list all node services.

Service	Key function
High Availability	Manages high availability virtual IP addresses for groups of Admin Nodes and Gateway Nodes. Note: This service is also found on Admin Nodes.
Load Balancer	Provides Layer 7 load balancing of S3 traffic from clients to Storage Nodes. This is the recommended load balancing mechanism. Note: This service is also found on Admin Nodes.
Server Status Monitor (SSM)	Monitors the operating system and underlying hardware.

What is an Archive Node?

Support for Archive Nodes has been removed.

For information about Archive Nodes, see [What is an Archive Node \(StorageGRID 11.8 doc site\)](#).


How StorageGRID manages data

What is an object

With object storage, the unit of storage is an object, rather than a file or a block. Unlike the tree-like hierarchy of a file system or block storage, object storage organizes data in a flat, unstructured layout.

Object storage decouples the physical location of the data from the method used to store and retrieve that data.

Each object in an object-based storage system has two parts: object data and object metadata.



What is object data?

Object data might be anything; for example, a photograph, a movie, or a medical record.

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.

Object metadata includes information such as the following:

- 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.
- The current storage location of each object copy or erasure-coded fragment.
- Any user metadata associated with the object.

Object metadata is customizable and expandable, making it flexible for applications to use.

For detailed information about how and where StorageGRID stores object metadata, go to [Manage object metadata storage](#).


How is object data protected?

The StorageGRID system provides you with two mechanisms to protect object data from loss: replication and erasure coding.

Replication

When StorageGRID matches objects to an information lifecycle management (ILM) rule that is configured to create replicated copies, the system creates exact copies of object data and stores them on Storage Nodes or Cloud Storage Pools. ILM rules dictate the number of copies made, where those copies are stored, and for how long they are retained by the system. If a copy is lost, for example, as a result of the loss of a Storage Node, the object is still available if a copy of it exists elsewhere in the StorageGRID system.


In the following example, the Make 2 Copies rule specifies that two replicated copies of each object be placed in a storage pool that contains three Storage Nodes.



Erasure coding

When StorageGRID matches objects to an ILM rule that is configured to create erasure-coded copies, it slices object data into data fragments, computes additional parity fragments, and stores each fragment on a different Storage Node. When an object is accessed, it is reassembled using the stored fragments. If a data or a parity fragment becomes corrupt or lost, the erasure coding algorithm can recreate that fragment using a subset of the remaining data and parity fragments. ILM rules and erasure-coding profiles determine the erasure-coding scheme used.

The following example illustrates the use of erasure coding on an object's data. In this example, the ILM rule uses a 4+2 erasure-coding scheme. Each object is sliced into four equal data fragments, and two parity fragments are computed from the object data. Each of the six fragments is stored on a different Storage Node across three data centers to provide data protection for node failures or site loss.



Related information

- [Manage objects with ILM](#)
- [Use information lifecycle management](#)

The life of an object

An object's life consists of various stages. Each stage represents the operations that occur with the object.

The life of an object includes the operations of ingest, copy management, retrieve, and delete.

- **Ingest:** The process of an S3 client application saving an object over HTTP to the StorageGRID system. At this stage, the StorageGRID system begins to manage the object.
- **Copy management:** The process of managing replicated and erasure-coded copies in StorageGRID, as described by the ILM rules in the active ILM policies. During the copy management stage, StorageGRID protects object data from loss by creating and maintaining the specified number and type of object copies on Storage Nodes or in a Cloud Storage Pool.
- **Retrieve:** The process of a client application accessing an object stored by the StorageGRID system. The client reads the object, which is retrieved from a Storage Node or Cloud Storage Pool.
- **Delete:** The process of removing all object copies from the grid. Objects can be deleted either as a result of the client application sending a delete request to the StorageGRID system, or as a result of an automatic process that StorageGRID performs when the object's lifetime expires.



Related information


- [Manage objects with ILM](#)
- [Use information lifecycle management](#)

Ingest data flow

An ingest, or save, operation consists of a defined data flow between the client and the StorageGRID system.

Data flow

When a client ingests an object to the StorageGRID system, the LDR service on Storage Nodes processes the request and stores the metadata and data to disk.



1. The client application creates the object and sends it to the StorageGRID system through an HTTP PUT request.
2. The object is evaluated against the system's ILM policy.
3. The LDR service saves the object data as a replicated copy or as an erasure-coded copy. (The diagram shows a simplified version of storing a replicated copy to disk.)
4. The LDR service sends the object metadata to the metadata store.
5. The metadata store saves the object metadata to disk.
6. The metadata store propagates copies of object metadata to other Storage Nodes. These copies are also saved to disk.

7. The LDR service returns an HTTP 200 OK response to the client to acknowledge that the object has been ingested.

Copy management

Object data is managed by the active ILM policies and associated ILM rules. ILM rules make replicated or erasure-coded copies to protect object data from loss.


Different types or locations of object copies might be required at different times in the object's life. ILM rules are periodically evaluated to ensure that objects are placed as required.

Object data is managed by the LDR service.

Content protection: replication

If an ILM rule's content placement instructions require replicated copies of object data, copies are made and stored to disk by the Storage Nodes that make up the configured storage pool.

The ILM engine in the LDR service controls replication and ensures that the correct number of copies are stored in the correct locations and for the correct amount of time.




1. The ILM engine queries the ADC service to determine the best destination LDR service within the storage pool specified by the ILM rule. It then sends that LDR service a command to initiate replication.
2. The destination LDR service queries the ADC service for the best source location. It then sends a replication request to the source LDR service.
3. The source LDR service sends a copy to the destination LDR service.
4. The destination LDR service notifies the ILM engine that the object data has been stored.
5. The ILM engine updates the metadata store with object location metadata.

Content protection: erasure coding

If an ILM rule includes instructions to make erasure-coded copies of object data, the applicable erasure-coding scheme breaks object data into data and parity fragments and distributes these fragments across the Storage Nodes configured in the erasure-coding profile.

The ILM engine, which is a component of the LDR service, controls erasure coding and ensures that the erasure-coding profile is applied to object data.




1. The ILM engine queries the ADC service to determine which DDS service can best perform the erasure coding operation. When determined, the ILM engine sends an "initiate" request to that service.
2. The DDS service instructs an LDR to erasure code the object data.
3. The source LDR service sends a copy to the LDR service selected for erasure coding.
4. After creating the appropriate number of parity and data fragments, the LDR service distributes these fragments across the Storage Nodes (Chunk services) that make up the erasure-coding profile's storage pool.
5. The LDR service notifies the ILM engine, confirming that object data is successfully distributed.
6. The ILM engine updates the metadata store with object location metadata.

Content protection: Cloud Storage Pool

If an ILM rule's content placement instructions require that a replicated copy of object data is stored on a Cloud Storage Pool, object data is duplicated to the external S3 bucket or Azure Blob storage container that was specified for the Cloud Storage Pool.

The ILM engine, which is a component of the LDR service, and the Data Mover service control the movement of objects to the Cloud Storage Pool.



1. The ILM engine selects a Data Mover service to replicate to the Cloud Storage Pool.
2. The Data Mover service sends the object data to the Cloud Storage Pool.
3. The Data Mover service notifies the ILM engine that the object data has been stored.
4. The ILM engine updates the metadata store with object location metadata.


Retrieve data flow

A retrieve operation consists of a defined data flow between the StorageGRID system and the client. The system uses attributes to track the retrieval of the object from a Storage Node or, if necessary, a Cloud Storage Pool.

The Storage Node's LDR service queries the metadata store for the location of the object data and retrieves it from the source LDR service. Preferentially, retrieval is from a Storage Node. If the object is not available on a Storage Node, the retrieval request is directed to a Cloud Storage Pool.



If the only object copy is on AWS Glacier storage or the Azure Archive tier, the client application must issue an S3 RestoreObject request to restore a retrievable copy to the Cloud Storage Pool.



1. The LDR service receives a retrieval request from the client application.
2. The LDR service queries the metadata store for the object data location and metadata.
3. LDR service forwards the retrieval request to the source LDR service.
4. The source LDR service returns the object data from the queried LDR service and the system returns the object to the client application.

Delete data flow

All object copies are removed from the StorageGRID system when a client performs a delete operation or when the object's lifetime expires, triggering its automatic removal. There is a defined data flow for object deletion.

Deletion hierarchy

StorageGRID provides several methods for controlling when objects are retained or deleted. Objects can be deleted by client request or automatically. StorageGRID always prioritizes any S3 Object Lock settings over client delete requests, which are prioritized over S3 bucket lifecycle and ILM placement instructions.

- **S3 Object Lock:** If the global S3 Object Lock setting is enabled for the grid, S3 clients can create buckets with S3 Object Lock enabled and then use the S3 REST API to specify retain-until-date and legal hold settings for each object version added to that bucket.
 - An object version that is under a legal hold can't be deleted by any method.
 - Before an object version's retain-until-date is reached, that version can't be deleted by any method.
 - Objects in buckets with S3 Object Lock enabled are retained by ILM "forever". However, after its retain-until-date is reached, an object version can be deleted by a client request or the expiration of the bucket lifecycle.
 - If S3 clients apply a default retain-until-date to the bucket, they don't need to specify a retain-until-date

for each object.


- **Client delete request:** An S3 client can issue a delete object request. When a client deletes an object, all copies of the object are removed from the StorageGRID system.
- **Delete objects in bucket:** Tenant Manager users can use this option to permanently remove all copies of the objects and object versions in selected buckets from the StorageGRID system.
- **S3 bucket lifecycle:** S3 clients can add a lifecycle configuration to their buckets that specifies an Expiration action. If a bucket lifecycle exists, StorageGRID automatically deletes all copies of an object when the date or number of days specified in the Expiration action are met, unless the client deletes the object first.
- **ILM placement instructions:** Assuming that the bucket does not have S3 Object Lock enabled and that there is no bucket lifecycle, StorageGRID automatically deletes an object when the last time period in the ILM rule ends and there are no further placements specified for the object.



When an S3 bucket lifecycle is configured, the lifecycle expiration actions override the ILM policy for objects that match the lifecycle filter. As a result, an object might be retained on the grid even after any ILM instructions for placing the object have lapsed.


See [How objects are deleted](#) for more information.

Data flow for client deletes



1. The LDR service receives a delete request from the client application.
2. The LDR service updates the metadata store so the object looks deleted to client requests, and instructs the ILM engine to remove all copies of object data.
3. The object is removed from the system. The metadata store is updated to remove object metadata.

Data flow for ILM deletes



1. The ILM engine determines that the object needs to be deleted.
2. The ILM engine notifies the metadata store. The metadata store updates object metadata so that the object looks deleted to client requests.
3. The ILM engine removes all copies of the object. The metadata store is updated to remove object metadata.

Information lifecycle management

You use information lifecycle management (ILM) to control the placement, duration, and ingest behavior for all objects in your StorageGRID system. ILM rules determine how StorageGRID stores objects over time. You configure one or more ILM rules and then add them to an ILM policy.

A grid has only one active policy at a time. A policy can contain multiple rules.

ILM rules define:

- Which objects should be stored. A rule can apply to all objects, or you can specify filters to identify which objects a rule applies to. For example, a rule can apply only to objects associated with certain tenant accounts, specific S3 buckets or Swift containers, or specific metadata values.
- The storage type and location. Objects can be stored on Storage Nodes or in Cloud Storage Pools.
- The type of object copies made. Copies can be replicated or erasure-coded.
- For replicated copies, the number of copies made.
- For erasure-coded copies, the erasure-coding scheme used.
- The changes over time to an object's storage location and type of copies.
- How object data is protected as objects are ingested into the grid (synchronous placement or dual commit).

Note that object metadata is not managed by ILM rules. Instead, object metadata is stored in a Cassandra database in what is known as a metadata store. Three copies of object metadata are automatically maintained at each site to protect the data from loss.

Example ILM rule

As an example, an ILM rule could specify the following:

- Apply only to the objects belonging to Tenant A.
- Make two replicated copies of those objects and store each copy at a different site.
- Retain the two copies "forever," which means that StorageGRID will not automatically delete them. Instead, StorageGRID will retain these objects until they are deleted by a client delete request or by the expiration of a bucket lifecycle.
- Use the Balanced option for ingest behavior: the two-site placement instruction is applied as soon as Tenant A saves an object to StorageGRID, unless it is not possible to immediately make both required copies.

For example, if Site 2 is unreachable when Tenant A saves an object, StorageGRID will make two interim copies on Storage Nodes at Site 1. As soon as Site 2 becomes available, StorageGRID will make the required copy at that site.

How an ILM policy evaluates objects

The active ILM policies for your StorageGRID system control the placement, duration, and ingest behavior of all objects.


When clients save objects to StorageGRID, the objects are evaluated against the ordered set of ILM rules in the active policy, as follows:

1. If the filters for the first rule in the policy match an object, the object is ingested according to that rule's ingest behavior and stored according to that rule's placement instructions.
2. If the filters for the first rule don't match the object, the object is evaluated against each subsequent rule in the policy until a match is made.
3. If no rules match an object, the ingest behavior and placement instructions for the default rule in the policy are applied. The default rule is the last rule in a policy and can't use any filters. It must apply to all tenants, all buckets, and all object versions.

Example ILM policy

As an example, an ILM policy could contain three ILM rules that specify the following:

- **Rule 1: Replicated copies for Tenant A**
 - Match all objects belonging to Tenant A.
 - Store these objects as three replicated copies at three sites.
 - Objects belonging to other tenants aren't matched by Rule 1, so they are evaluated against Rule 2.
- **Rule 2: Erasure coding for objects greater than 1 MB**
 - Match all objects from other tenants, but only if they are greater than 1 MB. These larger objects are stored using 6+3 erasure coding at three sites.
 - Does not match objects 1 MB or smaller, so these objects are evaluated against Rule 3.
- **Rule 3: 2 copies 2 data centers (default)**
 - Is the last and default rule in the policy. Does not use filters.
 - Make two replicated copies of all objects not matched by Rule 1 or Rule 2 (objects not belonging to Tenant A that are 1 MB or smaller).



Related information

- [Manage objects with ILM](#)

Explore StorageGRID

Explore the Grid Manager

The Grid Manager is the browser-based graphical interface that allows you to configure, manage, and monitor your StorageGRID system.



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


When you sign in to the Grid Manager, you are connecting to an Admin Node. Each StorageGRID system includes one primary Admin Node and any number of non-primary Admin Nodes. You can connect to any Admin Node, and each Admin Node displays a similar view of the StorageGRID system.

You can access the Grid Manager using a [supported web browser](#).

Grid Manager dashboard

When you first sign in to the Grid Manager, you can use the dashboard to [monitor system activities](#) at a glance.

The dashboard contains information about system health and performance, storage use, ILM processes, S3 operations, and the nodes in the grid. You can [configure the dashboard](#) by selecting from a collection of cards that contain the information you need to effectively monitor your system.



For an explanation of the information shown on each card, select the help icon for that card.

Search field

The **Search** field in the header bar allows you to quickly navigate to a specific page within Grid Manager. For example, you can enter **km** to access the Key management server (KMS) page.

You can use **Search** to find entries in the sidebar of the Grid Manager and on the Configuration, Maintenance, and Support menus. You can also search by name for items like grid nodes and tenant accounts.

Help menu

The help menu provides access to:

- The [FabricPool](#) and [S3 setup](#) wizard
- The StorageGRID documentation center for the current release
- [API documentation](#)
- Information about which version of StorageGRID is currently installed

Alerts menu

The Alerts menu provides an easy-to-use interface for detecting, evaluating, and resolving issues that might occur during StorageGRID operation.

From the Alerts menu, you can do the following to [manage alerts](#):

- Review current alerts
- Review resolved alerts

- Configure silences to suppress alert notifications
- Define alert rules for conditions that trigger alerts
- Configure the email server for alert notifications

Nodes page

The [Nodes page](#) displays information about the entire grid, each site in the grid, and each node at a site.

The Nodes home page displays combined metrics for the entire grid. To view information for a particular site or node, select the site or node.

Nodes

View the list and status of sites and grid nodes.

Q

Total node count: 14

Name  	Type 	Object data used  	Object metadata used  	CPU usage  
StorageGRID Deployment	Grid	0%	0%	—
 Data Center 1	Site	0%	0%	—
 DC1-ADM1	Primary Admin Node	—	—	21%
 DC1-ARC1	Archive Node	—	—	8%
 DC1-G1	Gateway Node	—	—	10%
 DC1-S1	Storage Node	0%	0%	29%

Tenants page

The [Tenants page](#) allows you to [create and monitor the storage tenant accounts](#) for your StorageGRID system. You must create at least one tenant account to specify who can store and retrieve objects and which functionality is available to them.

The Tenants page also provides usage details for each tenant, including the amount of storage used and the number of objects. If you set a quota when you created the tenant, you can see how much of that quota has been used.

Tenants

View information for each tenant account. Depending on the timing of ingests, network connectivity, and node status, the usage data shown might be out of date. To view more recent values, select the tenant name.

[Create](#) [Export to CSV](#) [Actions](#)

Displaying 2 results

☐	Name	Logical space used	Quota utilization	Quota	Object count	Sign in/Copy URL
☐	S3 Tenant	0 bytes	0%	100.00 GB	0	→ 📄
☐	Swift Tenant	0 bytes	0%	100.00 GB	0	→ 📄

← Previous 1 Next →

ILM menu

The [ILM menu](#) allows you to [configure the information lifecycle management \(ILM\) rules and policies](#) that govern data durability and availability. You can also enter an object identifier to view the metadata for that object.

From the ILM menu you can view and manage ILM:

- Rules
- Policies
- Policy tags
- Storage pools
- Storage grades
- Regions
- Object metadata lookup

Configuration menu

The Configuration menu allows you to specify network settings, security settings, system settings, monitoring options, and access control options.

Network tasks

Network tasks include:

- [Managing high availability groups](#)
- [Managing load balancer endpoints](#)
- [Configuring S3 endpoint domain names](#)
- [Managing traffic classification policies](#)
- [Configuring VLAN interfaces](#)

Security tasks

Security tasks include:

- [Managing security certificates](#)
- [Managing internal firewall controls](#)
- [Configuring key management servers](#)
- Configuring security settings including the [TLS and SSH policy](#), [network and object security options](#), and [interface security settings](#).
- Configuring the settings for a [storage proxy](#) or an [admin proxy](#)

System tasks

System tasks include:

- Using [grid federation](#) to clone tenant account information and replicate object data between two StorageGRID systems.
- Optionally, enabling the [Compress stored objects](#) option.
- [Managing S3 Object Lock](#)
- Understanding Storage options such as [object segmentation](#) and [storage volume watermarks](#).
- [Manage erasure-coding profiles](#).

Monitoring tasks

Monitoring tasks include:

- [Configuring audit messages and log destinations](#)
- [Using SNMP monitoring](#)

Access control tasks

Access control tasks include:

- [Managing admin groups](#)
- [Managing admin users](#)
- Changing the [provisioning passphrase](#) or [node console passwords](#)
- [Using identity federation](#)
- [Configuring SSO](#)

Maintenance menu

The Maintenance menu allows you to perform maintenance tasks, system maintenance, and network maintenance.

Tasks

Maintenance tasks include:

- [Decommission operations](#) to remove unused grid nodes and sites

- [Expansion operations](#) to add new grid nodes and sites
- [Grid node recovery procedures](#) to replace a failed node and restore data
- [Rename procedures](#) to change the display names of your grid, sites, and nodes
- [Object existence check operations](#) to verify the existence (although not the correctness) of object data
- Performing a [rolling reboot](#) to restart multiple grid nodes
- [Volume restoration operations](#)

System

System maintenance tasks you can perform include:

- [Viewing StorageGRID license information](#) or [updating license information](#)
- Generating and downloading the [Recovery Package](#)
- Performing StorageGRID software updates, including software upgrades, hotfixes, and updates to the SANtricity OS software on selected appliances
 - [Upgrade procedure](#)
 - [Hotfix procedure](#)
 - [Upgrade SANtricity OS on SG6000 storage controllers using Grid Manager](#)
 - [Upgrade SANtricity OS on SG5700 storage controllers using Grid Manager](#)

Network

Network maintenance tasks you can perform include:

- [Configuring DNS servers](#)
- [Updating Grid Network subnets](#)
- [Managing NTP servers](#)

Support menu

The Support menu provides options that help technical support analyze and troubleshoot your system.

Tools

From the Tools section of the Support menu, you can:

- [Configure AutoSupport](#)
- [Run diagnostics](#) on the current state of the grid
- [Access the Grid Topology tree](#) to view detailed information about grid nodes, services, and attributes
- [Collect log files and system data](#)
- [Review support metrics](#)



The tools available from the **Metrics** option are intended for use by technical support. Some features and menu items within these tools are intentionally non-functional.

Alarms (legacy)

The information about legacy alarms has been removed from this version of the documentation. Refer to [Manage alerts and alarms \(StorageGRID 11.8 documentation\)](#).

Other

From the Other section of the Support menu, you can:

- Manage [link cost](#)
- View [Network Management System \(NMS\)](#) entries
- Manage [storage watermarks](#)

Explore the Tenant Manager

The [Tenant Manager](#) is the browser-based graphical interface that tenant users access to configure, manage, and monitor their storage accounts.



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

When tenant users sign in to the Tenant Manager, they are connecting to an Admin Node.

Tenant Manager dashboard

After a grid administrator creates a tenant account using the Grid Manager or the Grid Management API, tenant users can sign in to the Tenant Manager.

The Tenant Manager dashboard allows tenant users to monitor storage usage at a glance. The Storage usage panel contains a list of the largest buckets (S3) or containers (Swift) for the tenant. The Space used value is the total amount of object data in the bucket or container. The bar chart represents the relative sizes of these buckets or containers.

The value shown above the bar chart is a sum of the space used for all of the tenant's buckets or containers. If the maximum number of gigabytes, terabytes, or petabytes available for the tenant was specified when the account was created, the amount of quota used and remaining are also shown.

Dashboard

16**Buckets**[View buckets](#)**2****Platform services****endpoints**[View endpoints](#)**0****Groups**[View groups](#)**1****User**[View users](#)

Storage usage [?](#)

6.5 TB of 7.2 TB used

0.7 TB (10.1%) remaining



Bucket name	Space used	Number of objects
Bucket-15	969.2 GB	913,425
Bucket-04	937.2 GB	576,806
Bucket-13	815.2 GB	957,389
Bucket-06	812.5 GB	193,843
Bucket-10	473.9 GB	583,245
Bucket-03	403.2 GB	981,226
Bucket-07	362.5 GB	420,726
Bucket-05	294.4 GB	785,190
8 other buckets	1.4 TB	3,007,036

Top buckets by capacity limit usage [?](#)

Bucket name	Usage
Bucket-10	82%
Bucket-03	57%
Bucket-15	20%

Tenant details [?](#)

Name: Tenant02

ID: 3341 1240 0546 8283 2208

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

Storage menu (S3)

The Storage menu is provided for S3 tenant accounts only. This menu allows S3 users to manage access keys; create, manage, and delete buckets; manage platform services endpoints; and view any grid federation connections they are permitted to use.

My access keys

S3 tenant users can manage access keys as follows:

- Users who have the Manage your own S3 credentials permission can create or remove their own S3 access keys.
- Users who have the Root access permission can manage the access keys for the S3 root account, their own account, and all other users. Root access keys also provide full access to the tenant's buckets and objects unless explicitly disabled by a bucket policy.



Managing the access keys for other users takes place from the Access Management menu.

Buckets

S3 tenant users with the appropriate permissions can perform the following tasks for their buckets:

- Create buckets
- Enable S3 Object Lock for a new bucket (assumes that S3 Object Lock is enabled for the StorageGRID)

system)

- Update consistency values
- Enable and disable last access time updates
- Enable or suspend object versioning
- Update S3 Object Lock default retention
- Configure cross-origin resource sharing (CORS)
- Delete all objects in a bucket
- Delete empty buckets
- Use the [S3 Console](#) to manage bucket objects

If a grid administrator has enabled the use of platform services for the tenant account, an S3 tenant user with the appropriate permissions can also perform these tasks:

- Configure S3 event notifications, which can be sent to a destination service that supports the Amazon Simple Notification Service.
- Configure CloudMirror replication, which enables the tenant to automatically replicate objects to an external S3 bucket.
- Configure search integration, which sends object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

Platform services endpoints

If a grid administrator has enabled the use of platform services for the tenant account, an S3 tenant user with the Manage endpoints permission can configure a destination endpoint for each platform service.

Grid federation connections

If a grid administrator has enabled the use of a grid federation connection for the tenant account, an S3 tenant user who has Root access permission can view the connection name, access the bucket details page for each bucket that has cross-grid replication enabled, and view the most recent error to occur when bucket data was being replicated to the other grid in the connection. See [View grid federation connections](#).

Access Management menu

The Access Management menu allows StorageGRID tenants to import user groups from a federated identity source and assign management permissions. Tenants can also manage local tenant groups and users, unless single sign-on (SSO) is in effect for the entire StorageGRID system.

Networking guidelines

Networking guidelines

Use these guidelines to learn about StorageGRID architecture and networking topologies and to learn the requirements for network configuration and provisioning.

About these instructions

These guidelines provide information you can use to create the StorageGRID networking infrastructure before deploying and configuring StorageGRID nodes. Use these guidelines to help ensure that communication can

occur among all the nodes in the grid and between the grid and external clients and services.

External clients and external services need to connect to StorageGRID networks to perform functions such as the following:

- Store and retrieve object data
- Receive email notifications
- Access the StorageGRID management interface (the Grid Manager and Tenant Manager)
- Access the audit share (optional)
- Provide services such as:
 - Network Time Protocol (NTP)
 - Domain name system (DNS)
 - Key Management Server (KMS)

StorageGRID networking must be configured appropriately to handle the traffic for these functions and more.

Before you begin

Configuring the networking for a StorageGRID system requires a high level of experience with Ethernet switching, TCP/IP networking, subnets, network routing, and firewalls.

Before you configure networking, become familiar with StorageGRID architecture as described in [Learn about StorageGRID](#).

After you determine which StorageGRID networks you want to use and how those networks will be configured, you can install and configure the StorageGRID nodes by following the appropriate instructions.

Install appliance nodes

- [Install appliance hardware](#)

Install software-based nodes

- [Install StorageGRID on Red Hat Enterprise Linux](#)
- [Install StorageGRID on Ubuntu or Debian](#)
- [Install StorageGRID on VMware](#)

Configure and administer StorageGRID software

- [Administer StorageGRID](#)
- [Release notes](#)

StorageGRID network types

The grid nodes in a StorageGRID system process *grid traffic*, *admin traffic*, and *client traffic*. You must configure the networking appropriately to manage these three types of traffic and to provide control and security.

Traffic types

Traffic type	Description	Network type
Grid traffic	The internal StorageGRID traffic that travels between all nodes in the grid. All grid nodes must be able to communicate with all other grid nodes over this network.	Grid Network (required)
Admin traffic	The traffic used for system administration and maintenance.	Admin Network (optional), VLAN network (optional)
Client traffic	The traffic that travels between external client applications and the grid, including all object storage requests from S3 clients.	Client Network (optional), VLAN network (optional)

You can configure networking in the following ways:

- Grid Network only
- Grid and Admin Networks
- Grid and Client Networks
- Grid, Admin, and Client Networks

The Grid Network is mandatory and can manage all grid traffic. The Admin and Client Networks can be included at the time of installation or added later to adapt to changes in requirements. Although the Admin Network and Client Network are optional, when you use these networks to handle administrative and client traffic, the Grid Network can be made isolated and secure.

Internal ports are only accessible over the Grid Network. External ports are accessible from all network types. This flexibility provides multiple options for designing a StorageGRID deployment and setting up external IP and port filtering in switches and firewalls. See [internal grid node communications](#) and [external communications](#).

Network interfaces

StorageGRID nodes are connected to each network using the following specific interfaces:

Network	Interface name
Grid Network (required)	eth0
Admin Network (optional)	eth1
Client Network (optional)	eth2

For details about mapping virtual or physical ports to node network interfaces, see the installation instructions:

Software-based nodes

- [Install StorageGRID on Red Hat Enterprise Linux](#)
- [Install StorageGRID on Ubuntu or Debian](#)
- [Install StorageGRID on VMware](#)

Appliance nodes

- [SG6160 storage appliance](#)
- [SGF6112 storage appliance](#)
- [SG6000 storage appliance](#)
- [SG5800 storage appliance](#)
- [SG5700 storage appliance](#)
- [SG110 and SG1100 services appliances](#)
- [SG100 and SG1000 services appliances](#)

Network information for each node

You must configure the following for each network you enable on a node:

- IP address
- Subnet mask
- Gateway IP address

You can only configure one IP address/mask/gateway combination for each of the three networks on each grid node. If you don't want to configure a gateway for a network, you should use the IP address as the gateway address.

High availability groups

High availability (HA) groups provide the ability to add virtual IP (VIP) addresses to the Grid or Client Network interface. For more information, see [Manage high availability groups](#).

Grid Network

The Grid Network is required. It is used for all internal StorageGRID traffic. The Grid Network provides connectivity among all nodes in the grid, across all sites and subnets. All nodes on the Grid Network must be able to communicate with all other nodes. The Grid Network can consist of multiple subnets. Networks containing critical grid services, such as NTP, can also be added as grid subnets.



StorageGRID does not support network address translation (NAT) between nodes.

The Grid Network can be used for all admin traffic and all client traffic, even if the Admin Network and Client Network are configured. The Grid Network gateway is the node default gateway unless the node has the Client Network configured.



When configuring the Grid Network, you must ensure that the network is secured from untrusted clients, such as those on the open internet.

Note the following requirements and details for the Grid Network gateway:

- The Grid Network gateway must be configured if there are multiple grid subnets.
- The Grid Network gateway is the node default gateway until grid configuration is complete.
- Static routes are generated automatically for all nodes to all subnets configured in the global Grid Network Subnet List.

- If a Client Network is added, the default gateway switches from the Grid Network gateway to the Client Network gateway when grid configuration is complete.

Admin Network

The Admin Network is optional. When configured, it can be used for system administration and maintenance traffic. The Admin Network is typically a private network and does not need to be routable between nodes.

You can choose which grid nodes should have the Admin Network enabled on them.

When you use the Admin Network, administrative and maintenance traffic does not need to travel across the Grid Network. Typical uses of the Admin Network include the following:

- Access to the Grid Manager and Tenant Manager user interfaces.
- Access to critical services such as NTP servers, DNS servers, external key management servers (KMS), and Lightweight Directory Access Protocol (LDAP) servers.
- Access to audit logs on Admin Nodes.
- Secure Shell Protocol (SSH) access for maintenance and support.

The Admin Network is never used for internal grid traffic. An Admin Network gateway is provided and allows the Admin Network to communicate with multiple external subnets. However, the Admin Network gateway is never used as the node default gateway.

Note the following requirements and details for the Admin Network gateway:

- The Admin Network gateway is required if connections will be made from outside of the Admin Network subnet or if multiple Admin Network subnets are configured.
- Static routes are created for each subnet configured in the node's Admin Network Subnet List.

Client Network

The Client Network is optional. When configured, it is used to provide access to grid services for client applications such as S3. If you plan to make StorageGRID data accessible to an external resource (for example, a Cloud Storage Pool or the StorageGRID CloudMirror replication service), the external resource can also use the Client Network. Grid nodes can communicate with any subnet reachable through the Client Network gateway.

You can choose which grid nodes should have the Client Network enabled on them. All nodes don't have to be on the same Client Network, and nodes will never communicate with each other over the Client Network. The Client Network does not become operational until grid installation is complete.

For added security, you can specify that a node's Client Network interface be untrusted so that the Client Network will be more restrictive of which connections are allowed. If a node's Client Network interface is untrusted, the interface accepts outbound connections such as those used by CloudMirror replication, but only accepts inbound connections on ports that have been explicitly configured as load balancer endpoints. See [Manage firewall controls](#) and [Configure load balancer endpoints](#).

When you use a Client Network, client traffic does not need to travel across the Grid Network. Grid Network traffic can be separated onto a secure, non-routable network. The following node types are often configured with a Client Network:

- Gateway Nodes, because these nodes provide access to the StorageGRID Load Balancer service and S3 client access to the grid.

- Storage Nodes, because these nodes provide access to the S3 protocol and to Cloud Storage Pools and the CloudMirror replication service.
- Admin Nodes, to ensure that tenant users can connect to the Tenant Manager without needing to use the Admin Network.

Note the following for the Client Network gateway:

- The Client Network gateway is required if the Client Network is configured.
- The Client Network gateway becomes the default route for the grid node when grid configuration is complete.

Optional VLAN networks

As required, you can optionally use virtual LAN (VLAN) networks for client traffic and for some types of admin traffic. Grid traffic, however, can't use a VLAN interface. The internal StorageGRID traffic between nodes must always use the Grid Network on eth0.

To support the use VLANs, you must configure one or more interfaces on a node as trunk interfaces at the switch. You can configure the Grid Network interface (eth0) or the Client Network interface (eth2) to be a trunk, or you can add trunk interfaces to the node.

If eth0 is configured as a trunk, Grid Network traffic flows over the trunk native interface, as configured on the switch. Similarly, if eth2 is configured as a trunk, and the Client Network is also configured on the same node, the Client Network uses the trunk port's native VLAN as configured on the switch.

Only inbound admin traffic, such as used for SSH, Grid Manager, or Tenant Manager traffic, is supported over VLAN networks. Outbound traffic, such as used for NTP, DNS, LDAP, KMS, and Cloud Storage Pools, is not supported over VLAN networks.



VLAN interfaces can be added to Admin Nodes and Gateway Nodes only. You can't use a VLAN interface for client or admin access to Storage Nodes.

See [Configure VLAN interfaces](#) for instructions and guidelines.

VLAN interfaces are only used in HA groups and are assigned VIP addresses on the active node. See [Manage high availability groups](#) for instructions and guidelines.

Network topology examples

Grid Network topology

The simplest network topology is created by configuring the Grid Network only.

When you configure the Grid Network, you establish the host IP address, subnet mask, and Gateway IP address for the eth0 interface for each grid node.

During configuration, you must add all Grid Network subnets to the Grid Network Subnet List (GNSL). This list includes all subnets for all sites, and might also include external subnets that provide access to critical services such as NTP, DNS, or LDAP.

At installation, the Grid Network interface applies static routes for all subnets in the GNSL and sets the node's default route to the Grid Network gateway if one is configured. The GNSL is not required if there is no Client Network and the Grid Network gateway is the node's default route. Host routes to all other nodes in the grid are


also generated.

In this example, all traffic shares the same network, including traffic related to S3 client requests and administrative and maintenance functions.



This topology is appropriate for single-site deployments that aren't externally available, proof-of-concept or test deployments, or when a third-party load balancer acts as the client access boundary. When possible, the Grid Network should be used exclusively for internal traffic. Both the Admin Network and the Client Network have additional firewall restrictions that block external traffic to internal services. Using the Grid Network for external client traffic is supported, but this use offers fewer layers of protection.

Topology example: Grid Network only



Provisioned		
GNSL → 200.200.200.0/24		
Grid Network		
Nodes	IP/mask	Gateway
Admin	200.200.200.32/24	200.200.200.1
Storage	200.200.200.33/24	200.200.200.1
Storage	200.200.200.34/24	200.200.200.1
Storage	200.200.200.35/24	200.200.200.1
Storage	200.200.200.36/24	200.200.200.1
Gateway	200.200.200.37/24	200.200.200.1

System Generated				
Nodes	Routes		Type	From
All	0.0.0.0/0	→ 200.200.200.1	Default	Grid Network gateway
	200.200.200.0/24	→ eth0	Link	Interface IP/mask

Admin Network topology


Having an Admin Network is optional. One way that you can use an Admin Network and a Grid Network is to configure a routable Grid Network and a bounded Admin Network for each node.

When you configure the Admin Network, you establish the host IP address, subnet mask, and Gateway IP address for the eth1 interface for each grid node.

The Admin Network can be unique to each node and can consist of multiple subnets. Each node can be configured with an Admin External Subnet List (AESL). The AESL lists the subnets reachable over the Admin Network for each node. The AESL must also include the subnets of any services the grid will access over the Admin Network, such as NTP, DNS, KMS, and LDAP. Static routes are applied for each subnet in the AESL.

In this example, the Grid Network is used for traffic related to S3 client requests and object management, while the Admin Network is used for administrative functions.

Topology example: Grid and Admin Networks



GNSL → 172.16.0.0/16

AESL (all) → 10.10.1.0/24 10.10.2.0/24 10.10.3.0/24

Nodes	Grid Network		Admin Network	
	IP/mask	Gateway	IP/mask	Gateway
Admin	172.16.200.32/24	172.16.200.1	10.10.1.10/24	10.10.1.1
Storage 1	172.16.200.33/24	172.16.200.1	10.10.1.11/24	10.10.1.1
Storage 2	172.16.200.34/24	172.16.200.1	10.10.3.65/24	10.10.3.1
Storage 3	172.16.200.35/24	172.16.200.1	10.10.1.12/24	10.10.1.1
Storage 4	172.16.200.36/24	172.16.200.1	10.10.1.13/24	10.10.1.1
Gateway	172.16.200.37/24	172.16.200.1	10.10.3.66/24	10.10.3.1

System Generated					
Nodes	Routes			Type	From
All	0.0.0.0/0	→	172.16.200.1	Default	Grid Network gateway
Admin, Storage 1, 3, and 4	172.16.0.0/16	→	eth0	Static	GNSL
	10.10.1.0/24	→	eth1	Link	Interface IP/mask
	10.10.2.0/24	→	10.10.1.1	Static	AESL
	10.10.3.0/24	→	10.10.1.1	Static	AESL
Storage 2, Gateway	172.16.0.0/16	→	eth0	Static	GNSL
	10.10.1.0/24	→	10.10.3.1	Static	AESL
	10.10.2.0/24	→	10.10.3.1	Static	AESL
	10.10.3.0/24	→	eth1	Link	Interface IP/mask

Client Network topology


Having a Client Network is optional. Using a Client Network allows client network traffic (for example, S3) to be separated from grid internal traffic, which allows grid networking to be more secure. Administrative traffic can be handled by either the Client or Grid Network when the Admin Network is not configured.

When you configure the Client Network, you establish the host IP address, subnet mask, and Gateway IP address for the eth2 interface for the configured node. Each node's Client Network can be independent of the Client Network on any other node.

If you configure a Client Network for a node during installation, the node's default gateway switches from the Grid Network gateway to the Client Network gateway when installation is complete. If a Client Network is added later, the node's default gateway switches in the same way.

In this example, the Client Network is used for S3 client requests and for administrative functions, while the Grid Network is dedicated to internal object management operations.

Topology example: Grid and Client Networks



GNSL → 172.16.0.0/16

Nodes	Grid Network	Client Network	
	IP/mask	IP/mask	Gateway
Admin	172.16.200.32/24	37.5.5.10/24	37.5.5.1
Storage	172.16.200.33/24	37.5.5.11/24	37.5.5.1
Storage	172.16.200.34/24	37.5.5.12/24	37.5.5.1
Storage	172.16.200.35/24	37.5.5.13/24	37.5.5.1
Storage	172.16.200.36/24	37.5.5.14/24	37.5.5.1
Gateway	172.16.200.37/24	37.5.5.15/24	37.5.5.1

System Generated

Nodes	Routes		Type	From
All	0.0.0.0/0	→ 37.5.5.1	Default	Client Network gateway
	172.16.0.0/16	→ eth0	Link	Interface IP/mask
	37.5.5.0/24	→ eth2	Link	Interface IP/mask

Related information

[Change node network configuration](#)


Topology for all three networks

You can configure all three networks into a network topology consisting of a private Grid Network, bounded site-specific Admin Networks, and open Client Networks. Using load balancer endpoints and untrusted Client Networks can provide additional security if needed.

In this example:

- The Grid Network is used for network traffic related to internal object management operations.
- The Admin Network is used for traffic related to administrative functions.
- The Client Network is used for traffic related to S3 client requests.

Topology example: Grid, Admin, and Client Networks



Networking requirements

You must verify that the current networking infrastructure and configuration can support the planned StorageGRID network design.

General networking requirements

All StorageGRID deployments must be able to support the following connections.

These connections can occur through the Grid, Admin, or Client Networks, or the combinations of these networks as illustrated in the network topology examples.

- **Management connections:** Inbound connections from an administrator to the node, usually through SSH. Web browser access to the Grid Manager, the Tenant Manager, and the StorageGRID Appliance Installer.
- **NTP server connections:** Outbound UDP connection that receives an inbound UDP response.

At least one NTP server must be reachable by the primary Admin Node.

- **DNS server connections:** Outbound UDP connection that receives an inbound UDP response.
- **LDAP/Active Directory server connections:** Outbound TCP connection from the Identity service on Storage Nodes.
- **AutoSupport:** Outbound TCP connection from the Admin Nodes to either `support.netapp.com` or a customer-configured proxy.
- **External key management server:** Outbound TCP connection from each appliance node with node encryption enabled.
- Inbound TCP connections from S3 clients.
- Outbound requests from StorageGRID platform services such as CloudMirror replication or from Cloud Storage Pools.

If StorageGRID is unable to contact any of the provisioned NTP or DNS servers using the default routing rules, it will automatically attempt contact on all networks (Grid, Admin, and Client) as long as the IP addresses of the DNS and NTP servers are specified. If the NTP or DNS servers can be reached on any network, StorageGRID will automatically create additional routing rules to ensure that network is used for all future attempts to connect to it.



Although you can use these automatically discovered host routes, in general you should manually configure the DNS and NTP routes to ensure connectivity in case automatic discovery fails.

If you aren't ready to configure the optional Admin and Client Networks during deployment, you can configure these networks when you approve grid nodes during the configuration steps. Additionally, you can configure these networks after installation, using the Change IP tool (see [Configure IP addresses](#)).

Only S3 client connections and SSH, Grid Manager, and Tenant Manager administrative connections are supported over VLAN interfaces. Outbound connections, such as to NTP, DNS, LDAP, AutoSupport, and KMS servers, must go over the Client, Admin, or Grid Network interfaces directly. If the interface is configured as a trunk to support VLAN interfaces, this traffic will flow over the interface's native VLAN, as configured at the switch.

Wide Area Networks (WANs) for multiple sites

When configuring a StorageGRID system with multiple sites, the WAN connection between sites must have a minimum bandwidth of 25 Mbit/second in each direction before accounting for client traffic. Data replication or erasure coding between sites, node or site expansion, node recovery, and other operations or configurations will require additional bandwidth.

Actual minimum WAN bandwidth requirements depend on client activity and the ILM protection scheme. For assistance estimating the minimum WAN bandwidth requirements, contact your NetApp Professional Services consultant.

Connections for Admin Nodes and Gateway Nodes

Admin Nodes must always be secured from untrusted clients, such as those on the open internet. You must ensure that no untrusted client can access any Admin Node on the Grid Network, the Admin Network, or the Client Network.

Admin Nodes and Gateway Nodes that you plan to add to high availability groups must be configured with a static IP address. For more information, see [Manage high availability groups](#).

Using network address translation (NAT)

Don't use network address translation (NAT) on the Grid Network between grid nodes or between StorageGRID sites. When you use private IPv4 addresses for the Grid Network, those addresses must be directly routable from every grid node at every site. As required, however, you can use NAT between external clients and grid nodes, such as to provide a public IP address for a Gateway Node. Using NAT to bridge a public network segment is supported only when you employ a tunneling application that is transparent to all nodes in the grid, meaning the grid nodes require no knowledge of public IP addresses.

Network-specific requirements

Follow the requirements for each StorageGRID network type.

Network gateways and routers

- If set, the gateway for a given network must be within the specific network's subnet.
- If you configure an interface using static addressing, you must specify a gateway address other than 0.0.0.0.
- If you don't have a gateway, the best practice is to set the gateway address to be the IP address of the network interface.

Subnets



Each network must be connected to its own subnet that does not overlap with any other network on the node.

The following restrictions are enforced by the Grid Manager during deployment. They are provided here to assist in pre-deployment network planning.

- The subnet mask for any network IP address can't be 255.255.255.254 or 255.255.255.255 (/31 or /32 in CIDR notation).
- The subnet defined by a network interface IP address and subnet mask (CIDR) can't overlap the subnet of any other interface configured on the same node.
- Do not use subnets that contain the following IPv4 addresses for the Grid Network, Admin Network, or Client Network of any node:
 - 192.168.130.101
 - 192.168.131.101
 - 192.168.130.102
 - 192.168.131.102
 - 198.51.100.2

- 198.51.100.4

For example, do not use the following subnet ranges for the Grid Network, Admin Network, or Client Network of any node:

- 192.168.130.0/24 because this subnet range contains the IP addresses 192.168.130.101 and 192.168.130.102
- 192.168.131.0/24 because this subnet range contains the IP addresses 192.168.131.101 and 192.168.131.102
- 198.51.100.0/24 because this subnet range contains the IP addresses 198.51.100.2 and 198.51.100.4
- The Grid Network subnet for each node must be included in the GNSL.
- The Admin Network subnet can't overlap the Grid Network subnet, the Client Network subnet, or any subnet in the GNSL.
- The subnets in the AESL can't overlap with any subnets in the GNSL.
- The Client Network subnet can't overlap the Grid Network subnet, the Admin Network subnet, any subnet in the GNSL, or any subnet in the AESL.

Grid Network

- At deployment time, each grid node must be attached to the Grid Network and must be able to communicate with the primary Admin Node using the networking configuration you specify when deploying the node.
- During normal grid operations, each grid node must be able to communicate with all other grid nodes over the Grid Network.



The Grid Network must be directly routable between each node. Network address translation (NAT) between nodes is not supported.

- If the Grid Network consists of multiple subnets, add them to the Grid Network Subnet List (GNSL). Static routes are created on all nodes for each subnet in the GNSL.
- If the Grid Network interface is configured as a trunk to support VLAN interfaces, the trunk native VLAN must be the VLAN used for Grid Network traffic. All grid nodes must be accessible over the trunk native VLAN.

Admin Network

The Admin Network is optional. If you plan to configure an Admin Network, follow these requirements and guidelines.

Typical uses of the Admin Network include management connections, AutoSupport, KMS, and connections to critical servers such as NTP, DNS, and LDAP if these connections aren't provided through the Grid Network or Client Network.



The Admin Network and AESL can be unique to each node, as long as the desired network services and clients are reachable.



You must define at least one subnet on the Admin Network to enable inbound connections from external subnets. Static routes are automatically generated on each node for each subnet in the AESL.

Client Network

The Client Network is optional. If you plan to configure a Client Network, note the following considerations.

- The Client Network is designed to support traffic from S3 clients. If configured, the Client Network gateway becomes the node's default gateway.
- If you use a Client Network, you can help secure StorageGRID from hostile attacks by accepting inbound client traffic only on explicitly configured load balancer endpoints. See [Configure load balancer endpoints](#).
- If the Client Network interface is configured as a trunk to support VLAN interfaces, consider whether configuring the Client Network interface (eth2) is necessary. If configured, Client Network traffic will flow over the trunk native VLAN, as configured in the switch.

Related information

[Change node network configuration](#)

Deployment-specific networking considerations

Linux deployments

For efficiency, reliability, and security, the StorageGRID system runs on Linux as a collection of container engines. Container engine-related network configuration is not required in a StorageGRID system.

Use a non-bond device, such as a VLAN or virtual Ethernet (veth) pair, for the container network interface. Specify this device as the network interface in the node configuration file.



Don't use bond or bridge devices directly as the container network interface. Doing so could prevent node start-up because of a kernel issue with the use of macvlan with bond and bridge devices in the container namespace.

See the installation instructions for [Red Hat Enterprise Linux](#) or [Ubuntu or Debian](#) deployments.

Host network configuration for container engine deployments

Before starting your StorageGRID deployment on a container engine platform, determine which networks (Grid, Admin, Client) each node will use. You must ensure that each node's network interface is configured on the correct virtual or physical host interface, and that each network has sufficient bandwidth.

Physical hosts

If you are using physical hosts to support grid nodes:

- Make sure all hosts use the same host interface for each node interface. This strategy simplifies host configuration and enables future node migration.
- Obtain an IP address for the physical host itself.



A physical interface on the host can be used by the host itself and one or more nodes running on the host. Any IP addresses assigned to the host or nodes using this interface must be unique. The host and the node can't share IP addresses.

- Open the required ports to the host.

- If you intend to use VLAN interfaces in StorageGRID, the host must have one or more trunk interfaces that provide access to the desired VLANs. These interfaces can be passed into the node container as eth0, eth2, or as additional interfaces. To add trunk or access interfaces, see the following:
 - **RHEL (before installing the node):** [Create node configuration files](#)
 - **Ubuntu or Debian (before installing the node):** [Create node configuration files](#)
 - **RHEL, Ubuntu, or Debian (after installing the node):** [Linux: Add trunk or access interfaces to a node](#)

Minimum bandwidth recommendations

The following table provides the minimum LAN bandwidth recommendations for each type of StorageGRID node and each type of network. You must provision each physical or virtual host with sufficient network bandwidth to meet the aggregate minimum bandwidth requirements for the total number and type of StorageGRID nodes you plan to run on that host.

Type of node	Type of network		
	Grid	Admin	Client
	Minimum LAN bandwidth		
Admin	10 Gbps	1 Gbps	1 Gbps
Gateway	10 Gbps	1 Gbps	10 Gbps
Storage	10 Gbps	1 Gbps	10 Gbps
Archive	10 Gbps	1 Gbps	10 Gbps



This table does not include SAN bandwidth, which is required for access to shared storage. If you are using shared storage accessed over Ethernet (iSCSI or FCoE), you should provision separate physical interfaces on each host to provide sufficient SAN bandwidth. To avoid introducing a bottleneck, SAN bandwidth for a given host should roughly match the aggregate Storage Node network bandwidth for all Storage Nodes running on that host.

Use the table to determine the minimum number of network interfaces to provision on each host, based on the number and type of StorageGRID nodes you plan to run on that host.

For example, to run one Admin Node, one Gateway Node, and one Storage Node on a single host:

- Connect the Grid and Admin Networks on the Admin Node (requires $10 + 1 = 11$ Gbps)
- Connect the Grid and Client Networks on the Gateway Node (requires $10 + 10 = 20$ Gbps)
- Connect the Grid Network on the Storage Node (requires 10 Gbps)

In this scenario, you should provide a minimum of $11 + 20 + 10 = 41$ Gbps of network bandwidth, which could be met by two 40 Gbps interfaces or five 10 Gbps interfaces, potentially aggregated into trunks and then shared by the three or more VLANs carrying the Grid, Admin, and Client subnets local to the physical data center containing the host.

For some recommended ways of configuring physical and network resources on the hosts in your

StorageGRID cluster to prepare for your StorageGRID deployment, see the following:

- [Configure the host network \(Red Hat Enterprise Linux\)](#)
- [Configure the host network \(Ubuntu or Debian\)](#)

Networking and ports for platform services and Cloud Storage Pools

If you plan to use StorageGRID platform services or Cloud Storage Pools, you must configure grid networking and firewalls to ensure that the destination endpoints can be reached.

Networking for platform services

As described in [Manage platform services for tenants](#) and [Manage platform services](#), platform services include external services that provide search integration, event notification, and CloudMirror replication.

Platform services require access from Storage Nodes that host the StorageGRID ADC service to the external service endpoints. Examples for providing access include:

- On the Storage Nodes with ADC services, configure unique Admin Networks with AESL entries that route to the target endpoints.
- Rely on the default route provided by a Client Network. If you use the default route, you can use the [untrusted Client Network feature](#) to restrict inbound connections.

Networking for Cloud Storage Pools

Cloud Storage Pools also require access from Storage Nodes to the endpoints provided by the external service used, such as Amazon S3 Glacier or Microsoft Azure Blob storage. For information, see [What is a Cloud Storage Pool](#).

Ports for platform services and Cloud Storage Pools

By default, platform services and Cloud Storage Pool communications use the following ports:

- **80**: For endpoint URIs that begin with `http`
- **443**: For endpoint URIs that begin with `https`

A different port can be specified when the endpoint is created or edited. See [Network port reference](#).

If you use a non-transparent proxy server, you must also [configure storage proxy settings](#) to allow messages to be sent to external endpoints, such as an endpoint on the internet.

VLANs and platform services and Cloud Storage Pools

You can't use VLAN networks for platform services or Cloud Storage Pools. The destination endpoints must be reachable over the Grid, Admin, or Client Network.

Appliance nodes

You can configure the network ports on StorageGRID appliances to use the port bond modes that meet your requirements for throughput, redundancy, and failover.

The 10/25-GbE ports on the StorageGRID appliances can be configured in Fixed or Aggregate bond mode for

connections to the Grid Network and Client Network.

The 1-GbE Admin Network ports can be configured in Independent or Active-Backup mode for connections to the Admin Network.

See the information about port bond modes for your appliance:

- [Port bond modes \(SG6160\)](#)
- [Port bond modes \(SGF6112\)](#)
- [Port bond modes \(SG6000-CN controller\)](#)
- [Port bond modes \(SG5800 controller\)](#)
- [Port bond modes \(E5700SG controller\)](#)
- [Port bond modes \(SG110 and SG1100\)](#)
- [Port bond modes \(SG100 and SG1000\)](#)

Network installation and provisioning

You must understand how the Grid Network and the optional Admin and Client Networks are used during node deployment and grid configuration.

Initial deployment of a node

When you first deploy a node, you must attach the node to the Grid Network and ensure it has access to the primary Admin Node. If the Grid Network is isolated, you can configure the Admin Network on the primary Admin Node for configuration and installation access from outside the Grid Network.

A Grid Network with a gateway configured becomes the default gateway for a node during deployment. The default gateway allows grid nodes on separate subnets to communicate with the primary Admin Node before the grid has been configured.

If necessary, subnets containing NTP servers or requiring access to the Grid Manager or API can also be configured as grid subnets.

Automatic node registration with primary Admin Node

After the nodes are deployed, they register themselves with the primary Admin Node using the Grid Network. You can then use the Grid Manager, the `configure-storagegrid.py` Python script, or the Installation API to configure the grid and approve the registered nodes. During grid configuration, you can configure multiple grid subnets. Static routes to these subnets through the Grid Network gateway will be created on each node when you complete grid configuration.

Disabling the Admin Network or Client Network

If you want to disable the Admin Network or Client Network, you can remove the configuration from them during the node approval process, or you can use the Change IP tool after installation is complete (see [Configure IP addresses](#)).

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You can't set up DHCP during configuration.



Nodes reboot when the Grid Network configuration is changed by DHCP, which can cause outages if a DHCP change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to reestablish these connections.

Network port reference

Internal grid node communications

The StorageGRID internal firewall allows incoming connections to specific ports on the Grid Network. Connections are also accepted on ports defined by load balancer endpoints.



NetApp recommends that you enable Internet Control Message Protocol (ICMP) traffic between grid nodes. Allowing ICMP traffic can improve failover performance when a grid node can't be reached.

In addition to ICMP and the ports listed in the table, StorageGRID uses the Virtual Router Redundancy Protocol (VRRP). VRRP is an internet protocol that uses IP protocol number 112. StorageGRID uses VRRP in unicast mode only. VRRP is required only if [high availability groups](#) are configured.

Guidelines for Linux-based nodes

If enterprise networking policies restrict access to any of these ports, you can remap ports at deployment time using a deployment configuration parameter. For more information about port remapping and deployment configuration parameters, see:

- [Install StorageGRID on Red Hat Enterprise Linux](#)
- [Install StorageGRID on Ubuntu or Debian](#)

Guidelines for VMware-based nodes

Configure the following ports only if you need to define firewall restrictions that are external to VMware networking.

If enterprise networking policies restrict access to any of these ports, you can remap ports when you deploy nodes using the VMware vSphere Web Client, or by using a configuration file setting when automating grid node deployment. For more information about port remapping and deployment configuration parameters, see [Install StorageGRID on VMware](#).

Guidelines for appliance nodes

If enterprise networking policies restrict access to any of these ports, you can remap ports using the

StorageGRID Appliance Installer. See [Optional: Remap network ports for appliance](#).

StorageGRID internal ports

Port	TCP or UDP	From	To	Details
22	TCP	Primary Admin Node	All nodes	For maintenance procedures, the primary Admin Node must be able to communicate with all other nodes using SSH on port 22. Allowing SSH traffic from other nodes is optional.
80	TCP	Appliances	Primary Admin Node	Used by StorageGRID appliances to communicate with the primary Admin Node to start the installation.
123	UDP	All nodes	All nodes	Network time protocol service. Every node synchronizes its time with every other node using NTP.
443	TCP	All nodes	Primary Admin Node	Used for communicating status to the primary Admin Node during installation and other maintenance procedures.
1055	TCP	All nodes	Primary Admin Node	Internal traffic for installation, expansion, recovery, and other maintenance procedures.
1139	TCP	Storage Nodes	Storage Nodes	Internal traffic between Storage Nodes.
1501	TCP	All nodes	Storage Nodes with ADC	Reporting, auditing, and configuration internal traffic.
1502	TCP	All nodes	Storage Nodes	S3- and Swift-related internal traffic.
1504	TCP	All nodes	Admin Nodes	NMS service reporting and configuration internal traffic.
1505	TCP	All nodes	Admin Nodes	AMS service internal traffic.
1506	TCP	All nodes	All nodes	Server status internal traffic.
1507	TCP	All nodes	Gateway Nodes	Load balancer internal traffic.
1508	TCP	All nodes	Primary Admin Node	Configuration management internal traffic.

Port	TCP or UDP	From	To	Details
1511	TCP	All nodes	Storage Nodes	Metadata internal traffic.
7001	TCP	Storage Nodes	Storage Nodes	Cassandra TLS inter-node cluster communication.
7443	TCP	All nodes	Primary Admin Node	Internal traffic for installation, expansion, recovery, other maintenance procedures, and error reporting.
8011	TCP	All nodes	Primary Admin Node	Internal traffic for installation, expansion, recovery, and other maintenance procedures.
8443	TCP	Primary Admin Node	Appliance nodes	Internal traffic related to the maintenance mode procedure.
9042	TCP	Storage Nodes	Storage Nodes	Cassandra client port.
9999	TCP	All nodes	All nodes	Internal traffic for multiple services. Includes maintenance procedures, metrics, and networking updates.
10226	TCP	Storage Nodes	Primary Admin Node	Used by StorageGRID appliances for forwarding AutoSupport packages from E-Series SANtricity System Manager to the primary Admin Node.
10342	TCP	All nodes	Primary Admin Node	Internal traffic for installation, expansion, recovery, and other maintenance procedures.
18000	TCP	Admin/Storage Nodes	Storage Nodes with ADC	Account service internal traffic.
18001	TCP	Admin/Storage Nodes	Storage Nodes with ADC	Identity Federation internal traffic.
18002	TCP	Admin/Storage Nodes	Storage Nodes	Internal API traffic related to object protocols.
18003	TCP	Admin/Storage Nodes	Storage Nodes with ADC	Platform services internal traffic.

Port	TCP or UDP	From	To	Details
18017	TCP	Admin/Storage Nodes	Storage Nodes	Data Mover service internal traffic for Cloud Storage Pools.
18019	TCP	Storage Nodes	Storage Nodes	Chunk service internal traffic for erasure coding.
18082	TCP	Admin/Storage Nodes	Storage Nodes	S3-related internal traffic.
18083	TCP	All nodes	Storage Nodes	Swift-related internal traffic.
18086	TCP	All grid nodes	All Storage Nodes	Internal traffic related to LDR service.
18200	TCP	Admin/Storage Nodes	Storage Nodes	Additional statistics about client requests.
19000	TCP	Admin/Storage Nodes	Storage Nodes with ADC	Keystone service internal traffic.

Related information

[External communications](#)

External communications

Clients need to communicate with grid nodes to ingest and retrieve content. The ports used depends on the object storage protocols chosen. These ports need to be accessible to the client.

Restricted access to ports

If enterprise networking policies restrict access to any of the ports, you can do one of the following:

- Use [load balancer endpoints](#) to allow access on user-defined ports.
- Remap ports when deploying nodes. However, you should not remap load balancer endpoints. See the information about port remapping for your StorageGRID node:
 - [Port remap keys for StorageGRID on Red Hat Enterprise Linux](#)
 - [Port remap keys for StorageGRID on Ubuntu or Debian](#)
 - [Remap ports for StorageGRID on VMware](#)
 - [Optional: Remap network ports for appliance](#)

Ports used for external communications

The following table shows the ports used for traffic into the nodes.



This list does not include ports that might be configured as [load balancer endpoints](#).

Port	TCP or UDP	Protocol	From	To	Details
22	TCP	SSH	Service laptop	All nodes	SSH or console access is required for procedures with console steps. Optionally, you can use port 2022 instead of 22.
25	TCP	SMTP	Admin Nodes	Email server	Used for alerts and email-based AutoSupport. You can override the default port setting of 25 using the Email Servers page.
53	TCP/ UDP	DNS	All nodes	DNS servers	Used for DNS.
67	UDP	DHCP	All nodes	DHCP service	Optionally used to support DHCP-based network configuration. The dhclient service does not run for statically-configured grids.
68	UDP	DHCP	DHCP service	All nodes	Optionally used to support DHCP-based network configuration. The dhclient service does not run for grids that use static IP addresses.
80	TCP	HTTP	Browser	Admin Nodes	Port 80 redirects to port 443 for the Admin Node user interface.
80	TCP	HTTP	Browser	Appliances	Port 80 redirects to port 8443 for the StorageGRID Appliance Installer.
80	TCP	HTTP	Storage Nodes with ADC	AWS	Used for platform services messages sent to AWS or other external services that use HTTP. Tenants can override the default HTTP port setting of 80 when creating an endpoint.
80	TCP	HTTP	Storage Nodes	AWS	Cloud Storage Pools requests sent to AWS targets that use HTTP. Grid administrators can override the default HTTP port setting of 80 when configuring a Cloud Storage Pool.

Port	TCP or UDP	Protocol	From	To	Details
111	TCP/ UDP	RPCBind	NFS client	Admin Nodes	Used by NFS-based audit export (portmap). Note: This port is required only if NFS-based audit export is enabled. Note: Support for NFS has been deprecated and will be removed in a future release.
123	UDP	NTP	Primary NTP nodes	External NTP	Network time protocol service. Nodes selected as primary NTP sources also synchronize clock times with the external NTP time sources.
161	TCP/ UDP	SNMP	SNMP client	All nodes	Used for SNMP polling. All nodes provide basic information; Admin Nodes also provide alert data. Defaults to UDP port 161 when configured. Note: This port is only required, and is only opened on the node firewall if SNMP is configured. If you plan to use SNMP, you can configure alternate ports. Note: For information about using SNMP with StorageGRID, contact your NetApp account representative.
162	TCP/ UDP	SNMP Notifications	All nodes	Notification destinations	Outbound SNMP notifications and traps default to UDP port 162. Note: This port is only required if SNMP is enabled and notification destinations are configured. If you plan to use SNMP, you can configure alternate ports. Note: For information about using SNMP with StorageGRID, contact your NetApp account representative.
389	TCP/ UDP	LDAP	Storage Nodes with ADC	Active Directory/LDAP	Used for connecting to an Active Directory or LDAP server for Identity Federation.

Port	TCP or UDP	Protocol	From	To	Details
443	TCP	HTTPS	Browser	Admin Nodes	Used by web browsers and management API clients for accessing the Grid Manager and Tenant Manager. Note: If you close Grid Manager ports 443 or 8443, any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list. See Configure firewall controls to configure privileged IP addresses.
443	TCP	HTTPS	Admin Nodes	Active Directory	Used by Admin Nodes connecting to Active Directory if single sign-on (SSO) is enabled.
443	TCP	HTTPS	Storage Nodes with ADC	AWS	Used for platform services messages sent to AWS or other external services that use HTTPS. Tenants can override the default HTTP port setting of 443 when creating an endpoint.
443	TCP	HTTPS	Storage Nodes	AWS	Cloud Storage Pools requests sent to AWS targets that use HTTPS. Grid administrators can override the default HTTPS port setting of 443 when configuring a Cloud Storage Pool.
903	TCP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export (<code>rpc.mountd</code>). Note: This port is required only if NFS-based audit export is enabled. Note: Support for NFS has been deprecated and will be removed in a future release.
2022	TCP	SSH	Service laptop	All nodes	SSH or console access is required for procedures with console steps. Optionally, you can use port 22 instead of 2022.
2049	TCP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export (<code>nfs</code>). Note: This port is required only if NFS-based audit export is enabled. Note: Support for NFS has been deprecated and will be removed in a future release.

Port	TCP or UDP	Protocol	From	To	Details
5353	UDP	mDNS	All nodes	All nodes	Provides the multicast DNS (mDNS) service that is used for full-grid IP changes and for primary Admin Node discovery during installation, expansion, and recovery.
5696	TCP	KMIP	Appliance	KMS	Key Management Interoperability Protocol (KMIP) external traffic from appliances configured for node encryption to the Key Management Server (KMS), unless a different port is specified on the KMS configuration page of the StorageGRID Appliance Installer.
8022	TCP	SSH	Service laptop	All nodes	SSH on port 8022 grants access to the base operating system on appliance and virtual node platforms for support and troubleshooting. This port is not used for Linux-based (bare metal) nodes and is not required to be accessible between grid nodes or during normal operations.
8443	TCP	HTTPS	Browser	Admin Nodes	Optional. Used by web browsers and management API clients for accessing the Grid Manager. Can be used to separate Grid Manager and Tenant Manager communications. Note: If you close Grid Manager ports 443 or 8443, any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list. See Configure firewall controls to configure privileged IP addresses.
9022	TCP	SSH	Service laptop	Appliances	Grants access to StorageGRID appliances in pre-configuration mode for support and troubleshooting. This port is not required to be accessible between grid nodes or during normal operations.
9091	TCP	HTTPS	External Grafana service	Admin Nodes	Used by external Grafana services for secure access to the StorageGRID Prometheus service. Note: This port is required only if certificate-based Prometheus access is enabled.

Port	TCP or UDP	Protocol	From	To	Details
9092	TCP	Kafka	Storage Nodes with ADC	Kafka cluster	Used for platform services messages sent to a Kafka cluster. Tenants can override the default Kafka port setting of 9092 when creating an endpoint.
9443	TCP	HTTPS	Browser	Admin Nodes	Optional. Used by web browsers and management API clients for accessing the Tenant Manager. Can be used to separate Grid Manager and Tenant Manager communications.
18082	TCP	HTTPS	S3 clients	Storage Nodes	S3 client traffic directly to Storage Nodes (HTTPS).
18083	TCP	HTTPS	Swift clients	Storage Nodes	Swift client traffic directly to Storage Nodes (HTTPS).
18084	TCP	HTTP	S3 clients	Storage Nodes	S3 client traffic directly to Storage Nodes (HTTP).
18085	TCP	HTTP	Swift clients	Storage Nodes	Swift client traffic directly to Storage Nodes (HTTP).
23000-23999	TCP	HTTPS	All nodes on the source grid for cross-grid replication	Admin Nodes and Gateway Nodes on the destination grid for cross-grid replication	This range of ports is reserved for grid federation connections. Both grids in a given connection use the same port.

Quick start for StorageGRID

Follow these high-level steps to configure and use any StorageGRID system.

1

Learn, plan, and collect data

Work with your NetApp account representative to understand the options and to plan your new StorageGRID system. Consider these types of questions:

- How much object data do you expect to store initially and over time?
- How many sites do you need?
- How many and what types of nodes do you need at each site?
- Which StorageGRID networks will you use?

- Who will use your grid to store objects? Which applications will they use?
- Do you have any special security or storage requirements?
- Do you need to comply with any legal or regulatory requirements?

Optionally, work with your NetApp Professional Services consultant to access the NetApp ConfigBuilder tool to complete a configuration workbook for use when installing and deploying your new system. You can also use this tool to help automate the configuration of any StorageGRID appliance. See [Automate appliance installation and configuration](#).

Review [Learn about StorageGRID](#) and the [Networking guidelines](#).

2

Install nodes

A StorageGRID system consists of individual hardware-based and software-based nodes. You first install the hardware for each appliance node and configure each Linux or VMware host.

To complete the installation, you install StorageGRID software on each appliance or software host and connect the nodes into a grid. During this step, you provide site and node names, subnet details, and the IP addresses for your NTP and DNS servers.

Learn how:

- [Install appliance hardware](#)
- [Install StorageGRID on Red Hat Enterprise Linux](#)
- [Install StorageGRID on Ubuntu or Debian](#)
- [Install StorageGRID on VMware](#)

3

Sign in and check system health

As soon as you install the primary Admin Node, you can sign in to the Grid Manager. From there, you can review the general health of your new system, enable AutoSupport and alert emails, and set up S3 endpoint domain names.

Learn how:

- [Sign in to the Grid Manager](#)
- [Monitor system health](#)
- [Configure AutoSupport](#)
- [Set up email notifications for alerts](#)
- [Configure S3 endpoint domain names](#)

4

Configure and manage

The configuration tasks you need to perform for a new StorageGRID system depend on how you will use your grid. At a minimum, you set up system access; use the FabricPool and S3 wizards; and manage various storage and security settings.

Learn how:

- [Control StorageGRID access](#)
- [Use S3 setup wizard](#)
- [Use FabricPool setup wizard](#)
- [Manage security](#)
- [System hardening](#)

5

Set up ILM

You control the placement and duration of every object in your StorageGRID system by configuring an information lifecycle management (ILM) policy that consists of one or more ILM rules. The ILM rules instruct StorageGRID how to create and distribute copies of object data and how to manage those copies over time.

Learn how: [Manage objects with ILM](#)

6

Use StorageGRID

After the initial configuration is complete, StorageGRID tenant accounts can use S3 client applications to ingest, retrieve, and delete objects.

Learn how:

- [Use a tenant account](#)
- [Use the S3 REST API](#)

7

Monitor and troubleshoot

When your system is up and running, you should monitor its activities on a regular basis and troubleshoot and resolve any alerts. You might also want to configure an external syslog server, use SNMP monitoring, or collect additional data.

Learn how:

- [Monitor StorageGRID](#)
- [Troubleshoot StorageGRID](#)

8

Expand, maintain, and recover

You can add nodes or sites to expand the capacity or functionality of your system. You can also perform various maintenance procedures to recover from failures or to keep your StorageGRID system up-to-date and performing efficiently.

Learn how:

- [Expand a grid](#)
- [Maintain your grid](#)
- [Recover nodes](#)

Install, upgrade, and hotfix StorageGRID

StorageGRID appliances

Go to [StorageGRID Appliance Documentation](#) to learn how to install, configure, and maintain StorageGRID storage and services appliances.

Install StorageGRID on Red Hat Enterprise Linux

Quick start for installing StorageGRID on Red Hat Enterprise Linux

Follow these high-level steps to install a Red Hat Enterprise Linux (RHEL) Linux StorageGRID node.

1

Preparation

- Learn about [StorageGRID architecture and network topology](#).
- Learn about the specifics of [StorageGRID networking](#).
- Gather and prepare the [Required information and materials](#).
- Prepare the required [CPU and RAM](#).
- Provide for [storage and performance requirements](#).
- [Prepare the Linux servers](#) that will host your StorageGRID nodes.

2

Deployment

Deploy grid nodes. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.

- To deploy software-based grid nodes on the hosts you prepared in step 1, use the Linux command line and [node configuration files](#).
- To deploy StorageGRID appliance nodes, follow the [Quick start for hardware installation](#).

3

Configuration

When all nodes have been deployed, use the Grid Manager to [configure the grid and complete the installation](#).

Automate the installation

To save time and provide consistency, you can automate the installation of the StorageGRID host service and the configuration of grid nodes.

- Use a standard orchestration framework such as Ansible, Puppet, or Chef to automate:
 - Installation of RHEL
 - Configuration of networking and storage

- Installation of the container engine and the StorageGRID host service
- Deployment of virtual grid nodes

See [Automate the installation and configuration of the StorageGRID host service](#).

- After you deploy grid nodes, [automate the configuration of the StorageGRID system](#) using the Python configuration script provided in the installation archive.
- [Automate the installation and configuration of appliance grid nodes](#)
- If you are an advanced developer of StorageGRID deployments, automate the installation of grid nodes by using the [installation REST API](#).

Plan and prepare for installation on Red Hat

Required information and materials

Before you install StorageGRID, gather and prepare the required information and materials.

Required information

Network plan

Which networks you intend to attach to each StorageGRID node. StorageGRID supports multiple networks for traffic separation, security, and administrative convenience.

See the StorageGRID [Networking guidelines](#).

Network information

IP addresses to assign to each grid node and the IP addresses of the DNS and NTP servers.

Servers for grid nodes

Identify a set of servers (physical, virtual, or both) that, in aggregate, provide sufficient resources to support the number and type of StorageGRID nodes you plan to deploy.



If your StorageGRID installation will not use StorageGRID appliance (hardware) Storage Nodes, you must use hardware RAID storage with battery-backed write cache (BBWC). StorageGRID does not support the use of virtual storage area networks (vSANs), software RAID, or no RAID protection.

Node migration (if needed)

Understand the [requirements for node migration](#), if you want to perform scheduled maintenance on physical hosts without any service interruption.

Related information

[NetApp Interoperability Matrix Tool](#)

Required materials

NetApp StorageGRID license

You must have a valid, digitally signed NetApp license.



A non-production license, which can be used for testing and proof of concept grids, is included in the StorageGRID installation archive.

StorageGRID installation archive

Download the [StorageGRID installation archive](#) and extract the files.

Service laptop

The StorageGRID system is installed through a service laptop.

The service laptop must have:

- Network port
- SSH client (for example, PuTTY)
- [Supported web browser](#)

StorageGRID documentation

- [Release notes](#)
- [Instructions for administering StorageGRID](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archive and extract the required files. Optionally, you can manually verify the files in the installation package.

Steps

1. Go to the [NetApp Downloads page for StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.
3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead statement appears, read it and select the checkbox.



You must apply any required hotfixes after you install the StorageGRID release. For more information, see the [hotfix procedure in the recovery and maintenance instructions](#).

5. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column, select the .tgz or .zip installation archive for Red Hat Enterprise Linux.



Select the .zip file if you are running Windows on the service laptop.

7. Save the installation archive.
8. If you need to verify the installation archive:
 - a. Download the StorageGRID code signature verification package. The file name for this package uses the format `StorageGRID_<version-number>_Code_Signature_Verification_Package.tar.gz`, where <version-number> is the StorageGRID software version.
 - b. Follow the steps to [manually verify the installation files](#).

9. Extract the files from the installation archive.

10. Choose the files you need.

The files you need depend on your planned grid topology and how you will deploy your StorageGRID system.



The paths listed in the table are relative to the top-level directory installed by the extracted installation archive

Path and file name	Description
<code>./rpms/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./rpms/NLF000000.txt</code>	A free license that does not provide any support entitlement for the product.
<code>./rpms/StorageGRID-Webscale-Images-version-SHA.rpm</code>	RPM package for installing the StorageGRID node images on your RHEL hosts.
<code>./rpms/StorageGRID-Webscale-Service-version-SHA.rpm</code>	RPM package for installing the StorageGRID host service on your RHEL hosts.
Deployment scripting tool	Description
<code>./rpms/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.
<code>./rpms/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./rpms/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./rpms/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate integration.
<code>./rpms/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./rpms/extras/ansible</code>	Example Ansible role and playbook for configuring RHEL hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.

Path and file name	Description
<code>./rpms/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
<code>./rpms/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.
<code>./rpms/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

Manually verify installation files (optional)

If necessary, you can manually verify the files in the StorageGRID installation archive.

Before you begin

You have [downloaded the verification package](#) from the [NetApp Downloads page for StorageGRID](#).

Steps

1. Extract the artifacts from the verification package:

```
tar -xf StorageGRID_11.9.0_Code_Signature_Verification_Package.tar.gz
```

2. Ensure that these artifacts were extracted:

- Leaf certificate: `Leaf-Cert.pem`
- Certificate chain: `CA-Int-Cert.pem`
- Time stamp response chain: `TS-Cert.pem`
- Checksum file: `sha256sum`
- Checksum signature: `sha256sum.sig`
- Time stamp response file: `sha256sum.sig.tsr`

3. Use the chain to verify the leaf certificate is valid.

Example: `openssl verify -CAfile CA-Int-Cert.pem Leaf-Cert.pem`

Expected output: `Leaf-Cert.pem: OK`

4. If step 2 failed because of an expired leaf certificate, use the `tsr` file to verify.

Example: `openssl ts -CAfile CA-Int-Cert.pem -untrusted TS-Cert.pem -verify -data sha256sum.sig -in sha256sum.sig.tsr`

Expected output includes: Verification: OK

5. Create a public key file from the leaf certificate.

Example: `openssl x509 -pubkey -noout -in Leaf-Cert.pem > Leaf-Cert.pub`

Expected output: *none*

6. Use the public key to verify the sha256sum file against sha256sum.sig.

Example: `openssl dgst -sha256 -verify Leaf-Cert.pub -signature sha256sum.sig sha256sum`

Expected output: Verified OK

7. Verify the sha256sum file content against newly created checksums.

Example: `sha256sum -c sha256sum`

Expected output: *<filename>*: OK
<filename> is the name of the archive file you downloaded.

8. [Complete the remaining steps](#) to extract and choose the appropriate files from the installation archive.

Software requirements for Red Hat Enterprise Linux

You can use a virtual machine to host any type of StorageGRID node. You need one virtual machine for each grid node.

To install StorageGRID on Red Hat Enterprise Linux (RHEL), you must install some third-party software packages. Some supported Linux distributions don't contain these packages by default. The software package versions that StorageGRID installations are tested on include those listed on this page.

If you select a Linux distribution and container runtime installation option that requires any of these packages, and they are not installed automatically by the Linux distribution, install one of the versions listed here if available from your provider or the supporting vendor for your Linux distribution. Otherwise, use the default package versions available from your vendor.

All installation options require either Podman or Docker. Do not install both packages. Install only the package required by your installation option.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

Python versions tested

- 3.5.2-2
- 3.6.8-2
- 3.6.8-38

- 3.6.9-1
- 3.7.3-1
- 3.8.10-0
- 3.9.2-1
- 3.9.10-2
- 3.9.16-1
- 3.10.6-1
- 3.11.2-6

Podman versions tested

- 3.2.3-0
- 3.4.4+ds1
- 4.1.1-7
- 4.2.0-11
- 4.3.1+ds1-8+b1
- 4.4.1-8
- 4.4.1-12

Docker versions tested



Docker support is deprecated and will be removed in a future release.

- Docker-CE 20.10.7
- Docker-CE 20.10.20-3
- Docker-CE 23.0.6-1
- Docker-CE 24.0.2-1
- Docker-CE 24.0.4-1
- Docker-CE 24.0.5-1
- Docker-CE 24.0.7-1
- 1.5-2

CPU and RAM requirements

Before installing StorageGRID software, verify and configure the hardware so that it is ready to support the StorageGRID system.

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: Dependent on the total RAM available and the amount of non-StorageGRID software running on the system
 - Generally, at least 24 GB per node, and 2 to 16 GB less than the total system RAM

- A minimum of 64 GB for each tenant that will have approximately 5,000 buckets

Software-based metadata-only node resources must match the existing Storage Nodes resources. For example:

- If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)
- If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0).

When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.

Ensure that the number of StorageGRID nodes you plan to run on each physical or virtual host does not exceed the number of CPU cores or the physical RAM available. If the hosts aren't dedicated to running StorageGRID (not recommended), be sure to consider the resource requirements of the other applications.



Monitor your CPU and memory usage regularly to ensure that these resources continue to accommodate your workload. For example, doubling the RAM and CPU allocation for virtual Storage Nodes would provide similar resources to those provided for StorageGRID appliance nodes. Additionally, if the amount of metadata per node exceeds 500 GB, consider increasing the RAM per node to 48 GB or more. For information about managing object metadata storage, increasing the Metadata Reserved Space setting, and monitoring CPU and memory usage, see the instructions for [administering](#), [monitoring](#), and [upgrading](#) StorageGRID.

If hyperthreading is enabled on the underlying physical hosts, you can provide 8 virtual cores (4 physical cores) per node. If hyperthreading is not enabled on the underlying physical hosts, you must provide 8 physical cores per node.

If you are using virtual machines as hosts and have control over the size and number of VMs, you should use a single VM for each StorageGRID node and size the VM accordingly.

For production deployments, you should not run multiple Storage Nodes on the same physical storage hardware or virtual host. Each Storage Node in a single StorageGRID deployment should be in its own isolated failure domain. You can maximize the durability and availability of object data if you ensure that a single hardware failure can only impact a single Storage Node.

See also [Storage and performance requirements](#).

Storage and performance requirements

You must understand the storage requirements for StorageGRID nodes, so you can provide enough space to support the initial configuration and future storage expansion.

StorageGRID nodes require three logical categories of storage:

- **Container pool** — Performance-tier (10K SAS or SSD) storage for the node containers, which will be

assigned to the container engine storage driver when you install and configure the container engine on the hosts that will support your StorageGRID nodes.

- **System data** — Performance-tier (10K SAS or SSD) storage for per-node persistent storage of system data and transaction logs, which the StorageGRID host services will consume and map into individual nodes.
- **Object data** — Performance-tier (10K SAS or SSD) storage and capacity-tier (NL-SAS/SATA) bulk storage for the persistent storage of object data and object metadata.

You must use RAID-backed block devices for all storage categories. Non-redundant disks, SSDs, or JBODs aren't supported. You can use shared or local RAID storage for any of the storage categories; however, if you want to use the node migration capability in StorageGRID, you must store both system data and object data on shared storage. For more information, see [Node container migration requirements](#).

Performance requirements

The performance of the volumes used for the container pool, system data, and object metadata significantly impacts the overall performance of the system. You should use performance-tier (10K SAS or SSD) storage for these volumes to ensure adequate disk performance in terms of latency, input/output operations per second (IOPS), and throughput. You can use capacity-tier (NL-SAS/SATA) storage for the persistent storage of object data.

The volumes used for the container pool, system data, and object data must have write-back caching enabled. The cache must be on a protected or persistent media.

Requirements for hosts that use NetApp ONTAP storage

If the StorageGRID node uses storage assigned from a NetApp ONTAP system, confirm that the volume does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

Number of hosts required

Each StorageGRID site requires a minimum of three Storage Nodes.



In a production deployment, don't run more than one Storage Node on a single physical or virtual host. Using a dedicated host for each Storage Node provides an isolated failure domain.

Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same hosts, or they can be deployed on their own dedicated hosts as required.

Number of storage volumes for each host

The following table shows the number of storage volumes (LUNs) required for each host and the minimum size required for each LUN, based on which nodes will be deployed on that host.

The maximum tested LUN size is 39 TB.



These numbers are for each host, not for the entire grid.

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
Container engine storage pool	Container pool	1	Total number of nodes × 100 GB
/var/local volume	System data	1 for each node on this host	90 GB
Storage Node	Object data	3 for each Storage Node on this host Note: A software-based Storage Node can have 1 to 48 storage volumes; at least 3 storage volumes are recommended.	12 TB (4 TB/LUN) See Storage requirements for Storage Nodes for more information.
Storage Node (metadata-only)	Object metadata	1	4 TB See Storage requirements for Storage Nodes for more information. Note: Only one rangedb is required for metadata-only Storage Nodes.
Admin Node audit logs	System data	1 for each Admin Node on this host	200 GB
Admin Node tables	System data	1 for each Admin Node on this host	200 GB



Depending on the audit level configured, the size of user inputs such as S3 object key name, and how much audit log data you need to preserve, you might need to increase the size of the audit log LUN on each Admin Node. Generally, a grid generates approximately 1 KB of audit data per S3 operation, which would mean that a 200 GB LUN would support 70 million operations per day or 800 operations per second for two to three days.

Minimum storage space for a host

The following table shows the minimum storage space required for each type of node. You can use this table to determine the minimum amount of storage you must provide to the host in each storage category, based on which nodes will be deployed on that host.



Disk snapshots can't be used to restore grid nodes. Instead, refer to the [grid node recovery](#) procedures for each type of node.

Type of node	Container pool	System data	Object data
Storage Node	100 GB	90 GB	4,000 GB

Type of node	Container pool	System data	Object data
Admin Node	100 GB	490 GB (3 LUNs)	<i>not applicable</i>
Gateway Node	100 GB	90 GB	<i>not applicable</i>

Example: Calculating the storage requirements for a host

Suppose you plan to deploy three nodes on the same host: one Storage Node, one Admin Node, and one Gateway Node. You should provide a minimum of nine storage volumes to the host. You will need a minimum of 300 GB of performance-tier storage for the node containers, 670 GB of performance-tier storage for system data and transaction logs, and 12 TB of capacity-tier storage for object data.

Type of node	LUN purpose	Number of LUNs	LUN size
Storage Node	Container engine storage pool	1	300 GB (100 GB/node)
Storage Node	<code>/var/local</code> volume	1	90 GB
Storage Node	Object data	3	12 TB (4 TB/LUN)
Admin Node	<code>/var/local</code> volume	1	90 GB
Admin Node	Admin Node audit logs	1	200 GB
Admin Node	Admin Node tables	1	200 GB
Gateway Node	<code>/var/local</code> volume	1	90 GB
Total		9	Container pool: 300 GB System data: 670 GB Object data: 12,000 GB


Storage requirements for Storage Nodes

A software-based Storage Node can have 1 to 48 storage volumes; 3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.



An appliance Storage Node can also have up to 48 storage volumes.

As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. Any remaining space on storage volume 0 and any other storage volumes in the Storage Node are used exclusively for 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. The three copies of object metadata are evenly distributed across all Storage Nodes at each site.

When installing a grid with metadata-only Storage Nodes, the grid must also contain a minimum number of nodes for object storage. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

- For a single-site grid, at least two Storage Nodes are configured for objects and metadata.
- For a multi-site grid, at least one Storage Node per site are configured for objects and metadata.

When you assign space to volume 0 of a new Storage Node, you must ensure there is adequate space for that node's portion of all object metadata.

- At a minimum, you must assign at least 4 TB to volume 0.



If you use only one storage volume for a Storage Node and you assign 4 TB or less to the volume, the Storage Node might enter the storage read-only state on startup and store object metadata only.



If you assign less than 500 GB to volume 0 (non-production use only), 10% of the storage volume's capacity is reserved for metadata.

- Software-based metadata-only node resources must match the existing Storage Nodes resources. For example:
 - If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)

- If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0)).

When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.

- If you are installing a new system (StorageGRID 11.6 or higher) and each Storage Node has 128 GB or more of RAM, assign 8 TB or more to volume 0. Using a larger value for volume 0 can increase the space allowed for metadata on each Storage Node.
- When configuring different Storage Nodes for a site, use the same setting for volume 0 if possible. If a site contains Storage Nodes of different sizes, the Storage Node with the smallest volume 0 will determine the metadata capacity of that site.

For details, go to [Manage object metadata storage](#).

Node container migration requirements

The node migration feature allows you to manually move a node from one host to another. Typically, both hosts are in the same physical data center.

Node migration allows you to perform physical host maintenance without disrupting grid operations. You move all StorageGRID nodes, one at a time, to another host before taking the physical host offline. Migrating nodes requires only a short downtime for each node and should not affect operation or availability of grid services.

If you want to use the StorageGRID node migration feature, your deployment must meet additional requirements:

- Consistent network interface names across hosts in a single physical data center
- Shared storage for StorageGRID metadata and object repository volumes that is accessible by all hosts in a single physical data center. For example, you might use NetApp E-Series storage arrays.

If you are using virtual hosts and the underlying hypervisor layer supports VM migration, you might want to use this capability instead of the node migration feature in StorageGRID. In this case, you can ignore these additional requirements.

Before performing migration or hypervisor maintenance, shut down the nodes gracefully. See the instructions for [shutting down a grid node](#).

VMware Live Migration not supported

When performing bare-metal installation on VMware VMs, OpenStack Live Migration and VMware live vMotion cause the virtual machine clock time to jump and aren't supported for grid nodes of any type. Though rare, incorrect clock times can result in loss of data or configuration updates.

Cold migration is supported. In cold migration, you shut down the StorageGRID nodes before migrating them between hosts. See the instructions for [shutting down a grid node](#).

Consistent network interface names

To move a node from one host to another, the StorageGRID host service needs to have some confidence that the external network connectivity the node has at its current location can be duplicated at the new location. It gets this confidence through the use of consistent network interface names in the hosts.

Suppose, for example, that StorageGRID NodeA running on Host1 has been configured with the following interface mappings:

eth0 → bond0.1001

eth1 → bond0.1002

eth2 → bond0.1003

The lefthand side of the arrows corresponds to the traditional interfaces as viewed from within a StorageGRID container (that is, the Grid, Admin, and Client Network interfaces, respectively). The righthand side of the arrows corresponds to the actual host interfaces providing these networks, which are three VLAN interfaces subordinate to the same physical interface bond.

Now, suppose you want to migrate NodeA to Host2. If Host2 also has interfaces named bond0.1001, bond0.1002, and bond0.1003, the system will allow the move, assuming that the like-named interfaces will provide the same connectivity on Host2 as they do on Host1. If Host2 does not have interfaces with the same names, the move will not be allowed.

There are many ways to achieve consistent network interface naming across multiple hosts; see [Configuring the host network](#) for some examples.

Shared storage

To achieve rapid, low-overhead node migrations, the StorageGRID node migration feature does not physically move node data. Instead, node migration is performed as a pair of export and import operations, as follows:

1. During the "node export" operation, a small amount of persistent state data is extracted from the node container running on HostA and cached on that node's system data volume. Then, the node container on HostA is deinstantiated.
2. During the "node import" operation, the node container on HostB that uses the same network interface and block storage mappings that were in effect on HostA is instantiated. Then, the cached persistent state data is inserted into the new instance.

Given this mode of operation, all of the node's system data and object storage volumes must be accessible from both HostA and HostB for the migration to be allowed, and to work. In addition, they must have been mapped into the node using names that are guaranteed to refer to the same LUNs on HostA and HostB.

The following example shows one solution for block device mapping for a StorageGRID Storage Node, where DM multipathing is in use on the hosts, and the alias field has been used in `/etc/multipath.conf` to provide consistent, friendly block device names available on all hosts.

`/var/local` → `/dev/mapper/sgws-sn1-var-local`
`rangedb0` → `/dev/mapper/sgws-sn1-rangedb0`
`rangedb1` → `/dev/mapper/sgws-sn1-rangedb1`
`rangedb2` → `/dev/mapper/sgws-sn1-rangedb2`
`rangedb3` → `/dev/mapper/sgws-sn1-rangedb3`

Prepare the hosts (Red Hat)

How host-wide settings change during installation

On bare metal systems, StorageGRID makes some changes to host-wide `sysctl` settings.

The following changes are made:

```
# Recommended Cassandra setting: CASSANDRA-3563, CASSANDRA-13008, DataStax
documentation
vm.max_map_count = 1048575

# core file customization
# Note: for cores generated by binaries running inside containers, this
# path is interpreted relative to the container filesystem namespace.
# External cores will go nowhere, unless /var/local/core also exists on
# the host.
kernel.core_pattern = /var/local/core/%e.core.%p

# Set the kernel minimum free memory to the greater of the current value
or
# 512MiB if the host has 48GiB or less of RAM or 1.83GiB if the host has
more than 48GiB of RTAM
vm.min_free_kbytes = 524288

# Enforce current default swappiness value to ensure the VM system has
some
# flexibility to garbage collect behind anonymous mappings. Bump
watermark_scale_factor
# to help avoid OOM conditions in the kernel during memory allocation
bursts. Bump
# dirty_ratio to 90 because we explicitly fsync data that needs to be
persistent, and
```

```

# so do not require the dirty_ratio safety net. A low dirty_ratio combined
with a large
# working set (nr_active_pages) can cause us to enter synchronous I/O mode
unnecessarily,
# with deleterious effects on performance.
vm.swappiness = 60
vm.watermark_scale_factor = 200
vm.dirty_ratio = 90

# Turn off slow start after idle
net.ipv4.tcp_slow_start_after_idle = 0

# Tune TCP window settings to improve throughput
net.core.rmem_max = 8388608
net.core.wmem_max = 8388608
net.ipv4.tcp_rmem = 4096 524288 8388608
net.ipv4.tcp_wmem = 4096 262144 8388608
net.core.netdev_max_backlog = 2500

# Turn on MTU probing
net.ipv4.tcp_mtu_probing = 1

# Be more liberal with firewall connection tracking
net.ipv4.netfilter.ip_conntrack_tcp_be_liberal = 1

# Reduce TCP keepalive time to reasonable levels to terminate dead
connections
net.ipv4.tcp_keepalive_time = 270
net.ipv4.tcp_keepalive_probes = 3
net.ipv4.tcp_keepalive_intvl = 30

# Increase the ARP cache size to tolerate being in a /16 subnet
net.ipv4.neigh.default.gc_thresh1 = 8192
net.ipv4.neigh.default.gc_thresh2 = 32768
net.ipv4.neigh.default.gc_thresh3 = 65536
net.ipv6.neigh.default.gc_thresh1 = 8192
net.ipv6.neigh.default.gc_thresh2 = 32768
net.ipv6.neigh.default.gc_thresh3 = 65536

# Disable IP forwarding, we are not a router
net.ipv4.ip_forward = 0

# Follow security best practices for ignoring broadcast ping requests
net.ipv4.icmp_echo_ignore_broadcasts = 1

# Increase the pending connection and accept backlog to handle larger
connection bursts.

```



```
net.core.somaxconn=4096
net.ipv4.tcp_max_syn_backlog=4096
```

Install Linux

You must install StorageGRID on all Red Hat Enterprise Linux grid hosts. For a list of supported versions, use the NetApp Interoperability Matrix Tool.

Before you begin

Ensure your operating system meets StorageGRID's minimum kernel version requirements, as listed below. Use the command `uname -r` to get your operating system's kernel version, or consult with your OS vendor.

Red Hat Enterprise Linux version	Minimum kernel version	Kernel package name
8.8 (deprecated)	4.18.0-477.10.1.el8_8.x86_64	kernel-4.18.0-477.10.1.el8_8.x86_64
8.10	4.18.0-553.el8_10.x86_64	kernel-4.18.0-553.el8_10.x86_64
9.0 (deprecated)	5.14.0-70.22.1.el9_0.x86_64	kernel-5.14.0-70.22.1.el9_0.x86_64
9.2 (deprecated)	5.14.0-284.11.1.el9_2.x86_64	kernel-5.14.0-284.11.1.el9_2.x86_64
9.4	5.14.0-427.18.1.el9_4.x86_64	kernel-5.14.0-427.18.1.el9_4.x86_64
9.6	5.14.0-570.18.1.el9_6.x86_64	kernel-5.14.0-570.18.1.el9_6.x86_64

Steps

1. Install Linux on all physical or virtual grid hosts according to the distributor's instructions or your standard procedure.



If you are using the standard Linux installer, select the "compute node" software configuration, if available, or "minimal install" base environment. Don't install any graphical desktop environments.

2. Ensure that all hosts have access to package repositories, including the Extras channel.

You might need these additional packages later in this installation procedure.

3. If swap is enabled:

- a. Run the following command: `$ sudo swapoff --all`
- b. Remove all swap entries from `/etc/fstab` to persist the settings.



Failing to disable swap entirely can severely lower performance.

Configure the host network (Red Hat Enterprise Linux)

After completing the Linux installation on your hosts, you might need to perform some additional configuration to prepare a set of network interfaces on each host that are suitable for mapping into the StorageGRID nodes you will deploy later.

Before you begin

- You have reviewed the [StorageGRID networking guidelines](#).
- You have reviewed the information about [node container migration requirements](#).
- If you are using virtual hosts, you have read the [considerations and recommendations for MAC address cloning](#) before configuring the host network.



If you are using VMs as hosts, you should select VMXNET 3 as the virtual network adapter. The VMware E1000 network adapter has caused connectivity issues with StorageGRID containers deployed on certain distributions of Linux.

About this task

Grid nodes must be able to access the Grid Network and, optionally, the Admin and Client Networks. You provide this access by creating mappings that associate the host's physical interface to the virtual interfaces for each grid node. When creating host interfaces, use friendly names to facilitate deployment across all hosts, and to enable migration.

The same interface can be shared between the host and one or more nodes. For example, you might use the same interface for host access and node Admin Network access, to facilitate host and node maintenance. Although the same interface can be shared between the host and individual nodes, all must have different IP addresses. IP addresses can't be shared between nodes or between the host and any node.

You can use the same host network interface to provide the Grid Network interface for all StorageGRID nodes on the host; you can use a different host network interface for each node; or you can do something in between. However, you would not typically provide the same host network interface as both the Grid and Admin Network interfaces for a single node, or as the Grid Network interface for one node and the Client Network interface for another.

You can complete this task in many ways. For example, if your hosts are virtual machines and you are deploying one or two StorageGRID nodes for each host, you can create the correct number of network interfaces in the hypervisor, and use a 1-to-1 mapping. If you are deploying multiple nodes on bare metal hosts for production use, you can leverage the Linux networking stack's support for VLAN and LACP for fault tolerance and bandwidth sharing. The following sections provide detailed approaches for both of these examples. You don't need to use either of these examples; you can use any approach that meets your needs.



Don't use bond or bridge devices directly as the container network interface. Doing so could prevent node start-up caused by a kernel issue with the use of MACVLAN with bond and bridge devices in the container namespace. Instead, use a non-bond device, such as a VLAN or virtual Ethernet (veth) pair. Specify this device as the network interface in the node configuration file.

Related information

[Creating node configuration files](#)

Considerations and recommendations for MAC address cloning

MAC address cloning causes the container to use the MAC address of the host, and the host to use the MAC address of either an address you specify or a randomly generated one. You should use MAC address cloning to avoid the use of promiscuous mode network configurations.

Enabling MAC cloning

In certain environments, security can be enhanced through MAC address cloning because it enables you to use a dedicated virtual NIC for the Admin Network, Grid Network, and Client Network. Having the container use the MAC address of the dedicated NIC on the host allows you to avoid using promiscuous mode network configurations.



MAC address cloning is intended to be used with virtual server installations and might not function properly with all physical appliance configurations.



If a node fails to start due to a MAC cloning targeted interface being busy, you might need to set the link to "down" before starting node. Additionally, it is possible that the virtual environment might prevent MAC cloning on a network interface while the link is up. If a node fails to set the MAC address and start due to an interface being busy, setting the link to "down" before starting the node might fix the issue.

MAC address cloning is disabled by default and must be set by node configuration keys. You should enable it when you install StorageGRID.

There is one key for each network:

- `ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC`
- `GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC`
- `CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC`

Setting the key to "true" causes the container to use the MAC address of the host's NIC. Additionally, the host will then use the MAC address of the specified container network. By default, the container address is a randomly generated address, but if you have set one using the `_NETWORK_MAC` node configuration key, that address is used instead. The host and container will always have different MAC addresses.



Enabling MAC cloning on a virtual host without also enabling promiscuous mode on the hypervisor might cause Linux host networking using the host's interface to stop working.

MAC cloning use cases

There are two use cases to consider with MAC cloning:

- **MAC cloning not enabled:** When the `_CLONE_MAC` key in the node configuration file is not set, or set to "false," the host will use the host NIC MAC and the container will have a StorageGRID-generated MAC unless a MAC is specified in the `_NETWORK_MAC` key. If an address is set in the `_NETWORK_MAC` key, the container will have the address specified in the `_NETWORK_MAC` key. This configuration of keys requires the use of promiscuous mode.
- **MAC cloning enabled:** When the `_CLONE_MAC` key in the node configuration file is set to "true," the container uses the host NIC MAC, and the host uses a StorageGRID-generated MAC unless a MAC is specified in the `_NETWORK_MAC` key. If an address is set in the `_NETWORK_MAC` key, the host uses the specified address instead of a generated one. In this configuration of keys, you should not use promiscuous mode.



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

To enable MAC cloning, see the [instructions for creating node configuration files](#).

MAC cloning example


Example of MAC cloning enabled with a host having MAC address of 11:22:33:44:55:66 for the interface ens256 and the following keys in the node configuration file:

- ADMIN_NETWORK_TARGET = ens256
- ADMIN_NETWORK_MAC = b2:9c:02:c2:27:10
- ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC = true

Result: the host MAC for ens256 is b2:9c:02:c2:27:10 and the Admin Network MAC is 11:22:33:44:55:66

Example 1: 1-to-1 mapping to physical or virtual NICs

Example 1 describes a simple physical interface mapping that requires little or no host-side configuration.



The Linux operating system creates the `ensXYZ` interfaces automatically during installation or boot, or when the interfaces are hot-added. No configuration is required other than ensuring that the interfaces are set to come up automatically after boot. You do have to determine which `ensXYZ` corresponds to which StorageGRID network (Grid, Admin, or Client) so you can provide the correct mappings later in the configuration process.

Note that the figure shows multiple StorageGRID nodes; however, you would normally use this configuration for

single-node VMs.

If Switch 1 is a physical switch, you should configure the ports connected to interfaces 10G1 through 10G3 for access mode, and place them on the appropriate VLANs.

Example 2: LACP bond carrying VLANs


About this task

Example 2 assumes you are familiar with bonding network interfaces and with creating VLAN interfaces on the Linux distribution you are using.

Example 2 describes a generic, flexible, VLAN-based scheme that facilitates the sharing of all available network bandwidth across all nodes on a single host. This example is particularly applicable to bare metal hosts.

To understand this example, suppose you have three separate subnets for the Grid, Admin, and Client Networks at each data center. The subnets are on separate VLANs (1001, 1002, and 1003) and are presented to the host on a LACP-bonded trunk port (bond0). You would configure three VLAN interfaces on the bond: bond0.1001, bond0.1002, and bond0.1003.

If you require separate VLANs and subnets for node networks on the same host, you can add VLAN interfaces on the bond and map them into the host (shown as bond0.1004 in the illustration).



Steps

1. Aggregate all physical network interfaces that will be used for StorageGRID network connectivity into a

single LACP bond.

Use the same name for the bond on every host. For example, `bond0`.

2. Create VLAN interfaces that use this bond as their associated "physical device," using the standard VLAN interface naming convention `physdev-name.VLAN ID`.

Note that steps 1 and 2 require appropriate configuration on the edge switches terminating the other ends of the network links. The edge switch ports must also be aggregated into a LACP port channel, configured as a trunk, and allowed to pass all required VLANs.

Sample interface configuration files for this per-host networking configuration scheme are provided.

Related information

[Example /etc/sysconfig/network-scripts](#)

Configure host storage

You must allocate block storage volumes to each host.

Before you begin

You have reviewed the following topics, which provide information you need to accomplish this task:

- [Storage and performance requirements](#)
- [Node container migration requirements](#)

About this task

When allocating block storage volumes (LUNs) to hosts, use the tables in "Storage requirements" to determine the following:

- Number of volumes required for each host (based on the number and types of nodes that will be deployed on that host)
- Storage category for each volume (that is, System Data or Object Data)
- Size of each volume

You will use this information as well as the persistent name assigned by Linux to each physical volume when you deploy StorageGRID nodes on the host.



You don't need to partition, format, or mount any of these volumes; you just need to ensure they are visible to the hosts.



Only one object-data LUN is required for metadata-only Storage Nodes.

Avoid using "raw" special device files (`/dev/sdb`, for example) as you compose your list of volume names. These files can change across reboots of the host, which will impact proper operation of the system. If you are using iSCSI LUNs and Device Mapper Multipathing, consider using multipath aliases in the `/dev/mapper` directory, especially if your SAN topology includes redundant network paths to the shared storage. Alternatively, you can use the system-created softlinks under `/dev/disk/by-path/` for your persistent device names.

For example:

```
ls -l
$ ls -l /dev/disk/by-path/
total 0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:00:07.1-ata-2 -> ../../sr0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0 ->
../../sda
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part1
-> ../../sda1
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part2
-> ../../sda2
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:1:0 ->
../../sdb
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:2:0 ->
../../sdc
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:3:0 ->
../../sdd
```

Results will differ for each installation.

Assign friendly names to each of these block storage volumes to simplify the initial StorageGRID installation and future maintenance procedures. If you are using the device mapper multipath driver for redundant access to shared storage volumes, you can use the `alias` field in your `/etc/multipath.conf` file.

For example:

```

multipaths {
    multipath {
        wwid 3600a09800059d6df00005df2573c2c30
        alias docker-storage-volume-hostA
    }
    multipath {
        wwid 3600a09800059d6df00005df3573c2c30
        alias sgws-adml-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df4573c2c30
        alias sgws-adml-audit-logs
    }
    multipath {
        wwid 3600a09800059d6df00005df5573c2c30
        alias sgws-adml-tables
    }
    multipath {
        wwid 3600a09800059d6df00005df6573c2c30
        alias sgws-gw1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-rangedb-0
    }
    ...
}

```

Using the alias field in this way causes the aliases to appear as block devices in the `/dev/mapper` directory on the host, allowing you to specify a friendly, easily-validated name whenever a configuration or maintenance operation requires specifying a block storage volume.



If you are setting up shared storage to support StorageGRID node migration and using Device Mapper Multipathing, you can create and install a common `/etc/multipath.conf` on all co-located hosts. Just make sure to use a different container engine storage volume on each host. Using aliases and including the target hostname in the alias for each container engine storage volume LUN will make this easy to remember and is recommended.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

Related information

[Configure container engine storage volume](#)

Configure container engine storage volume

Before installing the container engine (Docker or Podman), you might need to format the storage volume and mount it.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

About this task

You can skip these steps if you plan to use local storage for the Docker or Podman storage volume and have sufficient space available on the host partition containing `/var/lib/docker` for Docker and `/var/lib/containers` for Podman.



Podman is supported only on Red Hat Enterprise Linux (RHEL).

Steps

1. Create a file system on the container engine storage volume:

```
sudo mkfs.ext4 container-engine-storage-volume-device
```

2. Mount the container engine storage volume:

- For Docker:

```
sudo mkdir -p /var/lib/docker
sudo mount container-storage-volume-device /var/lib/docker
```

- For Podman:

```
sudo mkdir -p /var/lib/containers
sudo mount container-storage-volume-device /var/lib/containers
```

3. Add an entry for container-storage-volume-device to `/etc/fstab`.

This step ensures that the storage volume will remount automatically after host reboots.

Install Docker

The StorageGRID system runs on Red Hat Enterprise Linux as a collection of containers. If you have chosen to use the Docker container engine, follow these steps to install Docker. Otherwise, [install Podman](#).

Steps

1. Install Docker by following the instructions for your Linux distribution.



If Docker is not included with your Linux distribution, you can download it from the Docker website.

2. Ensure Docker has been enabled and started by running the following two commands:

```
sudo systemctl enable docker
```

```
sudo systemctl start docker
```

3. Confirm you have installed the expected version of Docker by entering the following:

```
sudo docker version
```

The Client and Server versions must be 1.11.0 or later.

Install Podman

The StorageGRID system runs on Red Hat Enterprise Linux as a collection of containers. If you have chosen to use the Podman container engine, follow these steps to install Podman. Otherwise, [install Docker](#).



Podman is supported only on Red Hat Enterprise Linux (RHEL).

Steps

1. Install Podman and Podman-Docker by following the instructions for your Linux distribution.



You must also install the Podman-Docker package when you install Podman.

2. Confirm you have installed the expected version of Podman and Podman-Docker by entering the following:

```
sudo docker version
```



The Podman-Docker package allows you to use Docker commands.

The Client and Server versions must be 3.2.3 or later.

```
Version: 3.2.3
API Version: 3.2.3
Go Version: go1.15.7
Built: Tue Jul 27 03:29:39 2021
OS/Arch: linux/amd64
```

Install StorageGRID host services

You use the StorageGRID RPM package to install the StorageGRID host services.

About this task

These instructions describe how to install the host services from the RPM packages. As an alternative, you can use the DNF repository metadata included in the installation archive to install the RPM packages remotely. See the DNF repository instructions for your Linux operating system.

Steps

1. Copy the StorageGRID RPM packages to each of your hosts, or make them available on shared storage.

For example, place them in the `/tmp` directory, so you can use the example command in the next step.

2. Log in to each host as root or using an account with sudo permission, and run the following commands in the order specified:

```
sudo dnf --nogpgcheck localinstall /tmp/StorageGRID-Webscale-Images-  
version-SHA.rpm
```

```
sudo dnf --nogpgcheck localinstall /tmp/StorageGRID-Webscale-Service-  
version-SHA.rpm
```



You must install the Images package first, and the Service package second.



If you placed the packages in a directory other than `/tmp`, modify the command to reflect the path you used.

Automate StorageGRID installation on Red Hat Enterprise Linux

You can automate the installation of the StorageGRID host service and the configuration of grid nodes.

Automating the deployment might be useful in any of the following cases:

- You already use a standard orchestration framework, such as Ansible, Puppet, or Chef, to deploy and configure physical or virtual hosts.
- You intend to deploy multiple StorageGRID instances.
- You are deploying a large, complex StorageGRID instance.

The StorageGRID host service is installed by a package and driven by configuration files. You can create the configuration files using one of these methods:

- [Create the configuration files](#) interactively during a manual installation.
- Prepare the configuration files ahead of time (or programmatically) to enable automated installation using standard orchestration frameworks, as describe in this article.

StorageGRID provides optional Python scripts for automating the configuration of StorageGRID appliances and the entire StorageGRID system (the "grid"). You can use these scripts directly, or you can inspect them to learn how to use the [StorageGRID installation REST API](#) in grid deployment and configuration tools you develop yourself.

Automate the installation and configuration of the StorageGRID host service

You can automate the installation of the StorageGRID host service using standard orchestration frameworks such as Ansible, Puppet, Chef, Fabric, or SaltStack.

The StorageGRID host service is packaged in an RPM and is driven by configuration files that you can prepare ahead of time (or programmatically) to enable automated installation. If you already use a standard orchestration framework to install and configure RHEL, adding StorageGRID to your playbooks or recipes should be straightforward.

See the example Ansible role and playbook in the `/extras` folder supplied with the installation archive. The Ansible playbook shows how the `storagegrid` role prepares the host and installs StorageGRID onto the target servers. You can customize the role or playbook as necessary.



The example playbook does not include the steps required to create network devices before starting the StorageGRID host service. Add these steps before finalizing and using the playbook.

You can automate all of the steps for preparing the hosts and deploying virtual grid nodes.

Example Ansible role and playbook

Example Ansible role and playbook are supplied with the installation archive in the `/extras` folder. The Ansible playbook shows how the `storagegrid` role prepares the hosts and installs StorageGRID onto the target servers. You can customize the role or playbook as necessary.

The installation tasks in the provided `storagegrid` role example use the `ansible.builtin.dnf` module to perform the installation from the local RPM files or a remote Yum repository. If the module is unavailable or not supported, you might need to edit the appropriate Ansible tasks in the following files to use the `yum` or `ansible.builtin.yum` module:

- `roles/storagegrid/tasks/rhel_install_from_repo.yml`
- `roles/storagegrid/tasks/rhel_install_from_local.yml`

Automate the configuration of StorageGRID

After deploying the grid nodes, you can automate the configuration of the StorageGRID system.

Before you begin

- You know the location of the following files from the installation archive.

Filename	Description
<code>configure-storagegrid.py</code>	Python script used to automate the configuration
<code>configure-storagegrid.sample.json</code>	Example configuration file for use with the script
<code>configure-storagegrid.blank.json</code>	Blank configuration file for use with the script

- You have created a `configure-storagegrid.json` configuration file. To create this file, you can modify the example configuration file (`configure-storagegrid.sample.json`) or the blank configuration file

```
(configure-storagegrid.blank.json).
```

About this task

You can use the `configure-storagegrid.py` Python script and the `configure-storagegrid.json` configuration file to automate the configuration of your StorageGRID system.



You can also configure the system using the Grid Manager or the Installation API.

Steps

1. Log in to the Linux machine you are using to run the Python script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/platform
```

where `platform` is `debs`, `rpms`, or `vsphere`.

3. Run the Python script and use the configuration file you created.

For example:

```
./configure-storagegrid.py ./configure-storagegrid.json --start-install
```

Result

A Recovery Package `.zip` file is generated during the configuration process, and it is downloaded to the directory where you are running the installation and configuration process. You must back up the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fails. For example, copy it to a secure, backed up network location and to a secure cloud storage location.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

If you specified that random passwords be generated, open the `Passwords.txt` file and look for the passwords required to access your StorageGRID system.

```
#####
##### The StorageGRID "Recovery Package" has been downloaded as: #####
#####      ./sgws-recovery-package-994078-rev1.zip      #####
##### Safeguard this file as it will be needed in case of a #####
#####      StorageGRID node recovery.      #####
#####
```

Your StorageGRID system is installed and configured when a confirmation message is displayed.

StorageGRID has been configured and installed.

Related information

[Installation REST API](#)

Deploy virtual grid nodes (Red Hat)

Create node configuration files for Red Hat Enterprise Linux deployments

Node configuration files are small text files that provide the information the StorageGRID host service needs to start a node and connect it to the appropriate network and block storage resources. Node configuration files are used for virtual nodes and aren't used for appliance nodes.

Location for node configuration files

Place the configuration file for each StorageGRID node in the `/etc/storagegrid/nodes` directory on the host where the node will run. For example, if you plan to run one Admin Node, one Gateway Node, and one Storage Node on HostA, you must place three node configuration files in `/etc/storagegrid/nodes` on HostA.

You can create the configuration files directly on each host using a text editor, such as vim or nano, or you can create them elsewhere and move them to each host.

Naming of node configuration files

The names of the configuration files are significant. The format is `node-name.conf`, where `node-name` is a name you assign to the node. This name appears in the StorageGRID Installer and is used for node maintenance operations, such as node migration.

Node names must follow these rules:

- Must be unique
- Must start with a letter
- Can contain the characters A through Z and a through z
- Can contain the numbers 0 through 9
- Can contain one or more hyphens (-)
- Must be no more than 32 characters, not including the `.conf` extension

Any files in `/etc/storagegrid/nodes` that don't follow these naming conventions will not be parsed by the host service.

If you have a multi-site topology planned for your grid, a typical node naming scheme might be:

```
site-nodetype-nodenum.conf
```

For example, you might use `dc1-adm1.conf` for the first Admin Node in Data Center 1, and `dc2-sn3.conf` for the third Storage Node in Data Center 2. However, you can use any scheme you like, as long as all node names follow the naming rules.

Contents of a node configuration file

A configuration file contains key/value pairs, with one key and one value per line. For each key/value pair, follow these rules:

- The key and the value must be separated by an equal sign (=) and optional whitespace.
- The keys can contain no spaces.
- The values can contain embedded spaces.
- Any leading or trailing whitespace is ignored.

The following table defines the values for all supported keys. Each key has one of the following designations:

- **Required:** Required for every node or for the specified node types
- **Best practice:** Optional, although recommended
- **Optional:** Optional for all nodes

Admin Network keys

ADMIN_IP

Value	Designation
Grid Network IPv4 address of the primary Admin Node for the grid to which this node belongs. Use the same value you specified for GRID_NETWORK_IP for the grid node with NODE_TYPE = VM_Admin_Node and ADMIN_ROLE = Primary. If you omit this parameter, the node attempts to discover a primary Admin Node using mDNS. How grid nodes discover the primary Admin Node Note: This value is ignored, and might be prohibited, on the primary Admin Node.	Best practice

ADMIN_NETWORK_CONFIG

Value	Designation
DHCP, STATIC, or DISABLED	Optional

ADMIN_NETWORK_ESL

Value	Designation
Comma-separated list of subnets in CIDR notation to which this node should communicate using the Admin Network gateway. Example: 172.16.0.0/21, 172.17.0.0/21	Optional

ADMIN_NETWORK_GATEWAY

Value	Designation
IPv4 address of the local Admin Network gateway for this node. Must be on the subnet defined by ADMIN_NETWORK_IP and ADMIN_NETWORK_MASK. This value is ignored for DHCP-configured networks. Examples: 1.1.1.1 10.224.4.81	Required if ADMIN_NETWORK_ESL is specified. Optional otherwise.

ADMIN_NETWORK_IP

Value	Designation
IPv4 address of this node on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 1.1.1.1 10.224.4.81	Required when ADMIN_NETWORK_CONFIG = STATIC. Optional otherwise.

ADMIN_NETWORK_MAC

Value	Designation
The MAC address for the Admin Network interface in the container. This field is optional. If omitted, a MAC address will be generated automatically. Must be 6 pairs of hexadecimal digits separated by colons. Example: b2:9c:02:c2:27:10	Optional

ADMIN_NETWORK_MASK

Value	Designation
IPv4 netmask for this node, on the Admin Network. Specify this key when ADMIN_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 255.255.255.0 255.255.248.0	Required if ADMIN_NETWORK_IP is specified and ADMIN_NETWORK_CONFIG = STATIC. Optional otherwise.

ADMIN_NETWORK_MTU

Value	Designation
The maximum transmission unit (MTU) for this node on the Admin Network. Don't specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur. Examples: 1500 8192	Optional

ADMIN_NETWORK_TARGET

Value	Designation
Name of the host device that you will use for Admin Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or CLIENT_NETWORK_TARGET. Note: Don't use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair. Best practice: Specify a value even if this node will not initially have an Admin Network IP address. Then you can add an Admin Network IP address later, without having to reconfigure the node on the host. Examples: <pre>bond0.1002</pre> <pre>ens256</pre>	Best practice

ADMIN_NETWORK_TARGET_TYPE

Value	Designation
Interface (This is the only supported value.)	Optional

ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Value	Designation
True or False Set the key to "true" to cause the StorageGRID container use the MAC address of the host host target interface on the Admin Network. Best practice: In networks where promiscuous mode would be required, use the ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead. For more details on MAC cloning: • Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux) • Considerations and recommendations for MAC address cloning (Ubuntu or Debian)	Best practice

ADMIN_ROLE

Value	Designation
Primary or non-primary This key is only required when NODE_TYPE = VM_Admin_Node; don't specify it for other node types.	Required when NODE_TYPE = VM_Admin_Node Optional otherwise.

Block device keys

BLOCK_DEVICE_AUDIT_LOGS

Value	Designation
Path and name of the block device special file this node will use for persistent storage of audit logs. Examples: <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-adml-audit-logs</pre>	Required for nodes with NODE_TYPE = VM_Admin_Node. Don't specify it for other node types.

BLOCK_DEVICE_RANGEDB_nnn

Value	Designation
Path and name of the block device special file this node will use for persistent object storage. This key is only required for nodes with <code>NODE_TYPE = VM_Storage_Node</code>; don't specify it for other node types. Only <code>BLOCK_DEVICE_RANGEDB_000</code> is required; the rest are optional. The block device specified for <code>BLOCK_DEVICE_RANGEDB_000</code> must be at least 4 TB; the others can be smaller. Don't leave gaps. If you specify <code>BLOCK_DEVICE_RANGEDB_005</code>, you must also specify <code>BLOCK_DEVICE_RANGEDB_004</code>. Note: For compatibility with existing deployments, two-digit keys are supported for upgraded nodes. Examples: <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-sn1-rangedb-000</pre>	Required: <code>BLOCK_DEVICE_RANGEDB_000</code> Optional: <code>BLOCK_DEVICE_RANGEDB_001</code> <code>BLOCK_DEVICE_RANGEDB_002</code> <code>BLOCK_DEVICE_RANGEDB_003</code> <code>BLOCK_DEVICE_RANGEDB_004</code> <code>BLOCK_DEVICE_RANGEDB_005</code> <code>BLOCK_DEVICE_RANGEDB_006</code> <code>BLOCK_DEVICE_RANGEDB_007</code> <code>BLOCK_DEVICE_RANGEDB_008</code> <code>BLOCK_DEVICE_RANGEDB_009</code> <code>BLOCK_DEVICE_RANGEDB_010</code> <code>BLOCK_DEVICE_RANGEDB_011</code> <code>BLOCK_DEVICE_RANGEDB_012</code> <code>BLOCK_DEVICE_RANGEDB_013</code> <code>BLOCK_DEVICE_RANGEDB_014</code> <code>BLOCK_DEVICE_RANGEDB_015</code>

BLOCK_DEVICE_TABLES

Value	Designation
Path and name of the block device special file this node will use for persistent storage of database tables. This key is only required for nodes with NODE_TYPE = VM_Admin_Node; don't specify it for other node types. Examples: <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-adml-tables</pre>	Required

BLOCK_DEVICE_VAR_LOCAL

Value	Designation
Path and name of the block device special file this node will use for its /var/local persistent storage. Examples: <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-sn1-var-local</pre>	Required

Client Network keys

CLIENT_NETWORK_CONFIG

Value	Designation
DHCP, STATIC, or DISABLED	Optional

CLIENT_NETWORK_GATEWAY

Value	Designation

IPv4 address of the local Client Network gateway for this node, which must be on the subnet defined by <code>CLIENT_NETWORK_IP</code> and <code>CLIENT_NETWORK_MASK</code>. This value is ignored for DHCP-configured networks. Examples: 1.1.1.1 10.224.4.81	Optional

CLIENT_NETWORK_IP

Value	Designation
IPv4 address of this node on the Client Network. This key is only required when <code>CLIENT_NETWORK_CONFIG = STATIC</code>; don't specify it for other values. Examples: 1.1.1.1 10.224.4.81	Required when <code>CLIENT_NETWORK_CONFIG = STATIC</code> Optional otherwise.

CLIENT_NETWORK_MAC

Value	Designation
The MAC address for the Client Network interface in the container. This field is optional. If omitted, a MAC address will be generated automatically. Must be 6 pairs of hexadecimal digits separated by colons. Example: <code>b2:9c:02:c2:27:20</code>	Optional

CLIENT_NETWORK_MASK

Value	Designation
IPv4 netmask for this node on the Client Network. Specify this key when CLIENT_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 255.255.255.0 255.255.248.0	Required if CLIENT_NETWORK_IP is specified and CLIENT_NETWORK_CONFIG = STATIC Optional otherwise.

CLIENT_NETWORK_MTU

Value	Designation
The maximum transmission unit (MTU) for this node on the Client Network. Don't specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur. Examples: 1500 8192	Optional

CLIENT_NETWORK_TARGET

Value	Designation
Name of the host device that you will use for Client Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or ADMIN_NETWORK_TARGET. Note: Don't use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair. Best practice: Specify a value even if this node will not initially have a Client Network IP address. Then you can add a Client Network IP address later, without having to reconfigure the node on the host. Examples: <pre>bond0.1003</pre> <pre>ens423</pre>	Best practice

CLIENT_NETWORK_TARGET_TYPE

Value	Designation
Interface (This is only supported value.)	Optional

CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Value	Designation
True or False Set the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Client Network. Best practice: In networks where promiscuous mode would be required, use the CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead. For more details on MAC cloning: • Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux) • Considerations and recommendations for MAC address cloning (Ubuntu or Debian)	Best practice

Grid Network keys

GRID_NETWORK_CONFIG

Value	Designation
STATIC or DHCP Defaults to STATIC if not specified.	Best practice

GRID_NETWORK_GATEWAY

Value	Designation
IPv4 address of the local Grid Network gateway for this node, which must be on the subnet defined by GRID_NETWORK_IP and GRID_NETWORK_MASK. This value is ignored for DHCP-configured networks. If the Grid Network is a single subnet with no gateway, use either the standard gateway address for the subnet (X.Y.Z.1) or this node's GRID_NETWORK_IP value; either value will simplify potential future Grid Network expansions.	Required

GRID_NETWORK_IP

Value	Designation
IPv4 address of this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 1.1.1.1 10.224.4.81	Required when GRID_NETWORK_CONFIG = STATIC Optional otherwise.

GRID_NETWORK_MAC

Value	Designation
The MAC address for the Grid Network interface in the container. Must be 6 pairs of hexadecimal digits separated by colons. Example: b2:9c:02:c2:27:30	Optional If omitted, a MAC address will be generated automatically.

GRID_NETWORK_MASK

Value	Designation
IPv4 netmask for this node on the Grid Network. Specify this key when GRID_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 255.255.255.0 255.255.248.0	Required when GRID_NETWORK_IP is specified and GRID_NETWORK_CONFIG = STATIC. Optional otherwise.

GRID_NETWORK_MTU

Value	Designation
The maximum transmission unit (MTU) for this node on the Grid Network. Don't specify if GRID_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur. IMPORTANT: For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The Grid Network MTU mismatch alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values don't have to be the same for all network types. Examples: 1500 8192	Optional

GRID_NETWORK_TARGET

Value	Designation
Name of the host device that you will use for Grid Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for ADMIN_NETWORK_TARGET or CLIENT_NETWORK_TARGET. Note: Don't use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair. Examples: <pre>bond0.1001</pre> <pre>ens192</pre>	Required

GRID_NETWORK_TARGET_TYPE

Value	Designation
Interface (This is the only supported value.)	Optional

GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Value	Designation
True or False Set the value of the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Grid Network. Best practice: In networks where promiscuous mode would be required, use the GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead. For more details on MAC cloning: • Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux) • Considerations and recommendations for MAC address cloning (Ubuntu or Debian)	Best practice

Installation password key (temporary)

CUSTOM_TEMPORARY_PASSWORD_HASH

Value	Designation
For the primary Admin Node, set a default temporary password for the StorageGRID Installation API during installation. Note: Set an installation password on the primary Admin Node only. If you attempt to set a password on another node type, validation of the node configuration file will fail. Setting this value has no effect when installation has completed. If this key is omitted, by default no temporary password is set. Alternatively, you can set a temporary password using the StorageGRID Installation API. Must be a <code>crypt()</code> SHA-512 password hash with format <code>\$6\$<salt>\$<password hash></code> for a password of at least 8 and no more than 32 characters. This hash can be generated using CLI tools, such as the <code>openssl passwd</code> command in SHA-512 mode.	Best practice

Interfaces key

INTERFACE_TARGET_nnnn

Value	Designation
Name and optional description for an extra interface you want to add to this node. You can add multiple extra interfaces to each node. For <i>nnnn</i>, specify a unique number for each INTERFACE_TARGET entry you are adding. For the value, specify the name of the physical interface on the bare-metal host. Then, optionally, add a comma and provide a description of the interface, which is displayed on the VLAN interfaces page and the HA groups page. Example: <code>INTERFACE_TARGET_0001=ens256, Trunk</code> If you add a trunk interface, you must configure a VLAN interface in StorageGRID. If you add an access interface, you can add the interface directly to an HA group; you don't need to configure a VLAN interface.	Optional

Maximum RAM key

MAXIMUM_RAM

Value	Designation
The maximum amount of RAM that this node is allowed to consume. If this key is omitted, the node has no memory restrictions. When setting this field for a production-level node, specify a value that is at least 24 GB and 16 to 32 GB less than the total system RAM. Note: The RAM value affects a node's actual metadata reserved space. See the description of what Metadata Reserved Space is. The format for this field is <i>numberunit</i>, where <i>unit</i> can be b, k, m, or g. Examples: 24g 38654705664b Note: If you want to use this option, you must enable kernel support for memory cgroups.	Optional

Node type keys

NODE_TYPE

Value	Designation
Type of node: • VM_Admin_Node • VM_Storage_Node • VM_Archive_Node • VM_API_Gateway	Required

STORAGE_TYPE

Value	Designation
Defines the type of objects a Storage Node contains. For more information, see Types of Storage Nodes. This key is only required for nodes with <code>NODE_TYPE = VM_Storage_Node</code>; don't specify it for other node types. Storage types: • combined • data • metadata Note: If the <code>STORAGE_TYPE</code> is not specified, the Storage Node type is set to combined (data and metadata) by default.	Optional

Port remap keys

PORT_REMAP

Value	Designation
Remaps any port used by a node for internal grid node communications or external communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID, as described in Internal grid node communications or External communications. IMPORTANT: Don't remap the ports you are planning to use to configure load balancer endpoints. Note: If only <code>PORT_REMAP</code> is set, the mapping that you specify is used for both inbound and outbound communications. If <code>PORT_REMAP_INBOUND</code> is also specified, <code>PORT_REMAP</code> applies only to outbound communications. The format used is: <i>network type/protocol/default port used by grid node/new port</i>, where <i>network type</i> is grid, admin, or client, and <i>protocol</i> is tcp or udp. Example: <code>PORT_REMAP = client/tcp/18082/443</code> You can also remap multiple ports using a comma-separated list. Example: <code>PORT_REMAP = client/tcp/18082/443, client/tcp/18083/80</code>	Optional

PORT_REMAP_INBOUND

Value	Designation
Remaps inbound communications to the specified port. If you specify <code>PORT_REMAP_INBOUND</code> but don't specify a value for <code>PORT_REMAP</code>, outbound communications for the port are unchanged. IMPORTANT: Don't remap the ports you are planning to use to configure load balancer endpoints. The format used is: <i>network type/protocol/remapped port/default port used by grid node</i>, where <i>network type</i> is <code>grid</code>, <code>admin</code>, or <code>client</code>, and <i>protocol</i> is <code>tcp</code> or <code>udp</code>. Example: <code>PORT_REMAP_INBOUND = grid/tcp/3022/22</code> You can also remap multiple inbound ports using a comma-separated list. Example: <code>PORT_REMAP_INBOUND = grid/tcp/3022/22, admin/tcp/3022/22</code>	Optional

How grid nodes discover the primary Admin Node

Grid nodes communicate with the primary Admin Node for configuration and management. Each grid node must know the IP address of the primary Admin Node on the Grid Network.

To ensure that a grid node can access the primary Admin Node, you can do either of the following when deploying the node:

- You can use the `ADMIN_IP` parameter to enter the primary Admin Node's IP address manually.
- You can omit the `ADMIN_IP` parameter to have the grid node discover the value automatically. Automatic discovery is especially useful when the Grid Network uses DHCP to assign the IP address to the primary Admin Node.

Automatic discovery of the primary Admin Node is accomplished using a multicast domain name system (mDNS). When the primary Admin Node first starts up, it publishes its IP address using mDNS. Other nodes on the same subnet can then query for the IP address and acquire it automatically. However, because multicast IP traffic is not normally routable across subnets, nodes on other subnets can't acquire the primary Admin Node's IP address directly.

If you use automatic discovery:



- You must include the `ADMIN_IP` setting for at least one grid node on any subnets that the primary Admin Node is not directly attached to. This grid node will then publish the primary Admin Node's IP address for other nodes on the subnet to discover with mDNS.
- Ensure that your network infrastructure supports passing multi-cast IP traffic within a subnet.

Example node configuration files

You can use the example node configuration files to help set up the node configuration

files for your StorageGRID system. The examples show node configuration files for all types of grid nodes.

For most nodes, you can add Admin and Client Network addressing information (IP, mask, gateway, and so on) when you configure the grid using the Grid Manager or the Installation API. The exception is the primary Admin Node. If you want to browse to the Admin Network IP of the primary Admin Node to complete grid configuration (because the Grid Network is not routed, for example), you must configure the Admin Network connection for the primary Admin Node in its node configuration file. This is shown in the example.



In the examples, the Client Network target has been configured as a best practice, even though the Client Network is disabled by default.

Example for primary Admin Node

Example file name: /etc/storagegrid/nodes/dcl-adm1.conf

Example file contents:

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
TEMPORARY_PASSWORD_TYPE = Use custom password
CUSTOM_TEMPORARY_PASSWORD = Passw0rd
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-adm1-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dcl-adm1-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dcl-adm1-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.2
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_IP = 192.168.100.2
ADMIN_NETWORK_MASK = 255.255.248.0
ADMIN_NETWORK_GATEWAY = 192.168.100.1
ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0.0/21,172.17.0.0/21
```

Example for Storage Node

Example file name: /etc/storagegrid/nodes/dcl-sn1.conf

Example file contents:


```
NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/dcl-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/dcl-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/dcl-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/dcl-sn1-rangedb-3
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Example for Gateway Node

Example file name: /etc/storagegrid/nodes/dcl-gw1.conf

Example file contents:

```
NODE_TYPE = VM_API_Gateway
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-gw1-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.5
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Example for a non-primary Admin Node

Example file name: /etc/storagegrid/nodes/dcl-adm2.conf

Example file contents:

```

NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Non-Primary
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adm2-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adm2-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dc1-adm2-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.6
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

```

Validate the StorageGRID configuration

After creating configuration files in `/etc/storagegrid/nodes` for each of your StorageGRID nodes, you must validate the contents of those files.

To validate the contents of the configuration files, run the following command on each host:

```
sudo storagegrid node validate all
```

If the files are correct, the output shows **PASSED** for each configuration file, as shown in the example.



When using only one LUN on metadata-only nodes, you might receive a warning message that can be ignored.

```

Checking for misnamed node configuration files... PASSED
Checking configuration file for node dc1-adm1... PASSED
Checking configuration file for node dc1-gw1... PASSED
Checking configuration file for node dc1-sn1... PASSED
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-sn3... PASSED
Checking for duplication of unique values between nodes... PASSED

```



For an automated installation, you can suppress this output by using the `-q` or `--quiet` options in the `storagegrid` command (for example, `storagegrid --quiet...`). If you suppress the output, the command will have a non-zero exit value if any configuration warnings or errors were detected.

If the configuration files are incorrect, the issues are shown as **WARNING** and **ERROR**, as shown in the example. If any configuration errors are found, you must correct them before you continue with the installation.

```

Checking for misnamed node configuration files...
WARNING: ignoring /etc/storagegrid/nodes/dcl-adml
WARNING: ignoring /etc/storagegrid/nodes/dcl-sn2.conf.keep
WARNING: ignoring /etc/storagegrid/nodes/my-file.txt
Checking configuration file for node dcl-adml...
ERROR: NODE_TYPE = VM_Foo_Node
      VM_Foo_Node is not a valid node type.  See *.conf.sample
ERROR: ADMIN_ROLE = Foo
      Foo is not a valid admin role.  See *.conf.sample
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-gw1-var-local
      /dev/mapper/sgws-gw1-var-local is not a valid block device
Checking configuration file for node dcl-gw1...
ERROR: GRID_NETWORK_TARGET = bond0.1001
      bond0.1001 is not a valid interface.  See `ip link show`
ERROR: GRID_NETWORK_IP = 10.1.3
      10.1.3 is not a valid IPv4 address
ERROR: GRID_NETWORK_MASK = 255.248.255.0
      255.248.255.0 is not a valid IPv4 subnet mask
Checking configuration file for node dcl-sn1...
ERROR: GRID_NETWORK_GATEWAY = 10.2.0.1
      10.2.0.1 is not on the local subnet
ERROR: ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0foo
      Could not parse subnet list
Checking configuration file for node dcl-sn2... PASSED
Checking configuration file for node dcl-sn3... PASSED
Checking for duplication of unique values between nodes...
ERROR: GRID_NETWORK_IP = 10.1.0.4
      dcl-sn2 and dcl-sn3 have the same GRID_NETWORK_IP
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn2-var-local
      dcl-sn2 and dcl-sn3 have the same BLOCK_DEVICE_VAR_LOCAL
ERROR: BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn2-rangedb-0
      dcl-sn2 and dcl-sn3 have the same BLOCK_DEVICE_RANGEDB_00

```

Start the StorageGRID host service

To start your StorageGRID nodes, and ensure they restart after a host reboot, you must enable and start the StorageGRID host service.

Steps

1. Run the following commands on each host:

```

sudo systemctl enable storagegrid
sudo systemctl start storagegrid

```

2. Run the following command to ensure the deployment is proceeding:

```
sudo storagegrid node status node-name
```

3. If any node returns a status of "Not Running" or "Stopped," run the following command:

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

4. If you have previously enabled and started the StorageGRID host service (or if you are unsure if the service has been enabled and started), also run the following command:

```
sudo systemctl reload-or-restart storagegrid
```

Configure the grid and complete installation (Red Hat)

Navigate to the Grid Manager

You use the Grid Manager to define all of the information required to configure your StorageGRID system.

Before you begin

The primary Admin Node must be deployed and have completed the initial startup sequence.

Steps

1. Open your web browser and navigate to:

```
https://primary_admin_node_ip
```

Alternatively, you can access the Grid Manager on port 8443:


```
https://primary_admin_node_ip:8443
```

You can use the IP address for the primary Admin Node IP on the Grid Network or on the Admin Network, as appropriate for your network configuration.

2. Manage a temporary installer password as needed:
 - If a password has already been set using one of these methods, enter the password to proceed.
 - A user set the password while accessing the installer previously
 - The password was automatically imported from the node config file at `/etc/storagegrid/nodes/<node_name>.conf`
 - If a password has not been set, optionally set a password to secure the StorageGRID installer.
3. Select **Install a StorageGRID system**.

The page used to configure a StorageGRID system appears.

Install



License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

Browse

Specify the StorageGRID license information

You must specify the name for your StorageGRID system and upload the license file provided by NetApp.

Steps

1. On the License page, enter a meaningful name for your StorageGRID system in the **Grid Name** field.

After installation, the name is displayed at the top of the Nodes menu.

2. Select **Browse**, locate the NetApp license file (*NLF-unique-id.txt*), and select **Open**.

The license file is validated, and the serial number is displayed.



The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product. You can update to a license that offers support after installation.

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File NLF-959007-Internal.txt

License Serial Number

3. Select **Next**.

Add sites

You must create at least one site when you are installing StorageGRID. You can create additional sites to increase the reliability and storage capacity of your StorageGRID system.

Steps

1. On the Sites page, enter the **Site Name**.
2. To add additional sites, click the plus sign next to the last site entry and enter the name in the new **Site Name** text box.

Add as many additional sites as required for your grid topology. You can add up to 16 sites.

NetApp® StorageGRID®

Help ▾

Install

1

License

2

Sites

3

Grid Network

4

Grid Nodes

5

NTP

6

DNS

7

Passwords

8

Summary

Sites

In a single-site deployment, infrastructure and operations are centralized in one site.

In a multi-site deployment, infrastructure can be distributed asymmetrically across sites, and proportional to the needs of each site. Typically, sites are located in geographically different locations. Having multiple sites also allows the use of distributed replication and erasure coding for increased availability and resiliency.

Site Name 1

Raleigh

✕

Site Name 2

Atlanta

+ ✕

3. Click **Next**.

Specify Grid Network subnets

You must specify the subnets that are used on the Grid Network.

About this task

The subnet entries include the subnets for the Grid Network for each site in your StorageGRID system, along with any subnets that need to be reachable through the Grid Network.

If you have multiple grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway.

Steps

1. Specify the CIDR network address for at least one Grid Network in the **Subnet 1** text box.
2. Click the plus sign next to the last entry to add an additional network entry. You must specify all subnets for all sites in the Grid Network.
 - If you have already deployed at least one node, click **Discover Grid Networks Subnets** to automatically populate the Grid Network Subnet List with the subnets reported by grid nodes that have registered with the Grid Manager.

- You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

NetApp® StorageGRID® Help ▾

Install

1 License 2 Sites 3 **Grid Network** 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Grid Network

You must specify the subnets that are used on the Grid Network. These entries typically include the subnets for the Grid Network for each site in your StorageGRID system. Select Discover Grid Networks to automatically add subnets based on the network configuration of all registered nodes.

Note: You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

Subnet 1 +

3. Click **Next**.

Approve pending grid nodes

You must approve each grid node before it can join the StorageGRID system.

Before you begin

You have deployed all virtual and StorageGRID appliance grid nodes.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

Steps

1. Review the Pending Nodes list, and confirm that it shows all of the grid nodes you deployed.



If a grid node is missing, confirm that it was deployed successfully and has the correct Grid Network IP of the primary admin node set for ADMIN_IP.

2. Select the radio button next to a pending node you want to approve.



Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

+ Approve		✖ Remove		Search		
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address	
☒	50:6b:4b:42:d7:00	NetApp-SGA	Storage Node	StorageGRID Appliance	172.16.5.20/21	
						◀ ▶

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

 Edit

 Reset

 Remove

Search



	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
☐	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
☐	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
☐	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
☐	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
☐	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21





3. Click **Approve**.

4. In General Settings, modify settings for the following properties, as necessary:

- **Site:** The system name of the site for this grid node.
- **Name:** The system name for the node. The name defaults to the name you specified when you configured the node.

System names are required for internal StorageGRID operations and can't be changed after you complete the installation. However, during this step of the installation process, you can change system names as required.

- **NTP Role:** The Network Time Protocol (NTP) role of the grid node. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the Primary role to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

- **Storage Type** (Storage Nodes only): Specify that a new Storage Node be used exclusively for data only, metadata only, or both. The options are **Data and metadata** ("combined"), **Data only**, and **Metadata only**.



See [Types of Storage Nodes](#) for information about requirements for these node types.

- **ADC service** (Storage Nodes only): Select **Automatic** to let the system determine whether the node requires the Administrative Domain Controller (ADC) service. The ADC service keeps track of the location and availability of grid services. At least three Storage Nodes at each site must include the ADC service. You can't add the ADC service to a node after it is deployed.

5. In Grid Network, modify settings for the following properties as necessary:

- **IPv4 Address (CIDR)**: The CIDR network address for the Grid Network interface (eth0 inside the container). For example: 192.168.1.234/21
- **Gateway**: The Grid Network gateway. For example: 192.168.0.1

The gateway is required if there are multiple grid subnets.



If you selected DHCP for the Grid Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

6. If you want to configure the Admin Network for the grid node, add or update the settings in the Admin Network section as necessary.

Enter the destination subnets of the routes out of this interface in the **Subnets (CIDR)** text box. If there are multiple Admin subnets, the Admin gateway is required.



If you selected DHCP for the Admin Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Admin Network was not configured during the initial installation using the StorageGRID Appliance Installer, it can't be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

Rebooting can take several minutes.

- Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- Select **Configure Networking > IP Configuration** and configure the enabled networks.
- Return to the Home page and click **Start Installation**.
- In the Grid Manager: If the node is listed in the Approved Nodes table, remove the node.

- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page of the Appliance Installer.

For additional information, see the installation instructions for your appliance model.

7. If you want to configure the Client Network for the grid node, add or update the settings in the Client Network section as necessary. If the Client Network is configured, the gateway is required, and it becomes the default gateway for the node after installation.



If you selected DHCP for the Client Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Client Network was not configured during the initial installation using the StorageGRID Appliance Installer, it can't be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

Rebooting can take several minutes.

- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, remove the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page of the Appliance Installer.

For additional information, see the installation instructions for your appliance.

8. Click **Save**.

The grid node entry moves to the Approved Nodes list.



Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

+

Approve

x

Remove

Search

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
No results found.				

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

Edit

Reset

x

Remove

Search

	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
☐	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
☐	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
☐	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
☐	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
☐	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21
☐	50:6b:4b:42:d7:00	NetApp-SGA	Raleigh	Storage Node	StorageGRID Appliance	172.16.5.20/21

9. Repeat these steps for each pending grid node you want to approve.

You must approve all nodes that you want in the grid. However, you can return to this page at any time before you click **Install** on the Summary page. You can modify the properties of an approved grid node by selecting its radio button and clicking **Edit**.

10. When you are done approving grid nodes, click **Next**.

Specify Network Time Protocol server information

You must specify the Network Time Protocol (NTP) configuration information for the StorageGRID system, so that operations performed on separate servers can be kept synchronized.

About this task

You must specify IPv4 addresses for the NTP servers.

You must specify external NTP servers. The specified NTP servers must use the NTP protocol.

You must specify four NTP server references of Stratum 3 or better to prevent issues with time drift.



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

[Support boundary to configure the Windows Time service for high-accuracy environments](#)

The external NTP servers are used by the nodes to which you previously assigned Primary NTP roles.



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

Steps

1. Specify the IPv4 addresses for at least four NTP servers in the **Server 1** to **Server 4** text boxes.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

The screenshot shows the NetApp StorageGRID installation wizard. At the top, there's a blue header with "NetApp® StorageGRID®" and a "Help" link. Below the header is a progress bar with eight steps: 1. License, 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP (highlighted in blue), 6. DNS, 7. Passwords, and 8. Summary. Below the progress bar, the "Network Time Protocol" section is active. It contains the instruction: "Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync." There are four input fields labeled "Server 1" through "Server 4". Server 1 contains "10.60.248.183", Server 2 contains "10.227.204.142", Server 3 contains "10.235.48.111", and Server 4 contains "0.0.0.0". To the right of the Server 4 field is a plus sign (+) to add more servers.

3. Select **Next**.

Specify DNS server information

You must specify DNS information for your StorageGRID system, so that you can access external servers using hostnames instead of IP addresses.

About this task

Specifying [DNS server information](#) allows you to use Fully Qualified Domain Name (FQDN) hostnames rather than IP addresses for email notifications and AutoSupport.

To ensure proper operation, specify two or three DNS servers. If you specify more than three, it is possible that only three will be used because of known OS limitations on some platforms. If you have routing restrictions in your environment, you can [customize the DNS server list](#) for individual nodes (typically all nodes at a site) to use a different set of up to three DNS servers.

If possible, use DNS servers that each site can access locally to ensure that an islanded site can resolve the FQDNs for external destinations.

Steps

1. Specify the IPv4 address for at least one DNS server in the **Server 1** text box.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

The screenshot shows the NetApp StorageGRID installation wizard. At the top is a blue header with "NetApp® StorageGRID®" and a "Help" dropdown. Below the header is a progress bar with eight steps: 1. License, 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP, 6. DNS (highlighted in blue), 7. Passwords, and 8. Summary. Below the progress bar is the "Domain Name Service" section. It contains a paragraph of instructions: "Enter the IP address for at least one Domain Name System (DNS) server, so that server hostnames can be used instead of IP addresses. Specifying at least two DNS servers is recommended. Configuring DNS enables server connectivity, email notifications, and NetApp AutoSupport." Below this text are two input fields. The first is labeled "Server 1" and contains the IP address "10.224.223.130". To its right is a red "X" icon. The second is labeled "Server 2" and contains the IP address "10.224.223.136". To its right is a red "+ X" icon.

The best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

Specify the StorageGRID system passwords

As part of installing your StorageGRID system, you need to enter the passwords to use to secure your system and perform maintenance tasks.

About this task

Use the Install passwords page to specify the provisioning passphrase and the grid management root user password.

- The provisioning passphrase is used as an encryption key and is not stored by the StorageGRID system.
- You must have the provisioning passphrase for installation, expansion, and maintenance procedures, including downloading the Recovery Package. Therefore, it is important that you store the provisioning passphrase in a secure location.
- You can change the provisioning passphrase from the Grid Manager if you have the current one.
- The grid management root user password can be changed using the Grid Manager.
- Randomly generated command line console and SSH passwords are stored in the `Passwords.txt` file in the Recovery Package.

Steps

1. In **Provisioning Passphrase**, enter the provisioning passphrase that will be required to make changes to the grid topology of your StorageGRID system.

Store the provisioning passphrase in a secure place.



If after the installation completes and you want to change the provisioning passphrase later, you can use the Grid Manager. Select **CONFIGURATION > Access control > Grid passwords**.

2. In **Confirm Provisioning Passphrase**, reenter the provisioning passphrase to confirm it.
3. In **Grid Management Root User Password**, enter the password to use to access the Grid Manager as the "root" user.

Store the password in a secure place.

4. In **Confirm Root User Password**, reenter the Grid Manager password to confirm it.

NetApp® StorageGRID® Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 **Passwords** 8 Summary

Passwords

Enter secure passwords that meet your organization's security policies. A text file containing the command line passwords must be downloaded during the final installation step.

Provisioning Passphrase

Confirm Provisioning Passphrase

Grid Management Root User Password

Confirm Root User Password

☒ Create random command line passwords.

5. If you are installing a grid for proof of concept or demo purposes, optionally clear the **Create random command line passwords** checkbox.

For production deployments, random passwords should always be used for security reasons. Clear **Create random command line passwords** only for demo grids if you want to use default passwords to access grid nodes from the command line using the "root" or "admin" account.



You are prompted to download the Recovery Package file (sgws-recovery-package-id-revision.zip) after you click **Install** on the Summary page. You must [download this file](#) to complete the installation. The passwords required to access the system are stored in the Passwords.txt file, contained in the Recovery Package file.

- Click **Next**.

Review your configuration and complete installation

You must carefully review the configuration information you have entered to ensure that the installation completes successfully.

Steps

- View the **Summary** page.

NetApp® StorageGRID®
Help

Install

1 License
2 Sites
3 Grid Network
4 Grid Nodes
5 NTP
6 DNS
7 Passwords
8 Summary

Summary

Verify that all of the grid configuration information is correct, and then click Install. You can view the status of each grid node as it installs. Click the Modify links to go back and change the associated information.

General Settings

Grid Name	Grid1	Modify License
Passwords	Auto-generated random command line passwords	Modify Passwords

Networking

NTP	10.60.248.183 10.227.204.142 10.235.48.111	Modify NTP
DNS	10.224.223.130 10.224.223.136	Modify DNS
Grid Network	172.16.0.0/21	Modify Grid Network

Topology

Topology	Atlanta	Modify Sites	Modify Grid Nodes
	Raleigh		
	dc1-adm1	dc1-g1	dc1-s1
	dc1-s2	dc1-s3	NetApp-SGA

- Verify that all of the grid configuration information is correct. Use the Modify links on the Summary page to go back and correct any errors.
- Click **Install**.



If a node is configured to use the Client Network, the default gateway for that node switches from the Grid Network to the Client Network when you click **Install**. If you lose connectivity, you must ensure that you are accessing the primary Admin Node through an accessible subnet. See [Networking guidelines](#) for details.

- Click **Download Recovery Package**.

When the installation progresses to the point where the grid topology is defined, you are prompted to download the Recovery Package file (.zip), and confirm that you can successfully access the contents of this file. You must download the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fail. The installation continues in the background, but you can't complete the

installation and access the StorageGRID system until you download and verify this file.

5. Verify that you can extract the contents of the .zip file, and then save it in two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

6. Select the **I have successfully downloaded and verified the Recovery Package file** checkbox, and click **Next**.

If the installation is still in progress, the status page appears. This page indicates the progress of the installation for each grid node.

Installation Status

If necessary, you may [Download the Recovery Package file](#) again.

						Search	
Name	Site	Grid Network IPv4 Address	Progress	Stage			
dc1-adm1	Site1	172.16.4.215/21		Starting services			
dc1-g1	Site1	172.16.4.216/21		Complete			
dc1-s1	Site1	172.16.4.217/21		Waiting for Dynamic IP Service peers			
dc1-s2	Site1	172.16.4.218/21		Downloading hotfix from primary Admin if needed			
dc1-s3	Site1	172.16.4.219/21		Downloading hotfix from primary Admin if needed			

When the Complete stage is reached for all grid nodes, the sign-in page for the Grid Manager appears.

7. Sign in to the Grid Manager using the "root" user and the password you specified during the installation.

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You can't set up DHCP during configuration.



Nodes reboot when the Grid Network configuration is changed by DHCP, which can cause outages if a DHCP change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to reestablish these connections.

Installation REST API

StorageGRID provides the StorageGRID Installation API for performing installation tasks.

The API uses the Swagger open source API platform to provide the API documentation. Swagger allows both

developers and non-developers to interact with the API in a user interface that illustrates how the API responds to parameters and options. This documentation assumes that you are familiar with standard web technologies and the JSON data format.



Any API operations you perform using the API Documentation webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Each REST API command includes the API's URL, an HTTP action, any required or optional URL parameters, and an expected API response.

StorageGRID Installation API

The StorageGRID Installation API is only available when you are initially configuring your StorageGRID system, and if you need to perform a primary Admin Node recovery. The Installation API can be accessed over HTTPS from the Grid Manager.

To access the API documentation, go to the installation web page on the primary Admin Node and select **Help > API documentation** from the menu bar.

The StorageGRID Installation API includes the following sections:

- **config** — Operations related to the product release and versions of the API. You can list the product release version and the major versions of the API supported by that release.
- **grid** — Grid-level configuration operations. You can get and update grid settings, including grid details, Grid Network subnets, grid passwords, and NTP and DNS server IP addresses.
- **nodes** — Node-level configuration operations. You can retrieve a list of grid nodes, delete a grid node, configure a grid node, view a grid node, and reset a grid node's configuration.
- **provision** — Provisioning operations. You can start the provisioning operation and view the status of the provisioning operation.
- **recovery** — Primary Admin Node recovery operations. You can reset information, upload the Recover Package, start the recovery, and view the status of the recovery operation.
- **recovery-package** — Operations to download the Recovery Package.
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.
- **temporary-password** — Operations on the temporary password to secure the mgmt-api during installation.

Where to go next

After completing an installation, perform the required integration and configuration tasks. You can perform the optional tasks as needed.

Required tasks

- [Create a tenant account](#) for the S3 client protocol that will be used to store objects on your StorageGRID system.
- [Control system access](#) by configuring groups and user accounts. Optionally, you can [configure a federated identity source](#) (such as Active Directory or OpenLDAP), so you can import administration groups and users. Or, you can [create local groups and users](#).
- Integrate and test the [S3 API](#) client applications you will use to upload objects to your StorageGRID

system.

- [Configure the information lifecycle management \(ILM\) rules and ILM policy](#) you want to use to protect object data.
- If your installation includes appliance Storage Nodes, use SANtricity OS to complete the following tasks:
 - Connect to each StorageGRID appliance.
 - Verify receipt of AutoSupport data.

See [Set up hardware](#).

- Review and follow the [StorageGRID system hardening guidelines](#) to eliminate security risks.
- [Configure email notifications for system alerts](#).

Optional tasks

- [Update grid node IP addresses](#) if they have changed since you planned your deployment and generated the Recovery Package.
- [Configure storage encryption](#), if required.
- [Configure storage compression](#) to reduce the size of stored objects, if required.
- [Configure VLAN interfaces](#) to isolate and partition network traffic, if required.
- [Configure high availability groups](#) to improve connection availability for the Grid Manager, Tenant Manager, and S3 clients, if required.
- [Configure load balancer endpoints](#) for S3 client connectivity, if required.

Troubleshoot installation issues

If any problems occur while installing your StorageGRID system, you can access the installation log files. Technical support might also need to use the installation log files to resolve issues.

The following installation log files are available from the container that is running each node:

- `/var/local/log/install.log` (found on all grid nodes)
- `/var/local/log/gdu-server.log` (found on the primary Admin Node)

The following installation log files are available from the host:

- `/var/log/storagegrid/daemon.log`
- `/var/log/storagegrid/nodes/node-name.log`

To learn how to access the log files, see [Collect log files and system data](#).

Related information

[Troubleshoot a StorageGRID system](#)

Example `/etc/sysconfig/network-scripts`

You can use the example files to aggregate four Linux physical interfaces into a single

LACP bond and then establish three VLAN interfaces subtending the bond for use as StorageGRID Grid, Admin, and Client Network interfaces.

Physical interfaces

Note that the switches at the other ends of the links must also treat the four ports as a single LACP trunk or port channel, and must pass at least the three referenced VLANs with tags.

/etc/sysconfig/network-scripts/ifcfg-ens160

```
TYPE=Ethernet
NAME=ens160
UUID=011b17dd-642a-4bb9-acae-d71f7e6c8720
DEVICE=ens160
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

/etc/sysconfig/network-scripts/ifcfg-ens192

```
TYPE=Ethernet
NAME=ens192
UUID=e28eb15f-76de-4e5f-9a01-c9200b58d19c
DEVICE=ens192
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

/etc/sysconfig/network-scripts/ifcfg-ens224

```
TYPE=Ethernet
NAME=ens224
UUID=b0e3d3ef-7472-4cde-902c-ef4f3248044b
DEVICE=ens224
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

/etc/sysconfig/network-scripts/ifcfg-ens256

```
TYPE=Ethernet
NAME=ens256
UUID=7cf7aabc-3e4b-43d0-809a-1e2378faa4cd
DEVICE=ens256
ONBOOT=yes
MASTER=bond0
SLAVE=yes
```

Bond interface

/etc/sysconfig/network-scripts/ifcfg-bond0

```
DEVICE=bond0
TYPE=Bond
BONDING_MASTER=yes
NAME=bond0
ONBOOT=yes
BONDING_OPTS=mode=802.3ad
```

VLAN interfaces

/etc/sysconfig/network-scripts/ifcfg-bond0.1001

```
VLAN=yes
TYPE=Vlan
DEVICE=bond0.1001
PHYSDEV=bond0
VLAN_ID=1001
REORDER_HDR=0
BOOTPROTO=none
UUID=296435de-8282-413b-8d33-c4dd40fca24a
ONBOOT=yes
```

/etc/sysconfig/network-scripts/ifcfg-bond0.1002

```
VLAN=yes
TYPE=Vlan
DEVICE=bond0.1002
PHYSDEV=bond0
VLAN_ID=1002
REORDER_HDR=0
BOOTPROTO=none
UUID=dbaaec72-0690-491c-973a-57b7dd00c581
ONBOOT=yes
```

/etc/sysconfig/network-scripts/ifcfg-bond0.1003

```
VLAN=yes
TYPE=Vlan
DEVICE=bond0.1003
PHYSDEV=bond0
VLAN_ID=1003
REORDER_HDR=0
BOOTPROTO=none
UUID=d1af4b30-32f5-40b4-8bb9-71a2fbf809a1
ONBOOT=yes
```

Install StorageGRID on Ubuntu or Debian

Quick start for installing StorageGRID on Ubuntu or Debian

Follow these high-level steps to install an Ubuntu or Debian StorageGRID node.

1

Preparation

- Learn about [StorageGRID architecture and network topology](#).
- Learn about the specifics of [StorageGRID networking](#).
- Gather and prepare the [Required information and materials](#).
- Prepare the required [CPU and RAM](#).
- Provide for [storage and performance requirements](#).
- [Prepare the Linux servers](#) that will host your StorageGRID nodes.

2

Deployment

Deploy grid nodes. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.

- To deploy software-based grid nodes on the hosts you prepared in step 1, use the Linux command line and [node configuration files](#).
- To deploy StorageGRID appliance nodes, follow the [Quick start for hardware installation](#).

3

Configuration

When all nodes have been deployed, use the Grid Manager to [configure the grid and complete the installation](#).

Automate the installation

To save time and provide consistency, you can automate the installation of the StorageGRID host service and the configuration of grid nodes.

- Use a standard orchestration framework such as Ansible, Puppet, or Chef to automate:
 - Installation of Ubuntu or Debian
 - Configuration of networking and storage
 - Installation of the container engine and the StorageGRID host service
 - Deployment of virtual grid nodes

See [Automate the installation and configuration of the StorageGRID host service](#).

- After you deploy grid nodes, [automate the configuration of the StorageGRID system](#) using the Python configuration script provided in the installation archive.
- [Automate the installation and configuration of appliance grid nodes](#)
- If you are an advanced developer of StorageGRID deployments, automate the installation of grid nodes by using the [installation REST API](#).

Plan and prepare for installation on Ubuntu or Debian

Required information and materials

Before you install StorageGRID, gather and prepare the required information and materials.

Required information

Network plan

Which networks you intend to attach to each StorageGRID node. StorageGRID supports multiple networks for traffic separation, security, and administrative convenience.

See the StorageGRID [Networking guidelines](#).

Network information

IP addresses to assign to each grid node and the IP addresses of the DNS and NTP servers.

Servers for grid nodes

Identify a set of servers (physical, virtual, or both) that, in aggregate, provide sufficient resources to support the number and type of StorageGRID nodes you plan to deploy.



If your StorageGRID installation will not use StorageGRID appliance (hardware) Storage Nodes, you must use hardware RAID storage with battery-backed write cache (BBWC). StorageGRID does not support the use of virtual storage area networks (vSANs), software RAID, or no RAID protection.

Node migration (if needed)

Understand the [requirements for node migration](#), if you want to perform scheduled maintenance on physical hosts without any service interruption.

Related information

[NetApp Interoperability Matrix Tool](#)

Required materials

NetApp StorageGRID license

You must have a valid, digitally signed NetApp license.



A non-production license, which can be used for testing and proof of concept grids, is included in the StorageGRID installation archive.

StorageGRID installation archive

[Download the StorageGRID installation archive and extract the files.](#)

Service laptop

The StorageGRID system is installed through a service laptop.

The service laptop must have:

- Network port
- SSH client (for example, PuTTY)
- [Supported web browser](#)

StorageGRID documentation

- [Release notes](#)
- [Instructions for administering StorageGRID](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archive and extract the required files. Optionally, you can manually verify the files in the installation package.

Steps

1. Go to the [NetApp Downloads page for StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.
3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead statement appears, read it and select the checkbox.



You must apply any required hotfixes after you install the StorageGRID release. For more information, see the [hotfix procedure in the recovery and maintenance instructions](#)

5. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column, select the .tgz or .zip installation archive for Ubuntu or Debian.



Select the .zip file if you are running Windows on the service laptop.

7. Save the installation archive.
8. If you need to verify the installation archive:
 - a. Download the StorageGRID code signature verification package. The file name for this package uses the format `StorageGRID_<version-number>_Code_Signature_Verification_Package.tar.gz`, where `<version-number>` is the StorageGRID software version.
 - b. Follow the steps to [manually verify the installation files](#).
9. Extract the files from the installation archive.
10. Choose the files you need.

The files you need depends on your planned grid topology and how you will deploy your StorageGRID system.



The paths listed in the table are relative to the top-level directory installed by the extracted installation archive.

Path and file name	Description
<code>./debs/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./debs/NLF000000.txt</code>	A non-production NetApp License File that you can use for testing and proof of concept deployments.
<code>./debs/storagegrid-webscale-images-version-SHA.deb</code>	DEB package for installing the StorageGRID node images on Ubuntu or Debian hosts.
<code>./debs/storagegrid-webscale-images-version-SHA.deb.md5</code>	MD5 checksum for the file <code>/debs/storagegrid-webscale-images-version-SHA.deb</code> .
<code>./debs/storagegrid-webscale-service-version-SHA.deb</code>	DEB package for installing the StorageGRID host service on Ubuntu or Debian hosts.
Deployment scripting tool	Description
<code>./debs/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.

Path and file name	Description
<code>./debs/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./debs/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate integration.
<code>./debs/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./debs/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./debs/extras/ansible</code>	Example Ansible role and playbook for configuring Ubuntu or Debian hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
<code>./debs/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
<code>./debs/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.
<code>./debs/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

Manually verify installation files (optional)

If necessary, you can manually verify the files in the StorageGRID installation archive.

Before you begin

You have [downloaded the verification package](#) from the [NetApp Downloads page for StorageGRID](#).

Steps

1. Extract the artifacts from the verification package:

```
tar -xf StorageGRID_11.9.0_Code_Signature_Verification_Package.tar.gz
```

2. Ensure that these artifacts were extracted:

- Leaf certificate: Leaf-Cert.pem
- Certificate chain: CA-Int-Cert.pem
- Time stamp response chain: TS-Cert.pem
- Checksum file: sha256sum
- Checksum signature: sha256sum.sig
- Time stamp response file: sha256sum.sig.tsr

3. Use the chain to verify the leaf certificate is valid.

Example: `openssl verify -CAfile CA-Int-Cert.pem Leaf-Cert.pem`

Expected output: Leaf-Cert.pem: OK

4. If step 2 failed because of an expired leaf certificate, use the tsr file to verify.

Example: `openssl ts -CAfile CA-Int-Cert.pem -untrusted TS-Cert.pem -verify -data sha256sum.sig -in sha256sum.sig.tsr`

Expected output includes: Verification: OK

5. Create a public key file from the leaf certificate.

Example: `openssl x509 -pubkey -noout -in Leaf-Cert.pem > Leaf-Cert.pub`

Expected output: none

6. Use the public key to verify the sha256sum file against sha256sum.sig.

Example: `openssl dgst -sha256 -verify Leaf-Cert.pub -signature sha256sum.sig sha256sum`

Expected output: Verified OK

7. Verify the sha256sum file content against newly created checksums.

Example: `sha256sum -c sha256sum`

Expected output: `<filename>: OK`
`<filename>` is the name of the archive file you downloaded.

8. [Complete the remaining steps](#) to extract and choose the appropriate installation files.

Software requirements for Ubuntu and Debian

You can use a virtual machine to host any type of StorageGRID node. You need one virtual machine for each grid node.

To install StorageGRID on Ubuntu or Debian, you must install some third-party software packages. Some supported Linux distributions don't contain these packages by default. The software package versions that

StorageGRID installations are tested on include those listed on this page.

If you select a Linux distribution and container runtime installation option that requires any of these packages, and they are not installed automatically by the Linux distribution, install one of the versions listed here if available from your provider or the supporting vendor for your Linux distribution. Otherwise, use the default package versions available from your vendor.

All installation options require either Podman or Docker. Do not install both packages. Install only the package required by your installation option.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

Python versions tested

- 3.5.2-2
- 3.6.8-2
- 3.6.8-38
- 3.6.9-1
- 3.7.3-1
- 3.8.10-0
- 3.9.2-1
- 3.9.10-2
- 3.9.16-1
- 3.10.6-1
- 3.11.2-6

Podman versions tested

- 3.2.3-0
- 3.4.4+ds1
- 4.1.1-7
- 4.2.0-11
- 4.3.1+ds1-8+b1
- 4.4.1-8
- 4.4.1-12

Docker versions tested



Docker support is deprecated and will be removed in a future release.

- Docker-CE 20.10.7
- Docker-CE 20.10.20-3
- Docker-CE 23.0.6-1
- Docker-CE 24.0.2-1

- Docker-CE 24.0.4-1
- Docker-CE 24.0.5-1
- Docker-CE 24.0.7-1
- 1.5-2

CPU and RAM requirements

Before installing StorageGRID software, verify and configure the hardware so that it is ready to support the StorageGRID system.

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: Dependent on the total RAM available and the amount of non-StorageGRID software running on the system
 - Generally, at least 24 GB per node, and 2 to 16 GB less than the total system RAM
 - A minimum of 64 GB for each tenant that will have approximately 5,000 buckets

Software-based metadata-only node resources must match the existing Storage Nodes resources. For example:

- If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)
- If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0).

When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.

Ensure that the number of StorageGRID nodes you plan to run on each physical or virtual host does not exceed the number of CPU cores or the physical RAM available. If the hosts aren't dedicated to running StorageGRID (not recommended), be sure to consider the resource requirements of the other applications.



Monitor your CPU and memory usage regularly to ensure that these resources continue to accommodate your workload. For example, doubling the RAM and CPU allocation for virtual Storage Nodes would provide similar resources to those provided for StorageGRID appliance nodes. Additionally, if the amount of metadata per node exceeds 500 GB, consider increasing the RAM per node to 48 GB or more. For information about managing object metadata storage, increasing the Metadata Reserved Space setting, and monitoring CPU and memory usage, see the instructions for [administering](#), [monitoring](#), and [upgrading](#) StorageGRID.

If hyperthreading is enabled on the underlying physical hosts, you can provide 8 virtual cores (4 physical cores) per node. If hyperthreading is not enabled on the underlying physical hosts, you must provide 8 physical cores per node.

If you are using virtual machines as hosts and have control over the size and number of VMs, you should use a single VM for each StorageGRID node and size the VM accordingly.

For production deployments, you should not run multiple Storage Nodes on the same physical storage hardware or virtual host. Each Storage Node in a single StorageGRID deployment should be in its own isolated failure domain. You can maximize the durability and availability of object data if you ensure that a single hardware failure can only impact a single Storage Node.

See also [Storage and performance requirements](#).

Storage and performance requirements

You must understand the storage requirements for StorageGRID nodes, so you can provide enough space to support the initial configuration and future storage expansion.

StorageGRID nodes require three logical categories of storage:

- **Container pool** — Performance-tier (10K SAS or SSD) storage for the node containers, which will be assigned to the Docker storage driver when you install and configure Docker on the hosts that will support your StorageGRID nodes.
- **System data** — Performance-tier (10K SAS or SSD) storage for per-node persistent storage of system data and transaction logs, which the StorageGRID host services will consume and map into individual nodes.
- **Object data** — Performance-tier (10K SAS or SSD) storage and capacity-tier (NL-SAS/SATA) bulk storage for the persistent storage of object data and object metadata.

You must use RAID-backed block devices for all storage categories. Non-redundant disks, SSDs, or JBODs aren't supported. You can use shared or local RAID storage for any of the storage categories; however, if you want to use the node migration capability in StorageGRID, you must store both system data and object data on shared storage. For more information, see [Node container migration requirements](#).

Performance requirements

The performance of the volumes used for the container pool, system data, and object metadata significantly impacts the overall performance of the system. You should use performance-tier (10K SAS or SSD) storage for these volumes to ensure adequate disk performance in terms of latency, input/output operations per second (IOPS), and throughput. You can use capacity-tier (NL-SAS/SATA) storage for the persistent storage of object data.

The volumes used for the container pool, system data, and object data must have write-back caching enabled. The cache must be on a protected or persistent media.

Requirements for hosts that use NetApp ONTAP storage

If the StorageGRID node uses storage assigned from a NetApp ONTAP system, confirm that the volume does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

Number of hosts required

Each StorageGRID site requires a minimum of three Storage Nodes.



In a production deployment, don't run more than one Storage Node on a single physical or virtual host. Using a dedicated host for each Storage Node provides an isolated failure domain.

Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same hosts, or they can be deployed on their own dedicated hosts as required.

Number of storage volumes for each host

The following table shows the number of storage volumes (LUNs) required for each host and the minimum size required for each LUN, based on which nodes will be deployed on that host.

The maximum tested LUN size is 39 TB.



These numbers are for each host, not for the entire grid.

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
Container engine storage pool	Container pool	1	Total number of nodes × 100 GB
/var/local volume	System data	1 for each node on this host	90 GB
Storage Node	Object data	3 for each Storage Node on this host Note: A software-based Storage Node can have 1 to 48 storage volumes; at least 3 storage volumes are recommended.	12 TB (4 TB/LUN) See Storage requirements for Storage Nodes for more information.
Storage Node (metadata-only)	Object metadata	1	4 TB See Storage requirements for Storage Nodes for more information. Note: Only one rangedb is required for metadata-only Storage Nodes.
Admin Node audit logs	System data	1 for each Admin Node on this host	200 GB
Admin Node tables	System data	1 for each Admin Node on this host	200 GB



Depending on the audit level configured, the size of user inputs such as S3 object key name, and how much audit log data you need to preserve, you might need to increase the size of the audit log LUN on each Admin Node. Generally, a grid generates approximately 1 KB of audit data per S3 operation, which would mean that a 200 GB LUN would support 70 million operations per day or 800 operations per second for two to three days.

Minimum storage space for a host

The following table shows the minimum storage space required for each type of node. You can use this table to determine the minimum amount of storage you must provide to the host in each storage category, based on which nodes will be deployed on that host.



Disk snapshots can't be used to restore grid nodes. Instead, refer to the [grid node recovery](#) procedures for each type of node.

Type of node	Container pool	System data	Object data
Storage Node	100 GB	90 GB	4,000 GB
Admin Node	100 GB	490 GB (3 LUNs)	<i>not applicable</i>
Gateway Node	100 GB	90 GB	<i>not applicable</i>

Example: Calculating the storage requirements for a host

Suppose you plan to deploy three nodes on the same host: one Storage Node, one Admin Node, and one Gateway Node. You should provide a minimum of nine storage volumes to the host. You will need a minimum of 300 GB of performance-tier storage for the node containers, 670 GB of performance-tier storage for system data and transaction logs, and 12 TB of capacity-tier storage for object data.

Type of node	LUN purpose	Number of LUNs	LUN size
Storage Node	Docker storage pool	1	300 GB (100 GB/node)
Storage Node	<code>/var/local</code> volume	1	90 GB
Storage Node	Object data	3	12 TB (4 TB/LUN)
Admin Node	<code>/var/local</code> volume	1	90 GB
Admin Node	Admin Node audit logs	1	200 GB
Admin Node	Admin Node tables	1	200 GB
Gateway Node	<code>/var/local</code> volume	1	90 GB

Type of node	LUN purpose	Number of LUNs	LUN size
Total		9	Container pool: 300 GB System data: 670 GB Object data: 12,000 GB


Storage requirements for Storage Nodes

A software-based Storage Node can have 1 to 48 storage volumes; 3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.



An appliance Storage Node can also have up to 48 storage volumes.

As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. Any remaining space on storage volume 0 and any other storage volumes in the Storage Node are used exclusively for 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. The three copies of object metadata are evenly distributed across all Storage Nodes at each site.

When installing a grid with metadata-only Storage Nodes, the grid must also contain a minimum number of nodes for object storage. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

- For a single-site grid, at least two Storage Nodes are configured for objects and metadata.
- For a multi-site grid, at least one Storage Node per site are configured for objects and metadata.

When you assign space to volume 0 of a new Storage Node, you must ensure there is adequate space for that node's portion of all object metadata.

- At a minimum, you must assign at least 4 TB to volume 0.



If you use only one storage volume for a Storage Node and you assign 4 TB or less to the volume, the Storage Node might enter the storage read-only state on startup and store object metadata only.



If you assign less than 500 GB to volume 0 (non-production use only), 10% of the storage volume's capacity is reserved for metadata.

- Software-based metadata-only node resources must match the existing Storage Nodes resources. For example:
 - If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)
 - If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0).

When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.

- If you are installing a new system (StorageGRID 11.6 or higher) and each Storage Node has 128 GB or more of RAM, assign 8 TB or more to volume 0. Using a larger value for volume 0 can increase the space allowed for metadata on each Storage Node.
- When configuring different Storage Nodes for a site, use the same setting for volume 0 if possible. If a site contains Storage Nodes of different sizes, the Storage Node with the smallest volume 0 will determine the metadata capacity of that site.

For details, go to [Manage object metadata storage](#).

Node container migration requirements

The node migration feature allows you to manually move a node from one host to another. Typically, both hosts are in the same physical data center.

Node migration allows you to perform physical host maintenance without disrupting grid operations. You move all StorageGRID nodes, one at a time, to another host before taking the physical host offline. Migrating nodes requires only a short downtime for each node and should not affect operation or availability of grid services.

If you want to use the StorageGRID node migration feature, your deployment must meet additional requirements:

- Consistent network interface names across hosts in a single physical data center
- Shared storage for StorageGRID metadata and object repository volumes that is accessible by all hosts in a single physical data center. For example, you might use NetApp E-Series storage arrays.

If you are using virtual hosts and the underlying hypervisor layer supports VM migration, you might want to use

this capability instead of the node migration feature in StorageGRID. In this case, you can ignore these additional requirements.

Before performing migration or hypervisor maintenance, shut down the nodes gracefully. See the instructions for [shutting down a grid node](#).

VMware Live Migration not supported

When performing bare-metal installation on VMware VMs, OpenStack Live Migration and VMware live vMotion cause the virtual machine clock time to jump and aren't supported for grid nodes of any type. Though rare, incorrect clock times can result in loss of data or configuration updates.

Cold migration is supported. In cold migration, you shut down the StorageGRID nodes before migrating them between hosts. See the instructions for [shutting down a grid node](#).

Consistent network interface names

To move a node from one host to another, the StorageGRID host service needs to have some confidence that the external network connectivity the node has at its current location can be duplicated at the new location. It gets this confidence through the use of consistent network interface names in the hosts.

Suppose, for example, that StorageGRID NodeA running on Host1 has been configured with the following interface mappings:

```
eth0  →  bond0.1001
eth1  →  bond0.1002
eth2  →  bond0.1003
```

The lefthand side of the arrows corresponds to the traditional interfaces as viewed from within a StorageGRID container (that is, the Grid, Admin, and Client Network interfaces, respectively). The righthand side of the arrows corresponds to the actual host interfaces providing these networks, which are three VLAN interfaces subordinate to the same physical interface bond.

Now, suppose you want to migrate NodeA to Host2. If Host2 also has interfaces named bond0.1001, bond0.1002, and bond0.1003, the system will allow the move, assuming that the like-named interfaces will provide the same connectivity on Host2 as they do on Host1. If Host2 does not have interfaces with the same names, the move will not be allowed.

There are many ways to achieve consistent network interface naming across multiple hosts; see [Configure the host network](#) for some examples.

Shared storage

To achieve rapid, low-overhead node migrations, the StorageGRID node migration feature does not physically move node data. Instead, node migration is performed as a pair of export and import operations, as follows:

Steps

1. During the "node export" operation, a small amount of persistent state data is extracted from the node container running on HostA and cached on that node's system data volume. Then, the node container on HostA is deinstantiated.

2. During the "node import" operation, the node container on HostB that uses the same network interface and block storage mappings that were in effect on HostA is instantiated. Then, the cached persistent state data is inserted into the new instance.

Given this mode of operation, all of the node's system data and object storage volumes must be accessible from both HostA and HostB for the migration to be allowed, and to work. In addition, they must have been mapped into the node using names that are guaranteed to refer to the same LUNs on HostA and HostB.

The following example shows one solution for block device mapping for a StorageGRID Storage Node, where DM multipathing is in use on the hosts, and the alias field has been used in `/etc/multipath.conf` to provide consistent, friendly block device names available on all hosts.

```
/var/local  → /dev/mapper/sgws-sn1-var-local
rangedb0   → /dev/mapper/sgws-sn1-rangedb0
rangedb1   → /dev/mapper/sgws-sn1-rangedb1
rangedb2   → /dev/mapper/sgws-sn1-rangedb2
rangedb3   → /dev/mapper/sgws-sn1-rangedb3
```

Prepare the hosts (Ubuntu or Debian)

How host-wide settings change during installation

On bare metal systems, StorageGRID makes some changes to host-wide `sysctl` settings.

The following changes are made:

```
# Recommended Cassandra setting: CASSANDRA-3563, CASSANDRA-13008, DataStax
documentation
vm.max_map_count = 1048575

# core file customization
# Note: for cores generated by binaries running inside containers, this
# path is interpreted relative to the container filesystem namespace.
# External cores will go nowhere, unless /var/local/core also exists on
# the host.
kernel.core_pattern = /var/local/core/%e.core.%p

# Set the kernel minimum free memory to the greater of the current value
or
# 512MiB if the host has 48GiB or less of RAM or 1.83GiB if the host has
more than 48GiB of RTAM
```

```

vm.min_free_kbytes = 524288

# Enforce current default swappiness value to ensure the VM system has
some
# flexibility to garbage collect behind anonymous mappings. Bump
watermark_scale_factor
# to help avoid OOM conditions in the kernel during memory allocation
bursts. Bump
# dirty_ratio to 90 because we explicitly fsync data that needs to be
persistent, and
# so do not require the dirty_ratio safety net. A low dirty_ratio combined
with a large
# working set (nr_active_pages) can cause us to enter synchronous I/O mode
unnecessarily,
# with deleterious effects on performance.
vm.swappiness = 60
vm.watermark_scale_factor = 200
vm.dirty_ratio = 90

# Turn off slow start after idle
net.ipv4.tcp_slow_start_after_idle = 0

# Tune TCP window settings to improve throughput
net.core.rmem_max = 8388608
net.core.wmem_max = 8388608
net.ipv4.tcp_rmem = 4096 524288 8388608
net.ipv4.tcp_wmem = 4096 262144 8388608
net.core.netdev_max_backlog = 2500

# Turn on MTU probing
net.ipv4.tcp_mtu_probing = 1

# Be more liberal with firewall connection tracking
net.ipv4.netfilter.ip_conntrack_tcp_be_liberal = 1

# Reduce TCP keepalive time to reasonable levels to terminate dead
connections
net.ipv4.tcp_keepalive_time = 270
net.ipv4.tcp_keepalive_probes = 3
net.ipv4.tcp_keepalive_intvl = 30

# Increase the ARP cache size to tolerate being in a /16 subnet
net.ipv4.neigh.default.gc_thresh1 = 8192
net.ipv4.neigh.default.gc_thresh2 = 32768
net.ipv4.neigh.default.gc_thresh3 = 65536
net.ipv6.neigh.default.gc_thresh1 = 8192

```

```

net.ipv6.neigh.default.gc_thresh2 = 32768
net.ipv6.neigh.default.gc_thresh3 = 65536

# Disable IP forwarding, we are not a router
net.ipv4.ip_forward = 0

# Follow security best practices for ignoring broadcast ping requests
net.ipv4.icmp_echo_ignore_broadcasts = 1

# Increase the pending connection and accept backlog to handle larger
connection bursts.
net.core.somaxconn=4096
net.ipv4.tcp_max_syn_backlog=4096

```

Install Linux

You must install StorageGRID on all Ubuntu or Debian grid hosts. For a list of supported versions, use the NetApp Interoperability Matrix Tool.

Before you begin

Ensure your operating system meets StorageGRID's minimum kernel version requirements, as listed below. Use the command `uname -r` to get your operating system's kernel version, or consult with your OS vendor.

Note: Support for Ubuntu versions 18.04 and 20.04 have been deprecated and will be removed in a future release.

Ubuntu version	Minimum kernel version	Kernel package name
18.04.6 (deprecated)	5.4.0-150-generic	linux-image-5.4.0-150-generic/bionic-updates,bionic-security,now 5.4.0-150.167~18.04.1
20.04.5 (deprecated)	5.4.0-131-generic	linux-image-5.4.0-131-generic/focal-updates,now 5.4.0-131.147
22.04.1	5.15.0-47-generic	linux-image-5.15.0-47-generic/jammy-updates,jammy-security,now 5.15.0-47.51
24.04	6.8.0-31-generic	linux-image-6.8.0-31-generic/noble,now 6.8.0-31.31

Note: Support for Debian version 11 has been deprecated and will be removed in a future release.

Debian version	Minimum kernel version	Kernel package name
11 (deprecated)	5.10.0-18-amd64	linux-image-5.10.0-18-amd64/stable,now 5.10.150-1

Debian version	Minimum kernel version	Kernel package name
12	6.1.0-9-amd64	linux-image-6.1.0-9-amd64/stable,now 6.1.27-1

Steps

1. Install Linux on all physical or virtual grid hosts according to the distributor's instructions or your standard procedure.



Don't install any graphical desktop environments. When installing Ubuntu, you must select **standard system utilities**. Selecting **OpenSSH server** is recommended to enable ssh access to your Ubuntu hosts. All other options can remain cleared.

2. Ensure that all hosts have access to Ubuntu or Debian package repositories.
3. If swap is enabled:
 - a. Run the following command: `$ sudo swapoff --all`
 - b. Remove all swap entries from `/etc/fstab` to persist the settings.



Failing to disable swap entirely can severely lower performance.

Understand AppArmor profile installation

If you are operating in a self-deployed Ubuntu environment and using the AppArmor mandatory access control system, the AppArmor profiles associated with packages you install on the base system might be blocked by the corresponding packages installed with StorageGRID.

By default, AppArmor profiles are installed for packages that you install on the base operating system. When you run these packages from the StorageGRID system container, the AppArmor profiles are blocked. The DHCP, MySQL, NTP, and tcdump base packages conflict with AppArmor, and other base packages might also conflict.

You have two choices for handling AppArmor profiles:

- Disable individual profiles for the packages installed on the base system that overlap with the packages in the StorageGRID system container. When you disable individual profiles, an entry appears in the StorageGRID log files indicating that AppArmor is enabled.

Use the following commands:

```
sudo ln -s /etc/apparmor.d/<profile.name> /etc/apparmor.d/disable/  
sudo apparmor_parser -R /etc/apparmor.d/<profile.name>
```

Example:

```
sudo ln -s /etc/apparmor.d/bin.ping /etc/apparmor.d/disable/  
sudo apparmor_parser -R /etc/apparmor.d/bin.ping
```

- Disable AppArmor altogether. For Ubuntu 9.10 or later, follow the instructions in the Ubuntu online community: [Disable AppArmor](#). Disabling AppArmor altogether might not be possible on newer Ubuntu versions.

After you disable AppArmor, no entries indicating that AppArmor is enabled will appear in the StorageGRID log files.

Configure the host network (Ubuntu or Debian)

After completing the Linux installation on your hosts, you might need to perform some additional configuration to prepare a set of network interfaces on each host that are suitable for mapping into the StorageGRID nodes you will deploy later.

Before you begin

- You have reviewed the [StorageGRID networking guidelines](#).
- You have reviewed the information about [node container migration requirements](#).
- If you are using virtual hosts, you have read the [considerations and recommendations for MAC address cloning](#) before configuring the host network.



If you are using VMs as hosts, you should select VMXNET 3 as the virtual network adapter. The VMware E1000 network adapter has caused connectivity issues with StorageGRID containers deployed on certain distributions of Linux.

About this task

Grid nodes must be able to access the Grid Network and, optionally, the Admin and Client Networks. You provide this access by creating mappings that associate the host's physical interface to the virtual interfaces for each grid node. When creating host interfaces, use friendly names to facilitate deployment across all hosts, and to enable migration.

The same interface can be shared between the host and one or more nodes. For example, you might use the same interface for host access and node Admin Network access, to facilitate host and node maintenance. Although the same interface can be shared between the host and individual nodes, all must have different IP addresses. IP addresses can't be shared between nodes or between the host and any node.

You can use the same host network interface to provide the Grid Network interface for all StorageGRID nodes on the host; you can use a different host network interface for each node; or you can do something in between. However, you would not typically provide the same host network interface as both the Grid and Admin Network interfaces for a single node, or as the Grid Network interface for one node and the Client Network interface for another.

You can complete this task in many ways. For example, if your hosts are virtual machines and you are deploying one or two StorageGRID nodes for each host, you can create the correct number of network interfaces in the hypervisor, and use a 1-to-1 mapping. If you are deploying multiple nodes on bare metal hosts for production use, you can leverage the Linux networking stack's support for VLAN and LACP for fault tolerance and bandwidth sharing. The following sections provide detailed approaches for both of these examples. You don't need to use either of these examples; you can use any approach that meets your needs.



Don't use bond or bridge devices directly as the container network interface. Doing so could prevent node start-up caused by a kernel issue with the use of MACVLAN with bond and bridge devices in the container namespace. Instead, use a non-bond device, such as a VLAN or virtual Ethernet (veth) pair. Specify this device as the network interface in the node configuration file.

Considerations and recommendations for MAC address cloning

MAC address cloning causes the container to use the MAC address of the host, and the host to use the MAC address of either an address you specify or a randomly generated one. You should use MAC address cloning to avoid the use of promiscuous mode network configurations.

Enabling MAC cloning

In certain environments, security can be enhanced through MAC address cloning because it enables you to use a dedicated virtual NIC for the Admin Network, Grid Network, and Client Network. Having the container use the MAC address of the dedicated NIC on the host allows you to avoid using promiscuous mode network configurations.



MAC address cloning is intended to be used with virtual server installations and might not function properly with all physical appliance configurations.



If a node fails to start due to a MAC cloning targeted interface being busy, you might need to set the link to "down" before starting node. Additionally, it is possible that the virtual environment might prevent MAC cloning on a network interface while the link is up. If a node fails to set the MAC address and start due to an interface being busy, setting the link to "down" before starting the node might fix the issue.

MAC address cloning is disabled by default and must be set by node configuration keys. You should enable it when you install StorageGRID.

There is one key for each network:

- `ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC`
- `GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC`
- `CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC`

Setting the key to "true" causes the container to use the MAC address of the host's NIC. Additionally, the host will then use the MAC address of the specified container network. By default, the container address is a randomly generated address, but if you have set one using the `_NETWORK_MAC` node configuration key, that address is used instead. The host and container will always have different MAC addresses.



Enabling MAC cloning on a virtual host without also enabling promiscuous mode on the hypervisor might cause Linux host networking using the host's interface to stop working.

MAC cloning use cases

There are two use cases to consider with MAC cloning:

- **MAC cloning not enabled:** When the `_CLONE_MAC` key in the node configuration file is not set, or set to "false," the host will use the host NIC MAC and the container will have a StorageGRID-generated MAC

unless a MAC is specified in the `_NETWORK_MAC` key. If an address is set in the `_NETWORK_MAC` key, the container will have the address specified in the `_NETWORK_MAC` key. This configuration of keys requires the use of promiscuous mode.

- **MAC cloning enabled:** When the `_CLONE_MAC` key in the node configuration file is set to "true," the container uses the host NIC MAC, and the host uses a StorageGRID-generated MAC unless a MAC is specified in the `_NETWORK_MAC` key. If an address is set in the `_NETWORK_MAC` key, the host uses the specified address instead of a generated one. In this configuration of keys, you should not use promiscuous mode.



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

To enable MAC cloning, see the [instructions for creating node configuration files](#).

MAC cloning example


Example of MAC cloning enabled with a host having MAC address of 11:22:33:44:55:66 for the interface `ens256` and the following keys in the node configuration file:

- `ADMIN_NETWORK_TARGET = ens256`
- `ADMIN_NETWORK_MAC = b2:9c:02:c2:27:10`
- `ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC = true`

Result: the host MAC for `ens256` is `b2:9c:02:c2:27:10` and the Admin Network MAC is `11:22:33:44:55:66`

Example 1: 1-to-1 mapping to physical or virtual NICs

Example 1 describes a simple physical interface mapping that requires little or no host-side configuration.



The Linux operating system creates the `ensXYZ` interfaces automatically during installation or boot, or when the interfaces are hot-added. No configuration is required other than ensuring that the interfaces are set to come up automatically after boot. You do have to determine which `ensXYZ` corresponds to which StorageGRID network (Grid, Admin, or Client) so you can provide the correct mappings later in the configuration process.

Note that the figure shows multiple StorageGRID nodes; however, you would normally use this configuration for single-node VMs.

If Switch 1 is a physical switch, you should configure the ports connected to interfaces $10G_1$ through $10G_3$ for access mode, and place them on the appropriate VLANs.

Example 2: LACP bond carrying VLANs


Example 2 assumes you are familiar with bonding network interfaces and with creating VLAN interfaces on the Linux distribution you are using.

About this task

Example 2 describes a generic, flexible, VLAN-based scheme that facilitates the sharing of all available network bandwidth across all nodes on a single host. This example is particularly applicable to bare metal hosts.

To understand this example, suppose you have three separate subnets for the Grid, Admin, and Client Networks at each data center. The subnets are on separate VLANs (1001, 1002, and 1003) and are presented to the host on a LACP-bonded trunk port (`bond0`). You would configure three VLAN interfaces on the bond: `bond0.1001`, `bond0.1002`, and `bond0.1003`.

If you require separate VLANs and subnets for node networks on the same host, you can add VLAN interfaces on the bond and map them into the host (shown as `bond0.1004` in the illustration).



Steps

1. Aggregate all physical network interfaces that will be used for StorageGRID network connectivity into a single LACP bond.

Use the same name for the bond on every host, for example, bond0.

2. Create VLAN interfaces that use this bond as their associated "physical device," using the standard VLAN interface naming convention `physdev-name.VLAN ID`.

Note that steps 1 and 2 require appropriate configuration on the edge switches terminating the other ends of the network links. The edge switch ports must also be aggregated into a LACP port channel, configured as a trunk, and allowed to pass all required VLANs.

Example interface configuration files for this per-host networking configuration scheme are provided.

Related information

[Example /etc/network/interfaces](#)

Configure host storage

You must allocate block storage volumes to each host.

Before you begin

You have reviewed the following topics, which provide information you need to accomplish this task:

- [Storage and performance requirements](#)
- [Node container migration requirements](#)

About this task

When allocating block storage volumes (LUNs) to hosts, use the tables in "Storage requirements" to determine the following:

- Number of volumes required for each host (based on the number and types of nodes that will be deployed on that host)
- Storage category for each volume (that is, System Data or Object Data)
- Size of each volume

You will use this information as well as the persistent name assigned by Linux to each physical volume when you deploy StorageGRID nodes on the host.



You don't need to partition, format, or mount any of these volumes; you just need to ensure they are visible to the hosts.



Only one object-data LUN is required for metadata-only Storage Nodes.

Avoid using "raw" special device files (`/dev/sdb`, for example) as you compose your list of volume names. These files can change across reboots of the host, which will impact proper operation of the system. If you are using iSCSI LUNs and Device Mapper Multipathing, consider using multipath aliases in the `/dev/mapper` directory, especially if your SAN topology includes redundant network paths to the shared storage. Alternatively, you can use the system-created softlinks under `/dev/disk/by-path/` for your persistent device names.

For example:

```
ls -l
$ ls -l /dev/disk/by-path/
total 0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:00:07.1-ata-2 -> ../../sr0
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0 ->
../../sda
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part1
-> ../../sda1
lrwxrwxrwx 1 root root 10 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:0:0-part2
-> ../../sda2
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:1:0 ->
../../sdb
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:2:0 ->
../../sdc
lrwxrwxrwx 1 root root 9 Sep 19 18:53 pci-0000:03:00.0-scsi-0:0:3:0 ->
../../sdd
```

Results will differ for each installation.

Assign friendly names to each of these block storage volumes to simplify the initial StorageGRID installation and future maintenance procedures. If you are using the device mapper multipath driver for redundant access to shared storage volumes, you can use the `alias` field in your `/etc/multipath.conf` file.

For example:

```
multipaths {
    multipath {
        wwid 3600a09800059d6df00005df2573c2c30
        alias docker-storage-volume-hostA
    }
    multipath {
        wwid 3600a09800059d6df00005df3573c2c30
        alias sgws-adml-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df4573c2c30
        alias sgws-adml-audit-logs
    }
    multipath {
        wwid 3600a09800059d6df00005df5573c2c30
        alias sgws-adml-tables
    }
    multipath {
        wwid 3600a09800059d6df00005df6573c2c30
        alias sgws-gw1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-var-local
    }
    multipath {
        wwid 3600a09800059d6df00005df7573c2c30
        alias sgws-sn1-rangedb-0
    }
    ...
}
```

Using the `alias` field in this way causes the aliases to appear as block devices in the `/dev/mapper` directory on the host, allowing you to specify a friendly, easily-validated name whenever a configuration or maintenance operation requires specifying a block storage volume.

If you are setting up shared storage to support StorageGRID node migration and using Device Mapper Multipathing, you can create and install a common `/etc/multipath.conf` on all co-located hosts. Just make sure to use a different Docker storage volume on each host. Using aliases and including the target hostname in the alias for each Docker storage volume LUN will make this easy to remember and is recommended.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

Related information

- [Storage and performance requirements](#)
- [Node container migration requirements](#)

Configure container engine storage volume

Before installing the container engine (Docker or Podman), you might need to format the storage volume and mount it.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

About this task

You can skip these steps if you plan to use local storage for the Docker storage volume and have sufficient space available on the host partition containing `/var/lib`.

Steps

1. Create a file system on the Docker storage volume:

```
sudo mkfs.ext4 docker-storage-volume-device
```

2. Mount the Docker storage volume:

```
sudo mkdir -p /var/lib/docker  
sudo mount docker-storage-volume-device /var/lib/docker
```

3. Add an entry for docker-storage-volume-device to `/etc/fstab`.

This step ensures that the storage volume will remount automatically after host reboots.

Install Docker

The StorageGRID system runs on Linux as a collection of Docker containers. Before you can install StorageGRID, you must install Docker.



Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release.

Steps

1. Install Docker by following the instructions for your Linux distribution.



If Docker is not included with your Linux distribution, you can download it from the Docker website.

2. Ensure Docker has been enabled and started by running the following two commands:

```
sudo systemctl enable docker
```

```
sudo systemctl start docker
```

3. Confirm you have installed the expected version of Docker by entering the following:

```
sudo docker version
```

The Client and Server versions must be 1.11.0 or later.

Related information

[Configure host storage](#)

Install StorageGRID host services

You use the StorageGRID DEB package to install the StorageGRID host services.

About this task

These instructions describe how to install the host services from the DEB packages. As an alternative, you can use the APT repository metadata included in the installation archive to install the DEB packages remotely. See the APT repository instructions for your Linux operating system.

Steps

1. Copy the StorageGRID DEB packages to each of your hosts, or make them available on shared storage.

For example, place them in the `/tmp` directory, so you can use the example command in the next step.

2. Log in to each host as root or using an account with sudo permission, and run the following commands.

You must install the `images` package first, and the `service` package second. If you placed the packages in a directory other than `/tmp`, modify the command to reflect the path you used.

```
sudo dpkg --install /tmp/storagegrid-webscale-images-version-SHA.deb
```

```
sudo dpkg --install /tmp/storagegrid-webscale-service-version-SHA.deb
```



Python 2.7 must already be installed before the StorageGRID packages can be installed. The `sudo dpkg --install /tmp/storagegrid-webscale-images-version-SHA.deb` command will fail until you have done so.

Automate the installation (Ubuntu or Debian)

You can automate the installation of the StorageGRID host service and the configuration of grid nodes.

About this task

Automating the deployment might be useful in any of the following cases:

- You already use a standard orchestration framework, such as Ansible, Puppet, or Chef, to deploy and configure physical or virtual hosts.
- You intend to deploy multiple StorageGRID instances.
- You are deploying a large, complex StorageGRID instance.

The StorageGRID host service is installed by a package and driven by configuration files that can be created interactively during a manual installation, or prepared ahead of time (or programmatically) to enable automated installation using standard orchestration frameworks. StorageGRID provides optional Python scripts for automating the configuration of StorageGRID appliances, and the whole StorageGRID system (the "grid"). You can use these scripts directly, or you can inspect them to learn how to use the StorageGRID Installation REST API in grid deployment and configuration tools you develop yourself.

Automate the installation and configuration of the StorageGRID host service

You can automate the installation of the StorageGRID host service using standard orchestration frameworks such as Ansible, Puppet, Chef, Fabric, or SaltStack.

The StorageGRID host service is packaged in a DEB and is driven by configuration files that can be prepared ahead of time (or programmatically) to enable automated installation. If you already use a standard orchestration framework to install and configure Ubuntu or Debian, adding StorageGRID to your playbooks or recipes should be straightforward.

You can automate these tasks:

1. Installing Linux
2. Configuring Linux
3. Configuring host network interfaces to meet StorageGRID requirements
4. Configuring host storage to meet StorageGRID requirements
5. Installing Docker
6. Installing the StorageGRID host service
7. Creating StorageGRID node configuration files in `/etc/storagegrid/nodes`
8. Validating StorageGRID node configuration files
9. Starting the StorageGRID host service

Example Ansible role and playbook

Example Ansible role and playbook are supplied with the installation archive in the `/extras` folder. The Ansible playbook shows how the `storagegrid` role prepares the hosts and installs StorageGRID onto the target servers. You can customize the role or playbook as necessary.

Automate the configuration of StorageGRID

After deploying the grid nodes, you can automate the configuration of the StorageGRID system.

Before you begin

- You know the location of the following files from the installation archive.

Filename	Description
<code>configure-storagegrid.py</code>	Python script used to automate the configuration
<code>configure-storagegrid.sample.json</code>	Example configuration file for use with the script
<code>configure-storagegrid.blank.json</code>	Blank configuration file for use with the script

- You have created a `configure-storagegrid.json` configuration file. To create this file, you can modify the example configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

About this task

You can use the `configure-storagegrid.py` Python script and the `configure-storagegrid.json` configuration file to automate the configuration of your StorageGRID system.



You can also configure the system using the Grid Manager or the Installation API.

Steps

1. Log in to the Linux machine you are using to run the Python script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/platform
```

where `platform` is `debs`, `rpms`, or `vsphere`.

3. Run the Python script and use the configuration file you created.

For example:

```
./configure-storagegrid.py ./configure-storagegrid.json --start-install
```

Result

A Recovery Package `.zip` file is generated during the configuration process, and it is downloaded to the directory where you are running the installation and configuration process. You must back up the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fails. For example, copy it to a secure, backed up network location and to a secure cloud storage location.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

If you specified that random passwords should be generated, open the `Passwords.txt` file and look for the passwords required to access your StorageGRID system.

```
#####  
##### The StorageGRID "Recovery Package" has been downloaded as: #####  
#####      ./sgws-recovery-package-994078-rev1.zip      #####  
#####   Safeguard this file as it will be needed in case of a   #####  
#####               StorageGRID node recovery.               #####  
#####
```

Your StorageGRID system is installed and configured when a confirmation message is displayed.

```
StorageGRID has been configured and installed.
```

Related information

[Installation REST API](#)

Deploy virtual grid nodes (Ubuntu or Debian)

Create node configuration files for Ubuntu or Debian deployments

Node configuration files are small text files that provide the information the StorageGRID host service needs to start a node and connect it to the appropriate network and block storage resources. Node configuration files are used for virtual nodes and aren't used for appliance nodes.

Location for node configuration files

Place the configuration file for each StorageGRID node in the `/etc/storagegrid/nodes` directory on the host where the node will run. For example, if you plan to run one Admin Node, one Gateway Node, and one Storage Node on HostA, you must place three node configuration files in `/etc/storagegrid/nodes` on HostA.

You can create the configuration files directly on each host using a text editor, such as `vim` or `nano`, or you can create them elsewhere and move them to each host.

Naming of node configuration files

The names of the configuration files are significant. The format is `node-name.conf`, where `node-name` is a name you assign to the node. This name appears in the StorageGRID Installer and is used for node maintenance operations, such as node migration.

Node names must follow these rules:

- Must be unique

- Must start with a letter
- Can contain the characters A through Z and a through z
- Can contain the numbers 0 through 9
- Can contain one or more hyphens (-)
- Must be no more than 32 characters, not including the `.conf` extension

Any files in `/etc/storagegrid/nodes` that don't follow these naming conventions will not be parsed by the host service.

If you have a multi-site topology planned for your grid, a typical node naming scheme might be:

```
site-nodetype-nodenum.conf
```

For example, you might use `dc1-adm1.conf` for the first Admin Node in Data Center 1, and `dc2-sn3.conf` for the third Storage Node in Data Center 2. However, you can use any scheme you like, as long as all node names follow the naming rules.

Contents of a node configuration file

A configuration file contains key/value pairs, with one key and one value per line. For each key/value pair, follow these rules:

- The key and the value must be separated by an equal sign (=) and optional whitespace.
- The keys can contain no spaces.
- The values can contain embedded spaces.
- Any leading or trailing whitespace is ignored.

The following table defines the values for all supported keys. Each key has one of the following designations:

- **Required:** Required for every node or for the specified node types
- **Best practice:** Optional, although recommended
- **Optional:** Optional for all nodes

Admin Network keys

ADMIN_IP

Value	Designation
Grid Network IPv4 address of the primary Admin Node for the grid to which this node belongs. Use the same value you specified for GRID_NETWORK_IP for the grid node with NODE_TYPE = VM_Admin_Node and ADMIN_ROLE = Primary. If you omit this parameter, the node attempts to discover a primary Admin Node using mDNS. How grid nodes discover the primary Admin Node Note: This value is ignored, and might be prohibited, on the primary Admin Node.	Best practice

ADMIN_NETWORK_CONFIG

Value	Designation
DHCP, STATIC, or DISABLED	Optional

ADMIN_NETWORK_ESL

Value	Designation
Comma-separated list of subnets in CIDR notation to which this node should communicate using the Admin Network gateway. Example: 172.16.0.0/21,172.17.0.0/21	Optional

ADMIN_NETWORK_GATEWAY

Value	Designation
IPv4 address of the local Admin Network gateway for this node. Must be on the subnet defined by ADMIN_NETWORK_IP and ADMIN_NETWORK_MASK. This value is ignored for DHCP-configured networks. Examples: 1.1.1.1 10.224.4.81	Required if ADMIN_NETWORK_ESL is specified. Optional otherwise.

ADMIN_NETWORK_IP

Value	Designation
IPv4 address of this node on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 1.1.1.1 10.224.4.81	Required when ADMIN_NETWORK_CONFIG = STATIC. Optional otherwise.

ADMIN_NETWORK_MAC

Value	Designation
The MAC address for the Admin Network interface in the container. This field is optional. If omitted, a MAC address will be generated automatically. Must be 6 pairs of hexadecimal digits separated by colons. Example: b2:9c:02:c2:27:10	Optional

ADMIN_NETWORK_MASK

Value	Designation
IPv4 netmask for this node, on the Admin Network. Specify this key when ADMIN_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 255.255.255.0 255.255.248.0	Required if ADMIN_NETWORK_IP is specified and ADMIN_NETWORK_CONFIG = STATIC. Optional otherwise.

ADMIN_NETWORK_MTU

Value	Designation
The maximum transmission unit (MTU) for this node on the Admin Network. Don't specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur. Examples: 1500 8192	Optional

ADMIN_NETWORK_TARGET

Value	Designation
Name of the host device that you will use for Admin Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or CLIENT_NETWORK_TARGET. Note: Don't use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair. Best practice: Specify a value even if this node will not initially have an Admin Network IP address. Then you can add an Admin Network IP address later, without having to reconfigure the node on the host. Examples: bond0.1002 ens256	Best practice

ADMIN_NETWORK_TARGET_TYPE

Value	Designation
Interface (This is the only supported value.)	Optional

ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Value	Designation
True or False Set the key to "true" to cause the StorageGRID container use the MAC address of the host host target interface on the Admin Network. Best practice: In networks where promiscuous mode would be required, use the ADMIN_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead. For more details on MAC cloning: • Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux) • Considerations and recommendations for MAC address cloning (Ubuntu or Debian)	Best practice

ADMIN_ROLE

Value	Designation
Primary or non-primary This key is only required when NODE_TYPE = VM_Admin_Node; don't specify it for other node types.	Required when NODE_TYPE = VM_Admin_Node Optional otherwise.

Block device keys

BLOCK_DEVICE_AUDIT_LOGS

Value	Designation
Path and name of the block device special file this node will use for persistent storage of audit logs. Examples: <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-adm1-audit-logs</pre>	Required for nodes with NODE_TYPE = VM_Admin_Node. Don't specify it for other node types.

BLOCK_DEVICE_RANGEDB_nnn

Value	Designation
Path and name of the block device special file this node will use for persistent object storage. This key is only required for nodes with <code>NODE_TYPE = VM_Storage_Node</code>; don't specify it for other node types. Only <code>BLOCK_DEVICE_RANGEDB_000</code> is required; the rest are optional. The block device specified for <code>BLOCK_DEVICE_RANGEDB_000</code> must be at least 4 TB; the others can be smaller. Don't leave gaps. If you specify <code>BLOCK_DEVICE_RANGEDB_005</code>, you must also specify <code>BLOCK_DEVICE_RANGEDB_004</code>. Note: For compatibility with existing deployments, two-digit keys are supported for upgraded nodes. Examples: <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-sn1-rangedb-000</pre>	Required: <code>BLOCK_DEVICE_RANGEDB_000</code> Optional: <code>BLOCK_DEVICE_RANGEDB_001</code> <code>BLOCK_DEVICE_RANGEDB_002</code> <code>BLOCK_DEVICE_RANGEDB_003</code> <code>BLOCK_DEVICE_RANGEDB_004</code> <code>BLOCK_DEVICE_RANGEDB_005</code> <code>BLOCK_DEVICE_RANGEDB_006</code> <code>BLOCK_DEVICE_RANGEDB_007</code> <code>BLOCK_DEVICE_RANGEDB_008</code> <code>BLOCK_DEVICE_RANGEDB_009</code> <code>BLOCK_DEVICE_RANGEDB_010</code> <code>BLOCK_DEVICE_RANGEDB_011</code> <code>BLOCK_DEVICE_RANGEDB_012</code> <code>BLOCK_DEVICE_RANGEDB_013</code> <code>BLOCK_DEVICE_RANGEDB_014</code> <code>BLOCK_DEVICE_RANGEDB_015</code>

BLOCK_DEVICE_TABLES

Value	Designation
Path and name of the block device special file this node will use for persistent storage of database tables. This key is only required for nodes with NODE_TYPE = VM_Admin_Node; don't specify it for other node types. Examples: <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-adml-tables</pre>	Required

BLOCK_DEVICE_VAR_LOCAL

Value	Designation
Path and name of the block device special file this node will use for its /var/local persistent storage. Examples: <pre>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</pre> <pre>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</pre> <pre>/dev/mapper/sgws-sn1-var-local</pre>	Required

Client Network keys

CLIENT_NETWORK_CONFIG

Value	Designation
DHCP, STATIC, or DISABLED	Optional

CLIENT_NETWORK_GATEWAY

Value	Designation

IPv4 address of the local Client Network gateway for this node, which must be on the subnet defined by <code>CLIENT_NETWORK_IP</code> and <code>CLIENT_NETWORK_MASK</code>. This value is ignored for DHCP-configured networks. Examples: 1.1.1.1 10.224.4.81	Optional

CLIENT_NETWORK_IP

Value	Designation
IPv4 address of this node on the Client Network. This key is only required when <code>CLIENT_NETWORK_CONFIG = STATIC</code>; don't specify it for other values. Examples: 1.1.1.1 10.224.4.81	Required when <code>CLIENT_NETWORK_CONFIG = STATIC</code> Optional otherwise.

CLIENT_NETWORK_MAC

Value	Designation
The MAC address for the Client Network interface in the container. This field is optional. If omitted, a MAC address will be generated automatically. Must be 6 pairs of hexadecimal digits separated by colons. Example: <code>b2:9c:02:c2:27:20</code>	Optional

CLIENT_NETWORK_MASK

Value	Designation
IPv4 netmask for this node on the Client Network. Specify this key when CLIENT_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 255.255.255.0 255.255.248.0	Required if CLIENT_NETWORK_IP is specified and CLIENT_NETWORK_CONFIG = STATIC Optional otherwise.

CLIENT_NETWORK_MTU

Value	Designation
The maximum transmission unit (MTU) for this node on the Client Network. Don't specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur. Examples: 1500 8192	Optional

CLIENT_NETWORK_TARGET

Value	Designation
Name of the host device that you will use for Client Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for GRID_NETWORK_TARGET or ADMIN_NETWORK_TARGET. Note: Don't use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair. Best practice: Specify a value even if this node will not initially have a Client Network IP address. Then you can add a Client Network IP address later, without having to reconfigure the node on the host. Examples: <pre>bond0.1003</pre> <pre>ens423</pre>	Best practice

CLIENT_NETWORK_TARGET_TYPE

Value	Designation
Interface (This is only supported value.)	Optional

CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Value	Designation
True or False Set the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Client Network. Best practice: In networks where promiscuous mode would be required, use the CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead. For more details on MAC cloning: • Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux) • Considerations and recommendations for MAC address cloning (Ubuntu or Debian)	Best practice

Grid Network keys

GRID_NETWORK_CONFIG

Value	Designation
STATIC or DHCP Defaults to STATIC if not specified.	Best practice

GRID_NETWORK_GATEWAY

Value	Designation
IPv4 address of the local Grid Network gateway for this node, which must be on the subnet defined by GRID_NETWORK_IP and GRID_NETWORK_MASK. This value is ignored for DHCP-configured networks. If the Grid Network is a single subnet with no gateway, use either the standard gateway address for the subnet (X.Y.Z.1) or this node's GRID_NETWORK_IP value; either value will simplify potential future Grid Network expansions.	Required

GRID_NETWORK_IP

Value	Designation
IPv4 address of this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 1.1.1.1 10.224.4.81	Required when GRID_NETWORK_CONFIG = STATIC Optional otherwise.

GRID_NETWORK_MAC

Value	Designation
The MAC address for the Grid Network interface in the container. Must be 6 pairs of hexadecimal digits separated by colons. Example: b2:9c:02:c2:27:30	Optional If omitted, a MAC address will be generated automatically.

GRID_NETWORK_MASK

Value	Designation
IPv4 netmask for this node on the Grid Network. Specify this key when GRID_NETWORK_CONFIG = STATIC; don't specify it for other values. Examples: 255.255.255.0 255.255.248.0	Required when GRID_NETWORK_IP is specified and GRID_NETWORK_CONFIG = STATIC. Optional otherwise.

GRID_NETWORK_MTU

Value	Designation
The maximum transmission unit (MTU) for this node on the Grid Network. Don't specify if GRID_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. IMPORTANT: The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur. IMPORTANT: For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The Grid Network MTU mismatch alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values don't have to be the same for all network types. Examples: 1500 8192	Optional

GRID_NETWORK_TARGET

Value	Designation
Name of the host device that you will use for Grid Network access by the StorageGRID node. Only network interface names are supported. Typically, you use a different interface name than what was specified for ADMIN_NETWORK_TARGET or CLIENT_NETWORK_TARGET. Note: Don't use bond or bridge devices as the network target. Either configure a VLAN (or other virtual interface) on top of the bond device, or use a bridge and virtual Ethernet (veth) pair. Examples: <pre>bond0.1001</pre> <pre>ens192</pre>	Required

GRID_NETWORK_TARGET_TYPE

Value	Designation
Interface (This is the only supported value.)	Optional

GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC

Value	Designation
True or False Set the value of the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Grid Network. Best practice: In networks where promiscuous mode would be required, use the GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead. For more details on MAC cloning: • Considerations and recommendations for MAC address cloning (Red Hat Enterprise Linux) • Considerations and recommendations for MAC address cloning (Ubuntu or Debian)	Best practice

Installation password key (temporary)

CUSTOM_TEMPORARY_PASSWORD_HASH

Value	Designation
For the primary Admin Node, set a default temporary password for the StorageGRID Installation API during installation. Note: Set an installation password on the primary Admin Node only. If you attempt to set a password on another node type, validation of the node configuration file will fail. Setting this value has no effect when installation has completed. If this key is omitted, by default no temporary password is set. Alternatively, you can set a temporary password using the StorageGRID Installation API. Must be a <code>crypt()</code> SHA-512 password hash with format <code>\$6\$<salt>\$<password hash></code> for a password of at least 8 and no more than 32 characters. This hash can be generated using CLI tools, such as the <code>openssl passwd</code> command in SHA-512 mode.	Best practice

Interfaces key

INTERFACE_TARGET_nnnn

Value	Designation
Name and optional description for an extra interface you want to add to this node. You can add multiple extra interfaces to each node. For <i>nnnn</i>, specify a unique number for each INTERFACE_TARGET entry you are adding. For the value, specify the name of the physical interface on the bare-metal host. Then, optionally, add a comma and provide a description of the interface, which is displayed on the VLAN interfaces page and the HA groups page. Example: <code>INTERFACE_TARGET_0001=ens256, Trunk</code> If you add a trunk interface, you must configure a VLAN interface in StorageGRID. If you add an access interface, you can add the interface directly to an HA group; you don't need to configure a VLAN interface.	Optional

Maximum RAM key

MAXIMUM_RAM

Value	Designation
The maximum amount of RAM that this node is allowed to consume. If this key is omitted, the node has no memory restrictions. When setting this field for a production-level node, specify a value that is at least 24 GB and 16 to 32 GB less than the total system RAM. Note: The RAM value affects a node's actual metadata reserved space. See the description of what Metadata Reserved Space is. The format for this field is <i>numberunit</i>, where <i>unit</i> can be b, k, m, or g. Examples: 24g 38654705664b Note: If you want to use this option, you must enable kernel support for memory cgroups.	Optional

Node type keys

NODE_TYPE

Value	Designation
Type of node: • VM_Admin_Node • VM_Storage_Node • VM_Archive_Node • VM_API_Gateway	Required

STORAGE_TYPE

Value	Designation
Defines the type of objects a Storage Node contains. For more information, see Types of Storage Nodes. This key is only required for nodes with <code>NODE_TYPE = VM_Storage_Node</code>; don't specify it for other node types. Storage types: • combined • data • metadata Note: If the <code>STORAGE_TYPE</code> is not specified, the Storage Node type is set to combined (data and metadata) by default.	Optional

Port remap keys

PORT_REMAP

Value	Designation
Remaps any port used by a node for internal grid node communications or external communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID, as described in Internal grid node communications or External communications. IMPORTANT: Don't remap the ports you are planning to use to configure load balancer endpoints. Note: If only <code>PORT_REMAP</code> is set, the mapping that you specify is used for both inbound and outbound communications. If <code>PORT_REMAP_INBOUND</code> is also specified, <code>PORT_REMAP</code> applies only to outbound communications. The format used is: <i>network type/protocol/default port used by grid node/new port</i>, where <i>network type</i> is grid, admin, or client, and <i>protocol</i> is tcp or udp. Example: <code>PORT_REMAP = client/tcp/18082/443</code> You can also remap multiple ports using a comma-separated list. Example: <code>PORT_REMAP = client/tcp/18082/443, client/tcp/18083/80</code>	Optional

PORT_REMAP_INBOUND

Value	Designation
Remaps inbound communications to the specified port. If you specify <code>PORT_REMAP_INBOUND</code> but don't specify a value for <code>PORT_REMAP</code>, outbound communications for the port are unchanged. IMPORTANT: Don't remap the ports you are planning to use to configure load balancer endpoints. The format used is: <i>network type/protocol/remapped port /default port used by grid node</i>, where <i>network type</i> is <code>grid</code>, <code>admin</code>, or <code>client</code>, and <i>protocol</i> is <code>tcp</code> or <code>udp</code>. Example: <code>PORT_REMAP_INBOUND = grid/tcp/3022/22</code> You can also remap multiple inbound ports using a comma-separated list. Example: <code>PORT_REMAP_INBOUND = grid/tcp/3022/22, admin/tcp/3022/22</code>	Optional

How grid nodes discover the primary Admin Node

Grid nodes communicate with the primary Admin Node for configuration and management. Each grid node must know the IP address of the primary Admin Node on the Grid Network.

To ensure that a grid node can access the primary Admin Node, you can do either of the following when deploying the node:

- You can use the `ADMIN_IP` parameter to enter the primary Admin Node's IP address manually.
- You can omit the `ADMIN_IP` parameter to have the grid node discover the value automatically. Automatic discovery is especially useful when the Grid Network uses DHCP to assign the IP address to the primary Admin Node.

Automatic discovery of the primary Admin Node is accomplished using a multicast domain name system (mDNS). When the primary Admin Node first starts up, it publishes its IP address using mDNS. Other nodes on the same subnet can then query for the IP address and acquire it automatically. However, because multicast IP traffic is not normally routable across subnets, nodes on other subnets can't acquire the primary Admin Node's IP address directly.

If you use automatic discovery:



- You must include the `ADMIN_IP` setting for at least one grid node on any subnets that the primary Admin Node is not directly attached to. This grid node will then publish the primary Admin Node's IP address for other nodes on the subnet to discover with mDNS.
- Ensure that your network infrastructure supports passing multi-cast IP traffic within a subnet.

Example node configuration files

You can use the example node configuration files to help set up the node configuration

files for your StorageGRID system. The examples show node configuration files for all types of grid nodes.

For most nodes, you can add Admin and Client Network addressing information (IP, mask, gateway, and so on) when you configure the grid using the Grid Manager or the Installation API. The exception is the primary Admin Node. If you want to browse to the Admin Network IP of the primary Admin Node to complete grid configuration (because the Grid Network is not routed, for example), you must configure the Admin Network connection for the primary Admin Node in its node configuration file. This is shown in the example.



In the examples, the Client Network target has been configured as a best practice, even though the Client Network is disabled by default.

Example for primary Admin Node

Example file name: /etc/storagegrid/nodes/dcl-adm1.conf

Example file contents:

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
TEMPORARY_PASSWORD_TYPE = Use custom password
CUSTOM_TEMPORARY_PASSWORD = Passw0rd
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-adm1-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dcl-adm1-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dcl-adm1-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.2
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_IP = 192.168.100.2
ADMIN_NETWORK_MASK = 255.255.248.0
ADMIN_NETWORK_GATEWAY = 192.168.100.1
ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0.0/21,172.17.0.0/21
```

Example for Storage Node

Example file name: /etc/storagegrid/nodes/dcl-sn1.conf

Example file contents:

```
NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/dcl-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/dcl-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/dcl-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/dcl-sn1-rangedb-3
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Example for Gateway Node

Example file name: /etc/storagegrid/nodes/dcl-gw1.conf

Example file contents:

```
NODE_TYPE = VM_API_Gateway
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-gw1-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.5
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Example for a non-primary Admin Node

Example file name: /etc/storagegrid/nodes/dcl-adm2.conf

Example file contents:

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Non-Primary
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adm2-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adm2-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dc1-adm2-tables
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.6
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

Validate the StorageGRID configuration

After creating configuration files in `/etc/storagegrid/nodes` for each of your StorageGRID nodes, you must validate the contents of those files.

To validate the contents of the configuration files, run the following command on each host:

```
sudo storagegrid node validate all
```

If the files are correct, the output shows **PASSED** for each configuration file, as shown in the example.



When using only one LUN on metadata-only nodes, you might receive a warning message that can be ignored.

```
Checking for misnamed node configuration files... PASSED
Checking configuration file for node dc1-adm1... PASSED
Checking configuration file for node dc1-gw1... PASSED
Checking configuration file for node dc1-sn1... PASSED
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-sn3... PASSED
Checking for duplication of unique values between nodes... PASSED
```



For an automated installation, you can suppress this output by using the `-q` or `--quiet` options in the `storagegrid` command (for example, `storagegrid --quiet...`). If you suppress the output, the command will have a non-zero exit value if any configuration warnings or errors were detected.

If the configuration files are incorrect, the issues are shown as **WARNING** and **ERROR**, as shown in the example. If any configuration errors are found, you must correct them before you continue with the installation.

```

Checking for misnamed node configuration files...
WARNING: ignoring /etc/storagegrid/nodes/dcl-adml
WARNING: ignoring /etc/storagegrid/nodes/dcl-sn2.conf.keep
WARNING: ignoring /etc/storagegrid/nodes/my-file.txt
Checking configuration file for node dcl-adml...
ERROR: NODE_TYPE = VM_Foo_Node
      VM_Foo_Node is not a valid node type.  See *.conf.sample
ERROR: ADMIN_ROLE = Foo
      Foo is not a valid admin role.  See *.conf.sample
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-gw1-var-local
      /dev/mapper/sgws-gw1-var-local is not a valid block device
Checking configuration file for node dcl-gw1...
ERROR: GRID_NETWORK_TARGET = bond0.1001
      bond0.1001 is not a valid interface.  See `ip link show`
ERROR: GRID_NETWORK_IP = 10.1.3
      10.1.3 is not a valid IPv4 address
ERROR: GRID_NETWORK_MASK = 255.248.255.0
      255.248.255.0 is not a valid IPv4 subnet mask
Checking configuration file for node dcl-sn1...
ERROR: GRID_NETWORK_GATEWAY = 10.2.0.1
      10.2.0.1 is not on the local subnet
ERROR: ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0foo
      Could not parse subnet list
Checking configuration file for node dcl-sn2... PASSED
Checking configuration file for node dcl-sn3... PASSED
Checking for duplication of unique values between nodes...
ERROR: GRID_NETWORK_IP = 10.1.0.4
      dcl-sn2 and dcl-sn3 have the same GRID_NETWORK_IP
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn2-var-local
      dcl-sn2 and dcl-sn3 have the same BLOCK_DEVICE_VAR_LOCAL
ERROR: BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn2-rangedb-0
      dcl-sn2 and dcl-sn3 have the same BLOCK_DEVICE_RANGEDB_00

```

Start the StorageGRID host service

To start your StorageGRID nodes, and ensure they restart after a host reboot, you must enable and start the StorageGRID host service.

Steps

1. Run the following commands on each host:

```

sudo systemctl enable storagegrid
sudo systemctl start storagegrid

```

2. Run the following command to ensure the deployment is proceeding:

```
sudo storagegrid node status node-name
```

3. If any node returns a status of "Not Running" or "Stopped," run the following command:

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

4. If you have previously enabled and started the StorageGRID host service (or if you are unsure if the service has been enabled and started), also run the following command:

```
sudo systemctl reload-or-restart storagegrid
```

Configure grid and complete installation (Ubuntu or Debian)

Navigate to the Grid Manager

You use the Grid Manager to define all of the information required to configure your StorageGRID system.

Before you begin

The primary Admin Node must be deployed and have completed the initial startup sequence.

Steps

1. Open your web browser and navigate to:

```
https://primary_admin_node_ip
```

Alternatively, you can access the Grid Manager on port 8443:


```
https://primary_admin_node_ip:8443
```

You can use the IP address for the primary Admin Node IP on the Grid Network or on the Admin Network, as appropriate for your network configuration.

2. Manage a temporary installer password as needed:
 - If a password has already been set using one of these methods, enter the password to proceed.
 - A user set the password while accessing the installer previously
 - The password was automatically imported from the node config file at `/etc/storagegrid/nodes/<node_name>.conf`
 - If a password has not been set, optionally set a password to secure the StorageGRID installer.
3. Select **Install a StorageGRID system**.

The page used to configure a StorageGRID system appears.

Install



License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

Browse

Specify the StorageGRID license information

You must specify the name for your StorageGRID system and upload the license file provided by NetApp.

Steps

1. On the License page, enter a meaningful name for your StorageGRID system in the **Grid Name** field.

After installation, the name is displayed at the top of the Nodes menu.

2. Select **Browse**, locate the NetApp license file (*NLF-unique-id.txt*), and select **Open**.

The license file is validated, and the serial number is displayed.



The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product. You can update to a license that offers support after installation.

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File NLF-959007-Internal.txt

License Serial Number

3. Select **Next**.

Add sites

You must create at least one site when you are installing StorageGRID. You can create additional sites to increase the reliability and storage capacity of your StorageGRID system.

Steps

1. On the Sites page, enter the **Site Name**.
2. To add additional sites, click the plus sign next to the last site entry and enter the name in the new **Site Name** text box.

Add as many additional sites as required for your grid topology. You can add up to 16 sites.

NetApp® StorageGRID®

Help ▾

Install

1

License

2

Sites

3

Grid Network

4

Grid Nodes

5

NTP

6

DNS

7

Passwords

8

Summary

Sites

In a single-site deployment, infrastructure and operations are centralized in one site.

In a multi-site deployment, infrastructure can be distributed asymmetrically across sites, and proportional to the needs of each site. Typically, sites are located in geographically different locations. Having multiple sites also allows the use of distributed replication and erasure coding for increased availability and resiliency.

Site Name 1

Raleigh

✕

Site Name 2

Atlanta

+ ✕

3. Click **Next**.

Specify Grid Network subnets

You must specify the subnets that are used on the Grid Network.

About this task

The subnet entries include the subnets for the Grid Network for each site in your StorageGRID system, along with any subnets that need to be reachable through the Grid Network.

If you have multiple grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway.

Steps

1. Specify the CIDR network address for at least one Grid Network in the **Subnet 1** text box.
2. Click the plus sign next to the last entry to add an additional network entry. You must specify all subnets for all sites in the Grid Network.
 - If you have already deployed at least one node, click **Discover Grid Networks Subnets** to automatically populate the Grid Network Subnet List with the subnets reported by grid nodes that have registered with the Grid Manager.

- You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

NetApp® StorageGRID® Help ▾

Install

1 License 2 Sites 3 **Grid Network** 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Grid Network

You must specify the subnets that are used on the Grid Network. These entries typically include the subnets for the Grid Network for each site in your StorageGRID system. Select Discover Grid Networks to automatically add subnets based on the network configuration of all registered nodes.

Note: You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

Subnet 1 +

3. Click **Next**.

Approve pending grid nodes

You must approve each grid node before it can join the StorageGRID system.

Before you begin

You have deployed all virtual and StorageGRID appliance grid nodes.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

Steps

1. Review the Pending Nodes list, and confirm that it shows all of the grid nodes you deployed.



If a grid node is missing, confirm that it was deployed successfully and has the correct Grid Network IP of the primary admin node set for ADMIN_IP.

2. Select the radio button next to a pending node you want to approve.



Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

+ Approve ✕ Remove		Search Q				
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address	
☒	50:6b:4b:42:d7:00	NetApp-SGA	Storage Node	StorageGRID Appliance	172.16.5.20/21	
						◀ ▶

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

✎ Edit ↺ Reset ✕ Remove		Search Q				
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
☐	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
☐	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
☐	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
☐	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
☐	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21
						◀ ▶

3. Click **Approve**.

4. In General Settings, modify settings for the following properties, as necessary:

- **Site:** The system name of the site for this grid node.
- **Name:** The system name for the node. The name defaults to the name you specified when you configured the node.

System names are required for internal StorageGRID operations and can't be changed after you complete the installation. However, during this step of the installation process, you can change system names as required.

- **NTP Role:** The Network Time Protocol (NTP) role of the grid node. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the Primary role to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

- **Storage Type** (Storage Nodes only): Specify that a new Storage Node be used exclusively for data only, metadata only, or both. The options are **Data and metadata** ("combined"), **Data only**, and **Metadata only**.



See [Types of Storage Nodes](#) for information about requirements for these node types.

- **ADC service** (Storage Nodes only): Select **Automatic** to let the system determine whether the node requires the Administrative Domain Controller (ADC) service. The ADC service keeps track of the location and availability of grid services. At least three Storage Nodes at each site must include the ADC service. You can't add the ADC service to a node after it is deployed.

5. In Grid Network, modify settings for the following properties as necessary:

- **IPv4 Address (CIDR)**: The CIDR network address for the Grid Network interface (eth0 inside the container). For example: 192.168.1.234/21
- **Gateway**: The Grid Network gateway. For example: 192.168.0.1

The gateway is required if there are multiple grid subnets.



If you selected DHCP for the Grid Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

6. If you want to configure the Admin Network for the grid node, add or update the settings in the Admin Network section as necessary.

Enter the destination subnets of the routes out of this interface in the **Subnets (CIDR)** text box. If there are multiple Admin subnets, the Admin gateway is required.



If you selected DHCP for the Admin Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Admin Network was not configured during the initial installation using the StorageGRID Appliance Installer, it can't be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

Rebooting can take several minutes.

- Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- Select **Configure Networking > IP Configuration** and configure the enabled networks.
- Return to the Home page and click **Start Installation**.
- In the Grid Manager: If the node is listed in the Approved Nodes table, remove the node.
- Remove the node from the Pending Nodes table.

- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page of the Appliance Installer.

For additional information, see the [Quick start for hardware installation](#) to locate instructions for your appliance.

7. If you want to configure the Client Network for the grid node, add or update the settings in the Client Network section as necessary. If the Client Network is configured, the gateway is required, and it becomes the default gateway for the node after installation.



If you selected DHCP for the Client Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Client Network was not configured during the initial installation using the StorageGRID Appliance Installer, it can't be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

Rebooting can take several minutes.

- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, remove the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page of the Appliance Installer.

To learn how to install StorageGRID appliances, see the [Quick start for hardware installation](#) to locate instructions for your appliance.

8. Click **Save**.

The grid node entry moves to the Approved Nodes list.



Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

+ Approve

✖ Remove

Search

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
No results found.				

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

Edit

Reset

✖ Remove

Search

	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
☐	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
☐	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
☐	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
☐	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
☐	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21
☐	50:6b:4b:42:d7:00	NetApp-SGA	Raleigh	Storage Node	StorageGRID Appliance	172.16.5.20/21

9. Repeat these steps for each pending grid node you want to approve.

You must approve all nodes that you want in the grid. However, you can return to this page at any time before you click **Install** on the Summary page. You can modify the properties of an approved grid node by selecting its radio button and clicking **Edit**.

10. When you are done approving grid nodes, click **Next**.

Specify Network Time Protocol server information

You must specify the Network Time Protocol (NTP) configuration information for the StorageGRID system, so that operations performed on separate servers can be kept synchronized.

About this task

You must specify IPv4 addresses for the NTP servers.

You must specify external NTP servers. The specified NTP servers must use the NTP protocol.

You must specify four NTP server references of Stratum 3 or better to prevent issues with time drift.



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

[Support boundary to configure the Windows Time service for high-accuracy environments](#)

The external NTP servers are used by the nodes to which you previously assigned Primary NTP roles.



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

Steps

1. Specify the IPv4 addresses for at least four NTP servers in the **Server 1** to **Server 4** text boxes.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

The screenshot shows the NetApp StorageGRID installation wizard. At the top, there's a blue header with "NetApp® StorageGRID®" and a "Help" link. Below the header is a progress bar with eight steps: 1. License, 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP (highlighted in blue), 6. DNS, 7. Passwords, and 8. Summary. Below the progress bar, the "Network Time Protocol" section is active. It contains the instruction: "Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync." There are four input fields labeled "Server 1" through "Server 4". Server 1 contains "10.60.248.183", Server 2 contains "10.227.204.142", Server 3 contains "10.235.48.111", and Server 4 contains "0.0.0.0". To the right of the Server 4 field is a plus sign (+) to add more servers.

3. Select **Next**.

Related information

[Networking guidelines](#)

Specify DNS server information

You must specify DNS information for your StorageGRID system, so that you can access external servers using hostnames instead of IP addresses.

About this task

Specifying [DNS server information](#) allows you to use Fully Qualified Domain Name (FQDN) hostnames rather than IP addresses for email notifications and AutoSupport.

To ensure proper operation, specify two or three DNS servers. If you specify more than three, it is possible that only three will be used because of known OS limitations on some platforms. If you have routing restrictions in your environment, you can [customize the DNS server list](#) for individual nodes (typically all nodes at a site) to use a different set of up to three DNS servers.

If possible, use DNS servers that each site can access locally to ensure that an islanded site can resolve the FQDNs for external destinations.

Steps

1. Specify the IPv4 address for at least one DNS server in the **Server 1** text box.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

The screenshot shows the NetApp StorageGRID installation wizard. At the top is a blue header with "NetApp® StorageGRID®" and a "Help" link. Below the header is a progress bar with eight steps: 1. License, 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP, 6. DNS (highlighted in blue), 7. Passwords, and 8. Summary. Below the progress bar is the "Domain Name Service" section. It contains a text box for "Server 1" with the value "10.224.223.130" and a red "X" icon. Below that is a text box for "Server 2" with the value "10.224.223.136" and a red "+ X" icon. The text below the text boxes reads: "Enter the IP address for at least one Domain Name System (DNS) server, so that server hostnames can be used instead of IP addresses. Specifying at least two DNS servers is recommended. Configuring DNS enables server connectivity, email notifications, and NetApp AutoSupport."

The best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

Specify the StorageGRID system passwords

As part of installing your StorageGRID system, you need to enter the passwords to use to secure your system and perform maintenance tasks.

About this task

Use the Install passwords page to specify the provisioning passphrase and the grid management root user password.

- The provisioning passphrase is used as an encryption key and is not stored by the StorageGRID system.
- You must have the provisioning passphrase for installation, expansion, and maintenance procedures, including downloading the Recovery Package. Therefore, it is important that you store the provisioning passphrase in a secure location.
- You can change the provisioning passphrase from the Grid Manager if you have the current one.
- The grid management root user password can be changed using the Grid Manager.

- Randomly generated command line console and SSH passwords are stored in the `Passwords.txt` file in the Recovery Package.

Steps

1. In **Provisioning Passphrase**, enter the provisioning passphrase that will be required to make changes to the grid topology of your StorageGRID system.

Store the provisioning passphrase in a secure place.



If after the installation completes and you want to change the provisioning passphrase later, you can use the Grid Manager. Select **CONFIGURATION > Access control > Grid passwords**.

2. In **Confirm Provisioning Passphrase**, reenter the provisioning passphrase to confirm it.
3. In **Grid Management Root User Password**, enter the password to use to access the Grid Manager as the "root" user.

Store the password in a secure place.

4. In **Confirm Root User Password**, reenter the Grid Manager password to confirm it.

NetApp® StorageGRID®
Help

Install

1 License
2 Sites
3 Grid Network
4 Grid Nodes
5 NTP
6 DNS
7 Passwords
8 Summary

Passwords

Enter secure passwords that meet your organization's security policies. A text file containing the command line passwords must be downloaded during the final installation step.

Provisioning Passphrase	
Confirm Provisioning Passphrase	
Grid Management Root User Password	
Confirm Root User Password	

☒ Create random command line passwords.

5. If you are installing a grid for proof of concept or demo purposes, optionally clear the **Create random command line passwords** checkbox.

For production deployments, random passwords should always be used for security reasons. Clear **Create random command line passwords** only for demo grids if you want to use default passwords to access grid nodes from the command line using the "root" or "admin" account.



You are prompted to download the Recovery Package file (`sgws-recovery-package-id-revision.zip`) after you click **Install** on the Summary page. You must [download this file](#) to complete the installation. The passwords required to access the system are stored in the `Passwords.txt` file, contained in the Recovery Package file.

6. Click **Next**.

Review your configuration and complete installation

You must carefully review the configuration information you have entered to ensure that the installation completes successfully.

Steps

1. View the **Summary** page.

NetApp® StorageGRID® Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 **Summary**

Summary

Verify that all of the grid configuration information is correct, and then click Install. You can view the status of each grid node as it installs. Click the Modify links to go back and change the associated information.

General Settings

Grid Name	Grid1	Modify License
Passwords	Auto-generated random command line passwords	Modify Passwords

Networking

NTP	10.60.248.183 10.227.204.142 10.235.48.111	Modify NTP
DNS	10.224.223.130 10.224.223.136	Modify DNS
Grid Network	172.16.0.0/21	Modify Grid Network

Topology

Topology	Atlanta	Modify Sites	Modify Grid Nodes
	Raleigh		
	dc1-adm1 dc1-g1 dc1-s1 dc1-s2 dc1-s3 NetApp-SGA		

2. Verify that all of the grid configuration information is correct. Use the Modify links on the Summary page to go back and correct any errors.

3. Click **Install**.



If a node is configured to use the Client Network, the default gateway for that node switches from the Grid Network to the Client Network when you click **Install**. If you lose connectivity, you must ensure that you are accessing the primary Admin Node through an accessible subnet. See [Networking guidelines](#) for details.

4. Click **Download Recovery Package**.

When the installation progresses to the point where the grid topology is defined, you are prompted to download the Recovery Package file (.zip), and confirm that you can successfully access the contents of this file. You must download the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fail. The installation continues in the background, but you can't complete the installation and access the StorageGRID system until you download and verify this file.

5. Verify that you can extract the contents of the .zip file, and then save it in two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

6. Select the **I have successfully downloaded and verified the Recovery Package file** checkbox, and click **Next**.

If the installation is still in progress, the status page appears. This page indicates the progress of the installation for each grid node.

Installation Status

If necessary, you may [Download the Recovery Package file](#) again.

Search

Name	IT	Site	IT	Grid Network IPv4 Address	Progress	IT	Stage	IT
dc1-adm1		Site1		172.16.4.215/21			Starting services	
dc1-g1		Site1		172.16.4.216/21			Complete	
dc1-s1		Site1		172.16.4.217/21			Waiting for Dynamic IP Service peers	
dc1-s2		Site1		172.16.4.218/21			Downloading hotfix from primary Admin if needed	
dc1-s3		Site1		172.16.4.219/21			Downloading hotfix from primary Admin if needed	

When the Complete stage is reached for all grid nodes, the sign-in page for the Grid Manager appears.

7. Sign in to the Grid Manager using the "root" user and the password you specified during the installation.

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You can't set up DHCP during configuration.



Nodes reboot when the Grid Network configuration is changed by DHCP, which can cause outages if a DHCP change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to reestablish these connections.

Installation REST API

StorageGRID provides the StorageGRID Installation API for performing installation tasks.

The API uses the Swagger open source API platform to provide the API documentation. Swagger allows both developers and non-developers to interact with the API in a user interface that illustrates how the API responds to parameters and options. This documentation assumes that you are familiar with standard web technologies and the JSON data format.



Any API operations you perform using the API Documentation webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Each REST API command includes the API's URL, an HTTP action, any required or optional URL parameters, and an expected API response.

StorageGRID Installation API

The StorageGRID Installation API is only available when you are initially configuring your StorageGRID system, and if you need to perform a primary Admin Node recovery. The Installation API can be accessed over HTTPS from the Grid Manager.

To access the API documentation, go to the installation web page on the primary Admin Node and select **Help > API documentation** from the menu bar.

The StorageGRID Installation API includes the following sections:

- **config** — Operations related to the product release and versions of the API. You can list the product release version and the major versions of the API supported by that release.
- **grid** — Grid-level configuration operations. You can get and update grid settings, including grid details, Grid Network subnets, grid passwords, and NTP and DNS server IP addresses.
- **nodes** — Node-level configuration operations. You can retrieve a list of grid nodes, delete a grid node, configure a grid node, view a grid node, and reset a grid node's configuration.
- **provision** — Provisioning operations. You can start the provisioning operation and view the status of the provisioning operation.
- **recovery** — Primary Admin Node recovery operations. You can reset information, upload the Recover Package, start the recovery, and view the status of the recovery operation.
- **recovery-package** — Operations to download the Recovery Package.
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.
- **temporary-password** — Operations on the temporary password to secure the mgmt-api during installation.

Related information

[Automating the installation](#)

Where to go next

After completing an installation, perform the required integration and configuration tasks. You can perform the optional tasks as needed.

Required tasks

- [Create a tenant account](#) for the S3 client protocol that will be used to store objects on your StorageGRID system.
- [Control system access](#) by configuring groups and user accounts. Optionally, you can [configure a federated identity source](#) (such as Active Directory or OpenLDAP), so you can import administration groups and users. Or, you can [create local groups and users](#).
- Integrate and test the [S3 API](#) client applications you will use to upload objects to your StorageGRID system.
- [Configure the information lifecycle management \(ILM\) rules and ILM policy](#) you want to use to protect object data.
- If your installation includes appliance Storage Nodes, use SANtricity OS to complete the following tasks:
 - Connect to each StorageGRID appliance.
 - Verify receipt of AutoSupport data.

See [Set up hardware](#).

- Review and follow the [StorageGRID system hardening guidelines](#) to eliminate security risks.
- [Configure email notifications for system alerts](#).

Optional tasks

- [Update grid node IP addresses](#) if they have changed since you planned your deployment and generated the Recovery Package.
- [Configure storage encryption](#), if required.
- [Configure storage compression](#) to reduce the size of stored objects, if required.
- [Configure VLAN interfaces](#) to isolate and partition network traffic, if required.
- [Configure high availability groups](#) to improve connection availability for the Grid Manager, Tenant Manager, and S3 clients, if required.
- [Configure load balancer endpoints](#) for S3 client connectivity, if required.

Troubleshoot installation issues

If any problems occur while installing your StorageGRID system, you can access the installation log files. Technical support might also need to use the installation log files to resolve issues.

The following installation log files are available from the container that is running each node:

- `/var/local/log/install.log` (found on all grid nodes)
- `/var/local/log/gdu-server.log` (found on the primary Admin Node)

The following installation log files are available from the host:

- `/var/log/storagegrid/daemon.log`
- `/var/log/storagegrid/nodes/<node-name>.log`

To learn how to access the log files, see [Collect log files and system data](#).

Related information

[Troubleshoot a StorageGRID system](#)

Example /etc/network/interfaces

The `/etc/network/interfaces` file includes three sections, which define the physical interfaces, bond interface, and VLAN interfaces. You can combine the three example sections into a single file, which will aggregate four Linux physical interfaces into a single LACP bond and then establish three VLAN interfaces subtending the bond for use as StorageGRID Grid, Admin, and Client Network interfaces.

Physical interfaces

Note that the switches at the other ends of the links must also treat the four ports as a single LACP trunk or port channel, and must pass at least the three referenced VLANs with tags.

```
# loopback interface
auto lo
iface lo inet loopback

# ens160 interface
auto ens160
iface ens160 inet manual
    bond-master bond0
    bond-primary en160

# ens192 interface
auto ens192
iface ens192 inet manual
    bond-master bond0

# ens224 interface
auto ens224
iface ens224 inet manual
    bond-master bond0

# ens256 interface
auto ens256
iface ens256 inet manual
    bond-master bond0
```

Bond interface

```
# bond0 interface
auto bond0
iface bond0 inet manual
    bond-mode 4
    bond-miimon 100
    bond-slaves ens160 ens192 end224 ens256
```

VLAN interfaces

```
# 1001 vlan
auto bond0.1001
iface bond0.1001 inet manual
vlan-raw-device bond0

# 1002 vlan
auto bond0.1002
iface bond0.1002 inet manual
vlan-raw-device bond0

# 1003 vlan
auto bond0.1003
iface bond0.1003 inet manual
vlan-raw-device bond0
```

Install StorageGRID on VMware

Quick start for installing StorageGRID on VMware

Follow these high-level steps to install a VMware StorageGRID node.

1

Preparation

- Learn about [StorageGRID architecture and network topology](#).
- Learn about the specifics of [StorageGRID networking](#).
- Gather and prepare the [Required information and materials](#).
- Install and configure [VMware vSphere Hypervisor, vCenter, and the ESX hosts](#).
- Prepare the required [CPU and RAM](#).
- Provide for [storage and performance requirements](#).

2

Deployment

Deploy grid nodes. When you deploy grid nodes, they are created as part of the StorageGRID system and

connected to one or more networks.

- Use the VMware vSphere Web Client, a .vmdk file, and a set of .ovf file templates to [deploy the software-based nodes as virtual machines \(VMs\)](#) on the servers you prepared in step 1.
- To deploy StorageGRID appliance nodes, follow the [Quick start for hardware installation](#).

3

Configuration

When all nodes have been deployed, use the Grid Manager to [configure the grid and complete the installation](#).

Automate the installation

To save time and provide consistency, you can automate the deployment and configuration of grid nodes and the configuration of the StorageGRID system.

- [Automate grid node deployment using VMware vSphere](#).
- After you deploy grid nodes, [automate the configuration of the StorageGRID system](#) using the Python configuration script provided in the installation archive.
- [Automate the installation and configuration of appliance grid nodes](#)
- If you are an advanced developer of StorageGRID deployments, automate the installation of grid nodes by using the [installation REST API](#).

Plan and prepare for installation on VMware

Required information and materials

Before you install StorageGRID, gather and prepare the required information and materials.

Required information

Network plan

Which networks you intend to attach to each StorageGRID node. StorageGRID supports multiple networks for traffic separation, security, and administrative convenience.

See the StorageGRID [Networking guidelines](#).

Network information

IP addresses to assign to each grid node and the IP addresses of the DNS and NTP servers.

Servers for grid nodes

Identify a set of servers (physical, virtual, or both) that, in aggregate, provide sufficient resources to support the number and type of StorageGRID nodes you plan to deploy.



If your StorageGRID installation will not use StorageGRID appliance (hardware) Storage Nodes, you must use hardware RAID storage with battery-backed write cache (BBWC). StorageGRID does not support the use of virtual storage area networks (vSANs), software RAID, or no RAID protection.

Related information

[NetApp Interoperability Matrix Tool](#)

Required materials

NetApp StorageGRID license

You must have a valid, digitally signed NetApp license.



A non-production license, which can be used for testing and proof of concept grids, is included in the StorageGRID installation archive.

StorageGRID installation archive

[Download the StorageGRID installation archive and extract the files.](#)

Service laptop

The StorageGRID system is installed through a service laptop.

The service laptop must have:

- Network port
- SSH client (for example, PuTTY)
- [Supported web browser](#)

StorageGRID documentation

- [Release notes](#)
- [Instructions for administering StorageGRID](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archives and extract the files. Optionally, you can manually verify the files in the installation package.

Steps

1. Go to the [NetApp Downloads page for StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.
3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead statement appears, read it and select the checkbox.



You must apply any required hotfixes after you install the StorageGRID release. For more information, see the [hotfix procedure in the recovery and maintenance instructions](#)

5. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column, select the .tgz or .zip installation archive for VMware.



Use the .zip file if you are running Windows on the service laptop.

7. Save the installation archive.

8. If you need to verify the installation archive:
 - a. Download the StorageGRID code signature verification package. The file name for this package uses the format `StorageGRID_<version-number>_Code_Signature_Verification_Package.tar.gz`, where `<version-number>` is the StorageGRID software version.
 - b. Follow the steps to [manually verify the installation files](#).
9. Extract the files from the installation archive.
10. Choose the files you need.

The files you need depend on your planned grid topology and how you will deploy your StorageGRID system.



The paths listed in the table are relative to the top-level directory installed by the extracted installation archive.

Path and file name	Description
<code>./vsphere/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./vsphere/NLF000000.txt</code>	A free license that does not provide any support entitlement for the product.
<code>./vsphere/NetApp-SG-version-SHA.vmdk</code>	The virtual machine disk file that is used as a template for creating grid node virtual machines.
<code>./vsphere/vsphere-primary-admin.ovf</code> <code>./vsphere/vsphere-primary-admin.mf</code>	The Open Virtualization Format template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying the primary Admin Node.
<code>./vsphere/vsphere-non-primary-admin.ovf</code> <code>./vsphere/vsphere-non-primary-admin.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying non-primary Admin Nodes.
<code>./vsphere/vsphere-gateway.ovf</code> <code>./vsphere/vsphere-gateway.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying Gateway Nodes.
<code>./vsphere/vsphere-storage.ovf</code> <code>./vsphere/vsphere-storage.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying virtual machine-based Storage Nodes.
Deployment scripting tool	Description
<code>./vsphere/deploy-vsphere-ovftool.sh</code>	A Bash shell script used to automate the deployment of virtual grid nodes.
<code>./vsphere/deploy-vsphere-ovftool-sample.ini</code>	An example configuration file for use with the <code>deploy-vsphere-ovftool.sh</code> script.

Path and file name	Description
<code>./vsphere/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.
<code>./vsphere/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./vsphere/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled. You can also use this script for Ping Federate integration.
<code>./vsphere/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./vsphere/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./vsphere/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
<code>./vsphere/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.
<code>./vsphere/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

Manually verify installation files (optional)

If necessary, you can manually verify the files in the StorageGRID installation archive.

Before you begin

You have [downloaded the verification package](#) from the [NetApp Downloads page for StorageGRID](#).

Steps

1. Extract the artifacts from the verification package:

```
tar -xf StorageGRID_11.9.0_Code_Signature_Verification_Package.tar.gz
```

2. Ensure that these artifacts were extracted:

- Leaf certificate: Leaf-Cert.pem
- Certificate chain: CA-Int-Cert.pem
- Time stamp response chain: TS-Cert.pem
- Checksum file: sha256sum
- Checksum signature: sha256sum.sig
- Time stamp response file: sha256sum.sig.tsr

3. Use the chain to verify the leaf certificate is valid.

Example: `openssl verify -CAfile CA-Int-Cert.pem Leaf-Cert.pem`

Expected output: Leaf-Cert.pem: OK

4. If step 2 failed because of an expired leaf certificate, use the `tsr` file to verify.

Example: `openssl ts -CAfile CA-Int-Cert.pem -untrusted TS-Cert.pem -verify -data sha256sum.sig -in sha256sum.sig.tsr`

Expected output includes: Verification: OK

5. Create a public key file from the leaf certificate.

Example: `openssl x509 -pubkey -noout -in Leaf-Cert.pem > Leaf-Cert.pub`

Expected output: *none*

6. Use the public key to verify the `sha256sum` file against `sha256sum.sig`.

Example: `openssl dgst -sha256 -verify Leaf-Cert.pub -signature sha256sum.sig sha256sum`

Expected output: Verified OK

7. Verify the `sha256sum` file content against newly created checksums.

Example: `sha256sum -c sha256sum`

Expected output: `<filename>: OK`
`<filename>` is the name of the archive file you downloaded.

8. [Complete the remaining steps](#) to extract and choose the appropriate installation files.

Software requirements for VMware

You can use a virtual machine to host any type of StorageGRID node. You need one virtual machine for each grid node.

VMware vSphere Hypervisor

You must install VMware vSphere Hypervisor on a prepared physical server. The hardware must be configured

correctly (including firmware versions and BIOS settings) before you install VMware software.

- Configure networking in the hypervisor as required to support networking for the StorageGRID system you are installing.

Networking guidelines

- Ensure that the datastore is large enough for the virtual machines and virtual disks that are required to host the grid nodes.
- If you create more than one datastore, name each datastore so that you can easily identify which datastore to use for each grid node when you create virtual machines.

ESX host configuration requirements



You must properly configure the network time protocol (NTP) on each ESX host. If the host time is incorrect, negative effects, including data loss, could occur.

VMware configuration requirements

You must install and configure VMware vSphere and vCenter before deploying StorageGRID nodes.

For supported versions of VMware vSphere Hypervisor and VMware vCenter Server software, see the [NetApp Interoperability Matrix Tool](#).

For the steps required to install these VMware products, see the VMware documentation.

CPU and RAM requirements

Before installing StorageGRID software, verify and configure the hardware so that it is ready to support the StorageGRID system.

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: Dependent on the total RAM available and the amount of non-StorageGRID software running on the system
 - Generally, at least 24 GB per node, and 2 to 16 GB less than the total system RAM
 - A minimum of 64 GB for each tenant that will have approximately 5,000 buckets

Software-based metadata-only node resources must match the existing Storage Nodes resources. For example:

- If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)
- If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0)).

When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.

VMware supports one node per virtual machine. Ensure that the StorageGRID node does not exceed the physical RAM available. Each virtual machine must be dedicated to running StorageGRID.



Monitor your CPU and memory usage regularly to ensure that these resources continue to accommodate your workload. For example, doubling the RAM and CPU allocation for virtual Storage Nodes would provide similar resources to those provided for StorageGRID appliance nodes. Additionally, if the amount of metadata per node exceeds 500 GB, consider increasing the RAM per node to 48 GB or more. For information about managing object metadata storage, increasing the Metadata Reserved Space setting, and monitoring CPU and memory usage, see the instructions for [administering](#), [monitoring](#), and [upgrading](#) StorageGRID.

If hyperthreading is enabled on the underlying physical hosts, you can provide 8 virtual cores (4 physical cores) per node. If hyperthreading is not enabled on the underlying physical hosts, you must provide 8 physical cores per node.

If you are using virtual machines as hosts and have control over the size and number of VMs, you should use a single VM for each StorageGRID node and size the VM accordingly.

See also [Storage and performance requirements](#).

Storage and performance requirements

You must understand the storage and performance requirements for StorageGRID nodes hosted by virtual machines, so you can provide enough space to support the initial configuration and future storage expansion.

Performance requirements

The performance of the OS volume and of the first storage volume significantly impacts the overall performance of the system. Ensure that these provide adequate disk performance in terms of latency, input/output operations per second (IOPS), and throughput.

All StorageGRID nodes require that the OS drive and all storage volumes have write-back caching enabled. The cache must be on a protected or persistent media.

Requirements for virtual machines that use NetApp ONTAP storage

If you are deploying a StorageGRID node as a virtual machine with storage assigned from a NetApp ONTAP system, you have confirmed that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running as an virtual machine on a VMware host, ensure the volume backing the datastore for the node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

Number of virtual machines required

Each StorageGRID site requires a minimum of three Storage Nodes.

Storage requirements by node type

In a production environment, the virtual machines for StorageGRID nodes must meet different requirements, depending on the types of nodes.



Disk snapshots can't be used to restore grid nodes. Instead, refer to the [grid node recovery](#) procedures for each type of node.

Node Type	Storage
Admin Node	100 GB LUN for OS 200 GB LUN for Admin Node tables 200 GB LUN for Admin Node audit log
Storage Node	100 GB LUN for OS 3 LUNs for each Storage Node on this host Note: A Storage Node can have 1 to 16 storage LUNs; at least 3 storage LUNs are recommended. Minimum size per LUN: 4 TB Maximum tested LUN size: 39 TB.
Storage Node (metadata-only)	100 GB LUN for OS 1 LUN Minimum size per LUN: 4 TB Maximum tested LUN size: 39 TB. Note: Only one rangedb is required for metadata-only Storage Nodes.
Gateway Node	100 GB LUN for OS



Depending on the audit level configured, the size of user inputs such as S3 object key name, and how much audit log data you need to preserve, you might need to increase the size of the audit log LUN on each Admin Node. Generally, a grid generates approximately 1 KB of audit data per S3 operation, which would mean that a 200 GB LUN would support 70 million operations per day or 800 operations per second for two to three days.


Storage requirements for Storage Nodes

A software-based Storage Node can have 1 to 16 storage volumes; 3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.



An appliance Storage Node can also have up to 48 storage volumes.

As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. Any remaining space on storage volume 0 and any other storage volumes in the Storage Node are used exclusively for 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. The three copies of object metadata are evenly distributed across all Storage Nodes at each site.

When installing a grid with metadata-only Storage Nodes, the grid must also contain a minimum number of nodes for object storage. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

- For a single-site grid, at least two Storage Nodes are configured for objects and metadata.
- For a multi-site grid, at least one Storage Node per site are configured for objects and metadata.

When you assign space to volume 0 of a new Storage Node, you must ensure there is adequate space for that node's portion of all object metadata.

- At a minimum, you must assign at least 4 TB to volume 0.



If you use only one storage volume for a Storage Node and you assign 4 TB or less to the volume, the Storage Node might enter the storage read-only state on startup and store object metadata only.



If you assign less than 500 GB to volume 0 (non-production use only), 10% of the storage volume's capacity is reserved for metadata.

- Software-based metadata-only node resources must match the existing Storage Nodes resources. For example:

- If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)
- If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0)).

When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.

- If you are installing a new system (StorageGRID 11.6 or higher) and each Storage Node has 128 GB or more of RAM, assign 8 TB or more to volume 0. Using a larger value for volume 0 can increase the space allowed for metadata on each Storage Node.
- When configuring different Storage Nodes for a site, use the same setting for volume 0 if possible. If a site contains Storage Nodes of different sizes, the Storage Node with the smallest volume 0 will determine the metadata capacity of that site.

For details, go to [Manage object metadata storage](#).

Automate the installation (VMware)

You can use the VMware OVF Tool to automate the deployment of grid nodes. You can also automate the configuration of StorageGRID.

Automate grid node deployment

Use the VMware OVF Tool to automate the deployment of grid nodes.

Before you begin

- You have access to a Linux/Unix system with Bash 3.2 or later.
- You have VMware vSphere with vCenter
- You have VMware OVF Tool 4.1 installed and correctly configured.
- You know the username and password to access VMware vSphere using the OVF Tool
- You have the sufficient permissions to deploy VMs from OVF files and power them on, and permissions to create additional volumes to attach to the VMs. See the `ovftool` documentation for details.
- You know the virtual infrastructure (VI) URL for the location in vSphere where you want to deploy the StorageGRID virtual machines. This URL will typically be a vApp, or Resource Pool. For example:
`vi://vcenter.example.com/vi/sgws`



You can use the VMware `ovftool` utility to determine this value (see the `ovftool` documentation for details).



If you are deploying to a vApp, the virtual machines will not start automatically the first time, and you must power them on manually.

- You have collected all the required information for the deployment configuration file. See [Collect information about your deployment environment](#) for information.
- You have access to the following files from the VMware installation archive for StorageGRID:

Filename	Description
NetApp-SG-version-SHA.vmdk	The virtual machine disk file that is used as a template for creating grid node virtual machines. Note: This file must be in the same folder as the .ovf and .mf files.
vsphere-primary-admin.ovf vsphere-primary-admin.mf	The Open Virtualization Format template file (.ovf) and manifest file (.mf) for deploying the primary Admin Node.
vsphere-non-primary-admin.ovf vsphere-non-primary-admin.mf	The template file (.ovf) and manifest file (.mf) for deploying non-primary Admin Nodes.
vsphere-gateway.ovf vsphere-gateway.mf	The template file (.ovf) and manifest file (.mf) for deploying Gateway Nodes.
vsphere-storage.ovf vsphere-storage.mf	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
deploy-vsphere-ovftool.sh	The Bash shell script used to automate the deployment of virtual grid nodes.
deploy-vsphere-ovftool-sample.ini	The example configuration file for use with the deploy-vsphere-ovftool.sh script.

Define the configuration file for your deployment

You specify the information needed to deploy virtual grid nodes for StorageGRID in a configuration file, which is used by the `deploy-vsphere-ovftool.sh` Bash script. You can modify an example configuration file, so that you don't have to create the file from scratch.

Steps

1. Make a copy of the example configuration file (`deploy-vsphere-ovftool.sample.ini`). Save the new file as `deploy-vsphere-ovftool.ini` in the same directory as `deploy-vsphere-ovftool.sh`.
2. Open `deploy-vsphere-ovftool.ini`.
3. Enter all of the information required to deploy VMware virtual grid nodes.

See [Configuration file settings](#) for information.

4. When you have entered and verified all of the necessary information, save and close the file.

Configuration file settings

The `deploy-vsphere-ovftool.ini` configuration file contains the settings that are required to deploy virtual grid nodes.

The configuration file first lists global parameters, and then lists node-specific parameters in sections defined by node name. When the file is used:

- *Global parameters* are applied to all grid nodes.
- *Node-specific parameters* override global parameters.

Global parameters

Global parameters are applied to all grid nodes, unless they are overridden by settings in individual sections. Place the parameters that apply to multiple nodes in the global parameter section, and then override these settings as necessary in the sections for individual nodes.

- **OVFTOOL_ARGUMENTS:** You can specify `OVFTOOL_ARGUMENTS` as global settings, or you can apply arguments individually to specific nodes. For example:

```
OVFTOOL_ARGUMENTS = --powerOn --noSSLVerify --diskMode=eagerZeroedThick  
--datastore='datastore_name'
```

You can use the `--powerOffTarget` and `--overwrite` options to shut down and replace existing virtual machines.



You should deploy nodes to different datastores and specify `OVFTOOL_ARGUMENTS` for each node, instead of globally.

- **SOURCE:** The path to the StorageGRID virtual machine template (`.vmdk`) file and the `.ovf` and `.mf` files for individual grid nodes. This defaults to the current directory.

```
SOURCE = /downloads/StorageGRID-Webscale-version/vsphere
```

- **TARGET:** The VMware vSphere virtual infrastructure (vi) URL for the location where StorageGRID will be deployed. For example:

```
TARGET = vi://vcenter.example.com/vm/sgws
```

- **GRID_NETWORK_CONFIG:** The method used to acquire IP addresses, either `STATIC` or `DHCP`. The default is `STATIC`. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_CONFIG = STATIC
```

- **GRID_NETWORK_TARGET:** The name of an existing VMware network to use for the Grid Network. If all or most of the nodes use the same network name, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_TARGET = SG Admin Network
```

- **GRID_NETWORK_MASK:** The network mask for the Grid Network. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_MASK = 255.255.255.0
```

- **GRID_NETWORK_GATEWAY:** The network gateway for the Grid Network. If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_GATEWAY = 10.1.0.1
```

- **GRID_NETWORK_MTU:** Optional. The maximum transmission unit (MTU) on the Grid Network. If specified, the value must be between 1280 and 9216. For example:

```
GRID_NETWORK_MTU = 9000
```

If omitted, 1400 is used.

If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.



The MTU value of the network must match the value configured on the virtual switch port in vSphere that the node is connected to. Otherwise, network performance issues or packet loss might occur.



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

- **ADMIN_NETWORK_CONFIG:** The method used to acquire IP addresses, either DISABLED, STATIC, or DHCP. The default is DISABLED. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_CONFIG = STATIC
```

- **ADMIN_NETWORK_TARGET:** The name of an existing VMware network to use for the Admin Network. This setting is required unless the Admin Network is disabled. If all or most of the nodes use the same network name, you can specify it here. Unlike the Grid Network, all nodes do not need to be connected to the same Admin Network. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_TARGET = SG Admin Network
```

- **ADMIN_NETWORK_MASK:** The network mask for the Admin Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_MASK = 255.255.255.0
```

- **ADMIN_NETWORK_GATEWAY:** The network gateway for the Admin Network. This setting is required if you are using static IP addressing and you specify external subnets in the ADMIN_NETWORK_ESL setting. (That is, it is not required if ADMIN_NETWORK_ESL is empty.) If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_GATEWAY = 10.3.0.1
```

- **ADMIN_NETWORK_ESL:** The external subnet list (routes) for the Admin Network, specified as a comma-separated list of CIDR route destinations. If all or most of the nodes use the same external subnet list, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_ESL = 172.16.0.0/21,172.17.0.0/21
```

- **ADMIN_NETWORK_MTU:** Optional. The maximum transmission unit (MTU) on the Admin Network. Don't specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1400 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. If all or most of the nodes use the same MTU for the Admin Network, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_MTU = 8192
```

- **CLIENT_NETWORK_CONFIG:** The method used to acquire IP addresses, either DISABLED, STATIC, or DHCP. The default is DISABLED. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_CONFIG = STATIC
```

- **CLIENT_NETWORK_TARGET:** The name of an existing VMware network to use for the Client Network. This setting is required unless the Client Network is disabled. If all or most of the nodes use the same network name, you can specify it here. Unlike the Grid Network, all nodes do not need to be connected to the same Client Network. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_TARGET = SG Client Network
```

- **CLIENT_NETWORK_MASK:** The network mask for the Client Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_MASK = 255.255.255.0
```

- **CLIENT_NETWORK_GATEWAY:** The network gateway for the Client Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_GATEWAY = 10.4.0.1
```

- **CLIENT_NETWORK_MTU:** Optional. The maximum transmission unit (MTU) on the Client Network. Don't specify if `CLIENT_NETWORK_CONFIG = DHCP`. If specified, the value must be between 1280 and 9216. If omitted, 1400 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. If all or most of the nodes use the same MTU for the Client Network, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_MTU = 8192
```

- **PORT_REMAP:** Remaps any port used by a node for internal grid node communications or external communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID. For the list of ports used by StorageGRID, see internal grid node communications and external communications in [Networking guidelines](#).



Don't remap the ports you are planning to use to configure load balancer endpoints.



If only `PORT_REMAP` is set, the mapping that you specify is used for both inbound and outbound communications. If `PORT_REMAP_INBOUND` is also specified, `PORT_REMAP` applies only to outbound communications.

The format used is: *network type/protocol/default port used by grid node/new port*, where network type is grid, admin, or client, and protocol is tcp or udp.

For example:

```
PORT_REMAP = client/tcp/18082/443
```

If used alone, this example setting symmetrically maps both inbound and outbound communications for the grid node from port 18082 to port 443. If used in conjunction with `PORT_REMAP_INBOUND`, this example setting maps outbound communications from port 18082 to port 443.

You can also remap multiple ports using a comma-separated list.

For example:

```
PORT_REMAP = client/tcp/18082/443, client/tcp/18083/80
```

- **PORT_REMAP_INBOUND:** Remaps inbound communications for the specified port. If you specify `PORT_REMAP_INBOUND` but don't specify a value for `PORT_REMAP`, outbound communications for the port are unchanged.



Don't remap the ports you are planning to use to configure load balancer endpoints.

The format used is: *network type/protocol/_default port used by grid node/new port*, where network type is grid, admin, or client, and protocol is tcp or udp.

For example:

```
PORT_REMAP_INBOUND = client/tcp/443/18082
```

This example takes traffic that is sent to port 443 to pass an internal firewall and directs it to port 18082, where the grid node is listening for S3 requests.

You can also remap multiple inbound ports using a comma-separated list.

For example:

```
PORT_REMAP_INBOUND = grid/tcp/3022/22, admin/tcp/3022/22
```

- **TEMPORARY_PASSWORD_TYPE:** The type of temporary installation password to be used when accessing the VM console or the StorageGRID Installation API, or using SSH, before the node joins the grid.



If all or most of the nodes use the same type of temporary installation password, specify the type in the global parameter section. Then, optionally use a different setting for an individual node. For example, if you select **Use Custom Password** globally, you can use **CUSTOM_TEMPORARY_PASSWORD=<password>** to set the password for each node.

TEMPORARY_PASSWORD_TYPE can be one of the following:

- **Use node name:** The node name is used as the temporary installation password and provides access to VM console, the StorageGRID Installation API, and SSH.
- **Disable password:** No temporary installation password will be used. If you need to access the VM to debug installation issues, see [Troubleshoot installation issues](#).
- **Use custom password:** The value provided with **CUSTOM_TEMPORARY_PASSWORD=<password>** is used as the temporary installation password and provides access to VM console, the StorageGRID Installation API, and SSH.



Optionally, you can omit the **TEMPORARY_PASSWORD_TYPE** parameter and only specify **CUSTOM_TEMPORARY_PASSWORD=<password>**.

• **CUSTOM_TEMPORARY_PASSWORD=<password>**

Optional. The temporary password to use during installation when accessing VM console, the StorageGRID Installation API, and SSH. Ignored if **TEMPORARY_PASSWORD_TYPE** is set to **Use node name** or **Disable password**.

Node-specific parameters

Each node is in its own section of the configuration file. Each node requires the following settings:

- The section head defines the node name that will be displayed in the Grid Manager. You can override that value by specifying the optional **NODE_NAME** parameter for the node.
- **NODE_TYPE:** VM_Admin_Node, VM_Storage_Node, or VM_API_Gateway_Node
- **STORAGE_TYPE:** combined, data, or metadata. This optional parameter for storage nodes defaults to combined (data and metadata) if it is not specified. For more information, see [Types of Storage Nodes](#).
- **GRID_NETWORK_IP:** The IP address for the node on the Grid Network.
- **ADMIN_NETWORK_IP:** The IP address for the node on the Admin Network. Required only if the node is attached to the Admin Network and **ADMIN_NETWORK_CONFIG** is set to **STATIC**.
- **CLIENT_NETWORK_IP:** The IP address for the node on the Client Network. Required only if the node is attached to the Client Network and **CLIENT_NETWORK_CONFIG** for this node is set to **STATIC**.
- **ADMIN_IP:** The IP address for the primary Admin node on the Grid Network. Use the value that you specify as the **GRID_NETWORK_IP** for the primary Admin Node. If you omit this parameter, the node attempts to discover the primary Admin Node IP using mDNS. For more information, see [How grid nodes discover the primary Admin Node](#).



The **ADMIN_IP** parameter is ignored for the primary Admin Node.

- Any parameters that were not set globally. For example, if a node is attached to the Admin Network and you did not specify **ADMIN_NETWORK** parameters globally, you must specify them for the node.

Primary Admin Node

The following additional settings are required for the primary Admin Node:

- **NODE_TYPE:** VM_Admin_Node
- **ADMIN_ROLE:** Primary

This example entry is for a primary Admin Node that is on all three networks:

```
[DC1-ADM1]
ADMIN_ROLE = Primary
NODE_TYPE = VM_Admin_Node
TEMPORARY_PASSWORD_TYPE = Use custom password
CUSTOM_TEMPORARY_PASSWORD = Passw0rd

GRID_NETWORK_IP = 10.1.0.2
ADMIN_NETWORK_IP = 10.3.0.2
CLIENT_NETWORK_IP = 10.4.0.2
```

The following additional setting is optional for the primary Admin Node:

- **DISK:** By default, Admin Nodes are assigned two additional 200 GB hard disks for audit and database use. You can increase these settings using the DISK parameter. For example:

```
DISK = INSTANCES=2, CAPACITY=300
```



For Admin nodes, INSTANCES must always equal 2.

Storage Node

The following additional setting is required for Storage Nodes:

- **NODE_TYPE:** VM_Storage_Node

This example entry is for a Storage Node that is on the Grid and Admin Networks, but not on the Client Network. This node uses the ADMIN_IP setting to specify the primary Admin Node's IP address on the Grid Network.

```
[DC1-S1]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.0.3
ADMIN_NETWORK_IP = 10.3.0.3

ADMIN_IP = 10.1.0.2
```

This second example entry is for a Storage Node on a Client Network where the customer's enterprise networking policy states that an S3 client application is only permitted to access the Storage Node using either port 80 or 443. The example configuration file uses PORT_REMAP to enable the Storage Node to send and receive S3 messages on port 443.

```
[DC2-S1]
  NODE_TYPE = VM_Storage_Node

  GRID_NETWORK_IP = 10.1.1.3
  CLIENT_NETWORK_IP = 10.4.1.3
  PORT_REMAP = client/tcp/18082/443

  ADMIN_IP = 10.1.0.2
```

The last example creates a symmetric remapping for ssh traffic from port 22 to port 3022, but explicitly sets the values for both inbound and outbound traffic.

```
[DC1-S3]
  NODE_TYPE = VM_Storage_Node

  GRID_NETWORK_IP = 10.1.1.3

  PORT_REMAP = grid/tcp/22/3022
  PORT_REMAP_INBOUND = grid/tcp/3022/22

  ADMIN_IP = 10.1.0.2
```

The following additional settings are optional for Storage Nodes:

- **DISK:** By default, Storage Nodes are assigned three 4 TB disks for RangeDB use. You can increase these settings with the DISK parameter. For example:

```
DISK = INSTANCES=16, CAPACITY=4096
```

- **STORAGE_TYPE:** By default, all new Storage Nodes are configured to store both object data and metadata, known as a *combined* Storage Node. You can change the Storage Node type to store only data or metadata with the STORAGE_TYPE parameter. For example:

```
STORAGE_TYPE = data
```

Gateway Node

The following additional setting is required for Gateway Nodes:

- **NODE_TYPE:** VM_API_Gateway

This example entry is for an example Gateway Node on all three networks. In this example, no Client Network parameters were specified in the global section of the configuration file, so they must be specified for the node:

```
[DC1-G1]
NODE_TYPE = VM_API_Gateway

GRID_NETWORK_IP = 10.1.0.5
ADMIN_NETWORK_IP = 10.3.0.5

CLIENT_NETWORK_CONFIG = STATIC
CLIENT_NETWORK_TARGET = SG Client Network
CLIENT_NETWORK_MASK = 255.255.255.0
CLIENT_NETWORK_GATEWAY = 10.4.0.1
CLIENT_NETWORK_IP = 10.4.0.5

ADMIN_IP = 10.1.0.2
```

Non-primary Admin Node

The following additional settings are required for non-primary Admin Nodes:

- **NODE_TYPE:** VM_Admin_Node
- **ADMIN_ROLE:** Non-Primary

This example entry is for a non-primary Admin Node that is not on the Client Network:

```
[DC2-ADM1]
ADMIN_ROLE = Non-Primary
NODE_TYPE = VM_Admin_Node

GRID_NETWORK_TARGET = SG Grid Network
GRID_NETWORK_IP = 10.1.0.6
ADMIN_NETWORK_IP = 10.3.0.6

ADMIN_IP = 10.1.0.2
```

The following additional setting is optional for non-primary Admin Nodes:

- **DISK:** By default, Admin Nodes are assigned two additional 200 GB hard disks for audit and database use. You can increase these settings using the DISK parameter. For example:

```
DISK = INSTANCES=2, CAPACITY=300
```



For Admin nodes, INSTANCES must always equal 2.

Run the Bash script

You can use the `deploy-vsphere-ovftool.sh` Bash script and the `deploy-vsphere-ovftool.ini` configuration

file you modified to automate the deployment of StorageGRID nodes in VMware vSphere.

Before you begin

You have created a `deploy-vsphere-ovftool.ini` configuration file for your environment.

You can use the help available with the Bash script by entering the help commands (`-h/--help`). For example:

```
./deploy-vsphere-ovftool.sh -h
```

or

```
./deploy-vsphere-ovftool.sh --help
```

Steps

1. Log in to the Linux machine you are using to run the Bash script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/vsphere
```

3. To deploy all grid nodes, run the Bash script with the appropriate options for your environment.

For example:

```
./deploy-vsphere-ovftool.sh --username=user --password=pwd ./deploy-vsphere-ovftool.ini
```

4. If a grid node failed to deploy because of an error, resolve the error and rerun the Bash script for only that node.

For example:

```
./deploy-vsphere-ovftool.sh --username=user --password=pwd --single-node="DC1-S3" ./deploy-vsphere-ovftool.ini
```

The deployment is complete when the status for each node is "Passed."

Deployment Summary

node	attempts	status
DC1-ADM1	1	Passed
DC1-G1	1	Passed
DC1-S1	1	Passed
DC1-S2	1	Passed
DC1-S3	1	Passed

Automate the configuration of StorageGRID

After deploying the grid nodes, you can automate the configuration of the StorageGRID system.

Before you begin

- You know the location of the following files from the installation archive.

Filename	Description
<code>configure-storagegrid.py</code>	Python script used to automate the configuration
<code>configure-storagegrid.sample.json</code>	Example configuration file for use with the script
<code>configure-storagegrid.blank.json</code>	Blank configuration file for use with the script

- You have created a `configure-storagegrid.json` configuration file. To create this file, you can modify the example configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

You can use the `configure-storagegrid.py` Python script and the `configure-storagegrid.json` grid configuration file to automate the configuration of your StorageGRID system.



You can also configure the system using the Grid Manager or the Installation API.

Steps

- Log in to the Linux machine you are using to run the Python script.
- Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/platform
```

where `platform` is `debs`, `rpms`, or `vsphere`.

3. Run the Python script and use the configuration file you created.

For example:

```
./configure-storagegrid.py ./configure-storagegrid.json --start-install
```

Result

A Recovery Package .zip file is generated during the configuration process, and it is downloaded to the directory where you are running the installation and configuration process. You must back up the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fails. For example, copy it to a secure, backed up network location and to a secure cloud storage location.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

If you specified that random passwords should be generated, open the `Passwords.txt` file and look for the passwords required to access your StorageGRID system.

```
#####
##### The StorageGRID "Recovery Package" has been downloaded as: #####
#####      ./sgws-recovery-package-994078-rev1.zip      #####
##### Safeguard this file as it will be needed in case of a #####
#####      StorageGRID node recovery.      #####
#####
```

Your StorageGRID system is installed and configured when a confirmation message is displayed.

```
StorageGRID has been configured and installed.
```

Related information

- [Navigate to the Grid Manager](#)
- [Installation REST API](#)

Deploy virtual machine grid nodes (VMware)

Collect information about your deployment environment

Before deploying grid nodes, you must collect information about your network configuration and VMware environment.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

VMware information

You must access the deployment environment and collect information about the VMware environment; the networks that were created for the Grid, Admin, and Client Networks; and the storage volume types you plan to use for Storage Nodes.

You must collect information about your VMware environment, including the following:

- The username and password for a VMware vSphere account that has appropriate permissions to complete the deployment.
- Host, datastore, and network configuration information for each StorageGRID node virtual machine.



VMware live vMotion causes the virtual machine clock time to jump and is not supported for grid nodes of any type. Though rare, incorrect clock times can result in loss of data or configuration updates.

Grid Network information

You must collect information about the VMware network created for the StorageGRID Grid Network (required), including:

- The network name.
- The method used to assign IP addresses, either static or DHCP.
 - If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).
 - If you are using DHCP, the IP address of the primary Admin Node on the Grid Network. See [How grid nodes discover the primary Admin Node](#) for more information.

Admin Network information

For nodes that will be connected to the optional StorageGRID Admin Network, you must collect information about the VMware network created for this network, including:

- The network name.
- The method used to assign IP addresses, either static or DHCP.
 - If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).
 - If you are using DHCP, the IP address of the primary Admin Node on the Grid Network. See [How grid nodes discover the primary Admin Node](#) for more information.
- The external subnet list (ESL) for the Admin Network.

Client Network information

For nodes that will be connected to the optional StorageGRID Client Network, you must collect information about the VMware network created for this network, including:

- The network name.
- The method used to assign IP addresses, either static or DHCP.
- If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).

Information about additional interfaces

You can optionally add trunk or access interfaces to the VM in vCenter after you install the node. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use VLAN interfaces to segregate the traffic belonging to different applications or tenants. Or, you might want to add an access interface to use in a high availability (HA) group.

The interfaces you add are displayed on the VLAN interfaces page and on the HA groups page in the Grid Manager.

- If you add a trunk interface, configure one or more VLAN interfaces for each new parent interface. See [configure VLAN interfaces](#).
- If you add an access interface, you must add it directly to HA groups. See [configure high availability groups](#).

Storage volumes for virtual Storage Nodes

You must collect the following information for virtual machine-based Storage Nodes:

- The number and size of storage volumes (storage LUNs) you plan to add. See [Storage and performance requirements](#).

Grid configuration information

You must collect information to configure your grid:

- Grid license
- Network Time Protocol (NTP) server IP addresses
- DNS server IP addresses

How grid nodes discover the primary Admin Node

Grid nodes communicate with the primary Admin Node for configuration and management. Each grid node must know the IP address of the primary Admin Node on the Grid Network.

To ensure that a grid node can access the primary Admin Node, you can do either of the following when deploying the node:

- You can use the ADMIN_IP parameter to enter the primary Admin Node's IP address manually.
- You can omit the ADMIN_IP parameter to have the grid node discover the value automatically. Automatic discovery is especially useful when the Grid Network uses DHCP to assign the IP address to the primary Admin Node.

Automatic discovery of the primary Admin Node is accomplished using a multicast domain name system (mDNS). When the primary Admin Node first starts up, it publishes its IP address using mDNS. Other nodes on the same subnet can then query for the IP address and acquire it automatically. However, because multicast IP traffic is not normally routable across subnets, nodes on other subnets can't acquire the primary Admin Node's IP address directly.

If you use automatic discovery:



- You must include the ADMIN_IP setting for at least one grid node on any subnets that the primary Admin Node is not directly attached to. This grid node will then publish the primary Admin Node's IP address for other nodes on the subnet to discover with mDNS.
- Ensure that your network infrastructure supports passing multi-cast IP traffic within a subnet.

Deploy a StorageGRID node as a virtual machine

You use VMware vSphere Web Client to deploy each grid node as a virtual machine. During deployment, each grid node is created and connected to one or more StorageGRID networks.

If you need to deploy any StorageGRID appliance Storage Nodes, see [Deploy appliance Storage Node](#).

Optionally, you can remap node ports or increase CPU or memory settings for the node before powering it on.

Before you begin

- You have reviewed how to [plan and prepare for installation](#), and you understand the requirements for software, CPU and RAM, and storage and performance.
- You are familiar with VMware vSphere Hypervisor and have experience deploying virtual machines in this environment.



The `open-vm-tools` package, an open-source implementation similar to VMware Tools, is included with the StorageGRID virtual machine. You don't need to install VMware Tools manually.

- You have downloaded and extracted the correct version of the StorageGRID installation archive for VMware.



If you are deploying the new node as part of an expansion or recovery operation, you must use the version of StorageGRID that is currently running on the grid.

- You have the StorageGRID Virtual Machine Disk (.vmdk) file:

```
NetApp-SG-version-SHA.vmdk
```

- You have the .ovf and .mf files for each type of grid node you are deploying:

Filename	Description
vsphere-primary-admin.ovf	The template file and manifest file for the primary Admin Node.
vsphere-primary-admin.mf	
vsphere-non-primary-admin.ovf	The template file and manifest file for a non-primary Admin Node.
vsphere-non-primary-admin.mf	

Filename	Description
vsphere-storage.ovf	The template file and manifest file for a Storage Node.
vsphere-storage.mf	
vsphere-gateway.ovf	The template file and manifest file for a Gateway Node.
vsphere-gateway.mf	

- The .vdmk, .ovf, and .mf files are all in the same directory.
- You have a plan to minimize failure domains. For example, you should not deploy all Gateway Nodes on a single vSphere ESXi host.



In a production deployment, don't run more than one Storage Node on a single virtual machine. Do not run multiple virtual machines on the same ESXi host if that would create an unacceptable failure-domain issue.

- If you are deploying a node as part of an expansion or recovery operation, you have the [instructions for expanding a StorageGRID system](#) or the [recovery and maintenance instructions](#).
- If you are deploying a StorageGRID node as a virtual machine with storage assigned from a NetApp ONTAP system, you have confirmed that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running as a virtual machine on a VMware host, ensure the volume backing the datastore for the node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

About this task

Follow these instructions to initially deploy VMware nodes, add a new VMware node in an expansion, or replace a VMware node as part of a recovery operation. Except as noted in the steps, the node deployment procedure is the same for all node types, including Admin Nodes, Storage Nodes, and Gateway Nodes.

If you are installing a new StorageGRID system:

- You can deploy nodes in any order.
- You must ensure that each virtual machine can connect to the primary Admin Node over the Grid Network.
- You must deploy all grid nodes before configuring the grid.

If you are performing an expansion or recovery operation:

- You must ensure that the new virtual machine can connect to all other nodes over the Grid Network.

If you need to remap any of the node's ports, don't power on the new node until the port remap configuration is complete.

Steps

1. Using VCenter, deploy an OVF template.

If you specify a URL, point to a folder containing the following files. Otherwise, select each of these files from a local directory.

```
NetApp-SG-version-SHA.vmdk  
vsphere-node.ovf  
vsphere-node.mf
```

For example, if this is the first node you are deploying, use these files to deploy the primary Admin Node for your StorageGRID system:

```
NetApp-SG-version-SHA.vmdk  
vsphere-primary-admin.ovf  
vsphere-primary-admin.mf
```

2. Provide a name for the virtual machine.

The standard practice is to use the same name for both the virtual machine and the grid node.

3. Place the virtual machine in the appropriate vApp or resource pool.

4. If you are deploying the primary Admin Node, read and accept the End User License Agreement.

Depending on your version of vCenter, the order of the steps will vary for accepting the End User License Agreement, specifying the name of the virtual machine, and selecting a datastore.

5. Select storage for the virtual machine.

If you are deploying a node as part of recovery operation, perform the instructions in the [storage recovery step](#) to add new virtual disks, reattach virtual hard disks from the failed grid node, or both.

When deploying a Storage Node, use 3 or more storage volumes, with each storage volume being 4 TB or larger. You must assign at least 4 TB to volume 0.



The Storage Node .ovf file defines several VMDKs for storage. Unless these VMDKs meet your storage requirements, you should remove them and assign appropriate VMDKs or RDMs for storage before powering up the node. VMDKs are more commonly used in VMware environments and are easier to manage, while RDMs might provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).



Some StorageGRID installations might use larger, more active storage volumes than typical virtualized workloads. You might need to tune some hypervisor parameters, such as `MaxAddressableSpaceTB`, to achieve optimal performance. If you encounter poor performance, contact your virtualization support resource to determine whether your environment could benefit from workload-specific configuration tuning.

6. Select networks.

Determine which StorageGRID networks the node will use by selecting a destination network for each source network.

- The Grid Network is required. You must select a destination network in the vSphere environment.

+

The Grid Network is used for all internal StorageGRID traffic. It provides connectivity among all nodes in the grid, across all sites and subnets. All nodes on the Grid Network must be able to communicate with all other nodes.

- If you use the Admin Network, select a different destination network in the vSphere environment. If you don't use the Admin Network, select the same destination you selected for the Grid Network.
- If you use the Client Network, select a different destination network in the vSphere environment. If you don't use the Client Network, select the same destination you selected for the Grid Network.
- If you use an Admin or Client network, nodes do not have to be on the same Admin or Client networks.

7. For **Customize Template**, configure the required StorageGRID node properties.

a. Enter the **Node name**.



If you are recovering a grid node, you must enter the name of the node you are recovering.

b. Use the **Temporary installation password** drop-down to specify a temporary installation password, so that you can access the VM console or the StorageGRID Installation API, or use SSH, before the new node joins the grid.



The temporary installation password is only used during node installation. After a node has been added to the grid, you can access it using the [node console password](#), which is listed in the `Passwords.txt` file in the Recovery Package.

- **Use node name:** The value you provided for the **Node name** field is used as the temporary installation password.
- **Use custom password:** A custom password is used as the temporary installation password.
- **Disable password:** No temporary installation password will be used. If you need to access the VM to debug installation issues, see [Troubleshoot installation issues](#).

c. If you selected **Use custom password**, specify the temporary installation password you want to use in the **Custom password** field.

d. In the **Grid Network (eth0)** section, select STATIC or DHCP for the **Grid network IP configuration**.

- If you select STATIC, enter the **Grid network IP**, **Grid network mask**, **Grid network gateway**, and **Grid network MTU**.
- If you select DHCP, the **Grid network IP**, **Grid network mask**, and **Grid network gateway** are automatically assigned.

e. In the **Primary Admin IP** field, enter the IP address of the primary Admin Node for the Grid Network.



This step does not apply if the node you are deploying is the primary Admin Node.

If you omit the primary Admin Node IP address, the IP address will be automatically discovered if the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet. However, it is recommended to set the primary Admin Node IP address here.

f. In the **Admin Network (eth1)** section, select STATIC, DHCP, or DISABLED for the **Admin network IP configuration**.

- If you don't want to use the Admin Network, select **DISABLED** and enter **0.0.0.0** for the Admin Network IP. You can leave the other fields blank.
 - If you select **STATIC**, enter the **Admin network IP**, **Admin network mask**, **Admin network gateway**, and **Admin network MTU**.
 - If you select **STATIC**, enter the **Admin network external subnet list**. You must also configure a gateway.
 - If you select **DHCP**, the **Admin network IP**, **Admin network mask**, and **Admin network gateway** are automatically assigned.
- g. In the **Client Network (eth2)** section, select **STATIC**, **DHCP**, or **DISABLED** for the **Client network IP configuration**.
- If you don't want to use the Client Network, select **DISABLED** and enter **0.0.0.0** for the Client Network IP. You can leave the other fields blank.
 - If you select **STATIC**, enter the **Client network IP**, **Client network mask**, **Client network gateway**, and **Client network MTU**.
 - If you select **DHCP**, the **Client network IP**, **Client network mask**, and **Client network gateway** are automatically assigned.
8. Review the virtual machine configuration and make any changes necessary.
9. When you are ready to complete, select **Finish** to start the upload of the virtual machine.
10. If you deployed this node as part of recovery operation and this is not a full-node recovery, perform these steps after deployment is complete:
- a. Right-click the virtual machine, and select **Edit Settings**.
 - b. Select each default virtual hard disk that has been designated for storage, and select **Remove**.
 - c. Depending on your data recovery circumstances, add new virtual disks according to your storage requirements, reattach any virtual hard disks preserved from the previously removed failed grid node, or both.

Note the following important guidelines:

- If you are adding new disks you should use the same type of storage device that was in use before node recovery.
 - The Storage Node .ovf file defines several VMDKs for storage. Unless these VMDKs meet your storage requirements, you should remove them and assign appropriate VMDKs or RDMs for storage before powering up the node. VMDKs are more commonly used in VMware environments and are easier to manage, while RDMs might provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).
11. If you need to remap the ports used by this node, follow these steps.

You might need to remap a port if your enterprise networking policies restrict access to one or more ports that are used by StorageGRID. See the [networking guidelines](#) for the ports used by StorageGRID.



Don't remap the ports used in load balancer endpoints.

- a. Select the new VM.
- b. From the Configure tab, select **Settings > vApp Options**. The location of **vApp Options** depends on the version of vCenter.
- c. In the **Properties** table, locate **PORT_REMAP_INBOUND** and **PORT_REMAP**.

- d. To symmetrically map both inbound and outbound communications for a port, select **PORT_REMAP**.



If only PORT_REMAP is set, the mapping that you specify applies to both inbound and outbound communications. If PORT_REMAP_INBOUND is also specified, PORT_REMAP applies only to outbound communications.

- i. Select **Set Value**.

- ii. Enter the port mapping:

```
<network type>/<protocol>/<default port used by grid node>/<new port>
```

<network type> is grid, admin, or client, and <protocol> is tcp or udp.

For example, to remap ssh traffic from port 22 to port 3022, enter:

```
client/tcp/22/3022
```

You can remap multiple ports using a comma-separated list.

For example:

```
client/tcp/18082/443, client/tcp/18083/80
```

- iii. Select **OK**.

- e. To specify the port used for inbound communications to the node, select **PORT_REMAP_INBOUND**.



If you specify PORT_REMAP_INBOUND and don't specify a value for PORT_REMAP, outbound communications for the port are unchanged.

- i. Select **Set Value**.

- ii. Enter the port mapping:

```
<network type>/<protocol>/<remapped inbound port>/<default inbound port used by grid node>
```

<network type> is grid, admin, or client, and <protocol> is tcp or udp.

For example, to remap inbound SSH traffic that is sent to port 3022 so that it is received at port 22 by the grid node, enter the following:

```
client/tcp/3022/22
```

You can remap multiple inbound ports using a comma-separated list.

For example:

```
grid/tcp/3022/22, admin/tcp/3022/22
```

- iii. Select **OK**

12. If you want to increase the CPU or memory for the node from the default settings:

- a. Right-click the virtual machine, and select **Edit Settings**.

- b. Change the number of CPUs or the amount of memory as required.

Set the **Memory Reservation** to the same size as the **Memory** allocated to the virtual machine.

- c. Select **OK**.

13. Power on the virtual machine.

After you finish

If you deployed this node as part of an expansion or recovery procedure, return to those instructions to complete the procedure.

Configure the grid and complete installation (VMware)

Navigate to the Grid Manager

You use the Grid Manager to define all of the information required to configure your StorageGRID system.

Before you begin

The primary Admin Node must be deployed and have completed the initial startup sequence.

Steps

1. Open your web browser and navigate to:

```
https://primary_admin_node_ip
```

Alternatively, you can access the Grid Manager on port 8443:


```
https://primary_admin_node_ip:8443
```

You can use the IP address for the primary Admin Node IP on the Grid Network or on the Admin Network, as appropriate for your network configuration. You might need to use the security/advanced option in your browser to navigate to an untrusted certificate.

2. Manage a temporary installer password as needed:
 - If a password has already been set using one of these methods, enter the password to proceed.
 - A user set the password while accessing the installer previously
 - The SSH/console password was automatically imported from the OVF properties
 - If a password has not been set, optionally set a password to secure the StorageGRID installer.
3. Select **Install a StorageGRID system**.

The page used to configure a StorageGRID grid appears.

Install



License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

Browse

Specify the StorageGRID license information

You must specify the name for your StorageGRID system and upload the license file provided by NetApp.

Steps

1. On the License page, enter a meaningful name for your StorageGRID system in the **Grid Name** field.

After installation, the name is displayed at the top of the Nodes menu.

2. Select **Browse**, locate the NetApp license file (*NLF-unique-id.txt*), and select **Open**.

The license file is validated, and the serial number is displayed.



The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product. You can update to a license that offers support after installation.

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File NLF-959007-Internal.txt

License Serial Number

3. Select **Next**.

Add sites

You must create at least one site when you are installing StorageGRID. You can create additional sites to increase the reliability and storage capacity of your StorageGRID system.

Steps

1. On the Sites page, enter the **Site Name**.
2. To add additional sites, click the plus sign next to the last site entry and enter the name in the new **Site Name** text box.

Add as many additional sites as required for your grid topology. You can add up to 16 sites.

NetApp® StorageGRID®

Help ▾

Install

1

2

3

4

5

6

7

8

License

Sites

Grid Network

Grid Nodes

NTP

DNS

Passwords

Summary

Sites

In a single-site deployment, infrastructure and operations are centralized in one site.

In a multi-site deployment, infrastructure can be distributed asymmetrically across sites, and proportional to the needs of each site. Typically, sites are located in geographically different locations. Having multiple sites also allows the use of distributed replication and erasure coding for increased availability and resiliency.

Site Name 1

Raleigh

✕

Site Name 2

Atlanta

+ ✕

3. Click **Next**.

Specify Grid Network subnets

You must specify the subnets that are used on the Grid Network.

About this task

The subnet entries include the subnets for the Grid Network for each site in your StorageGRID system, along with any subnets that need to be reachable through the Grid Network.

If you have multiple grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway.

Steps

1. Specify the CIDR network address for at least one Grid Network in the **Subnet 1** text box.
2. Click the plus sign next to the last entry to add an additional network entry. You must specify all subnets for all sites in the Grid Network.
 - If you have already deployed at least one node, click **Discover Grid Networks Subnets** to automatically populate the Grid Network Subnet List with the subnets reported by grid nodes that have registered with the Grid Manager.

- You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

NetApp® StorageGRID® Help ▾

Install

1 License 2 Sites 3 **Grid Network** 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Grid Network

You must specify the subnets that are used on the Grid Network. These entries typically include the subnets for the Grid Network for each site in your StorageGRID system. Select Discover Grid Networks to automatically add subnets based on the network configuration of all registered nodes.

Note: You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

Subnet 1 +

3. Click **Next**.

Approve pending grid nodes

You must approve each grid node before it can join the StorageGRID system.

Before you begin

You have deployed all virtual and StorageGRID appliance grid nodes.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

Steps

1. Review the Pending Nodes list, and confirm that it shows all of the grid nodes you deployed.



If a grid node is missing, confirm that it was deployed successfully and has the correct Grid Network IP of the primary admin node set for ADMIN_IP.

2. Select the radio button next to a pending node you want to approve.



Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

+ Approve		✖ Remove		Search		
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address	
☒	50:6b:4b:42:d7:00	NetApp-SGA	Storage Node	StorageGRID Appliance	172.16.5.20/21	
						◀ ▶

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

Edit		Reset		✖ Remove		Search	
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address	
☐	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21	
☐	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21	
☐	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21	
☐	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21	
☐	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21	
							◀ ▶

3. Click **Approve**.

4. In General Settings, modify settings for the following properties, as necessary:

- **Site:** The system name of the site for this grid node.
- **Name:** The system name for the node. The name defaults to the name you specified when you configured the node.

System names are required for internal StorageGRID operations and can't be changed after you complete the installation. However, during this step of the installation process, you can change system names as required.



For a VMware node, you can change the name here, but this action will not change the name of the virtual machine in vSphere.

- **NTP Role:** The Network Time Protocol (NTP) role of the grid node. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the Primary role to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

- **Storage Type** (Storage Nodes only): Specify that a new Storage Node be used exclusively for data only, metadata only, or both. The options are **Data and metadata** ("combined"), **Data only**, and **Metadata only**.



See [Types of Storage Nodes](#) for information about requirements for these node types.

- **ADC service** (Storage Nodes only): Select **Automatic** to let the system determine whether the node requires the Administrative Domain Controller (ADC) service. The ADC service keeps track of the location and availability of grid services. At least three Storage Nodes at each site must include the ADC service. You can't add the ADC service to a node after it is deployed.

5. In Grid Network, modify settings for the following properties as necessary:

- **IPv4 Address (CIDR)**: The CIDR network address for the Grid Network interface (eth0 inside the container). For example: 192.168.1.234/21
- **Gateway**: The Grid Network gateway. For example: 192.168.0.1



The gateway is required if there are multiple grid subnets.



If you selected DHCP for the Grid Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

6. If you want to configure the Admin Network for the grid node, add or update the settings in the Admin Network section as necessary.

Enter the destination subnets of the routes out of this interface in the **Subnets (CIDR)** text box. If there are multiple Admin subnets, the Admin gateway is required.



If you selected DHCP for the Admin Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Admin Network was not configured during the initial installation using the StorageGRID Appliance Installer, it can't be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

Rebooting can take several minutes.

- Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- Select **Configure Networking > IP Configuration** and configure the enabled networks.
- Return to the Home page and click **Start Installation**.
- In the Grid Manager: If the node is listed in the Approved Nodes table, remove the node.
- Remove the node from the Pending Nodes table.

- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page of the Appliance Installer.

For additional information, see the [Quick start for hardware installation](#) to locate instructions for your appliance.

- 7. If you want to configure the Client Network for the grid node, add or update the settings in the Client Network section as necessary. If the Client Network is configured, the gateway is required, and it becomes the default gateway for the node after installation.



If you selected DHCP for the Client Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the configured IP address is not within a DHCP address pool.

Appliances: For a StorageGRID appliance, if the Client Network was not configured during the initial installation using the StorageGRID Appliance Installer, it can't be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

Rebooting can take several minutes.

- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, remove the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page of the Appliance Installer.

For additional information, see the [Quick start for hardware installation](#) to locate instructions for your appliance.

- 8. Click **Save**.

The grid node entry moves to the Approved Nodes list.



Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

+ Approve

✖ Remove

Search

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
No results found.				

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

Edit

Reset

✖ Remove

Search

	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
☐	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
☐	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
☐	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
☐	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
☐	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21
☐	50:6b:4b:42:d7:00	NetApp-SGA	Raleigh	Storage Node	StorageGRID Appliance	172.16.5.20/21

9. Repeat these steps for each pending grid node you want to approve.

You must approve all nodes that you want in the grid. However, you can return to this page at any time before you click **Install** on the Summary page. You can modify the properties of an approved grid node by selecting its radio button and clicking **Edit**.

10. When you are done approving grid nodes, click **Next**.

Specify Network Time Protocol server information

You must specify the Network Time Protocol (NTP) configuration information for the StorageGRID system, so that operations performed on separate servers can be kept synchronized.

About this task

You must specify IPv4 addresses for the NTP servers.

You must specify external NTP servers. The specified NTP servers must use the NTP protocol.

You must specify four NTP server references of Stratum 3 or better to prevent issues with time drift.



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

[Support boundary to configure the Windows Time service for high-accuracy environments](#)

The external NTP servers are used by the nodes to which you previously assigned Primary NTP roles.



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

Perform additional checks for VMware, such as ensuring that the hypervisor uses the same NTP source as the virtual machine, and using VMTools to disable the time sync between the hypervisor and StorageGRID virtual machines.

Steps

- 1. Specify the IPv4 addresses for at least four NTP servers in the **Server 1** to **Server 4** text boxes.
- 2. If necessary, select the plus sign next to the last entry to add additional server entries.

NetApp® StorageGRID®

Help ▾

Install

1

License

2

Sites

3

Grid Network

4

Grid Nodes

5

NTP

6

DNS

7

Passwords

8

Summary

Network Time Protocol

Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync.

Server 1

10.60.248.183

Server 2

10.227.204.142

Server 3

10.235.48.111

Server 4

0.0.0.0

+

- 3. Select **Next**.

Specify DNS server information

You must specify DNS information for your StorageGRID system, so that you can access external servers using hostnames instead of IP addresses.

About this task

Specifying [DNS server information](#) allows you to use Fully Qualified Domain Name (FQDN) hostnames rather than IP addresses for email notifications and AutoSupport.

To ensure proper operation, specify two or three DNS servers. If you specify more than three, it is possible that only three will be used because of known OS limitations on some platforms. If you have routing restrictions in your environment, you can [customize the DNS server list](#) for individual nodes (typically all nodes at a site) to use a different set of up to three DNS servers.

If possible, use DNS servers that each site can access locally to ensure that an islanded site can resolve the FQDNs for external destinations.

Steps

1. Specify the IPv4 address for at least one DNS server in the **Server 1** text box.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

The screenshot shows the NetApp StorageGRID installation wizard. At the top is a blue header with "NetApp® StorageGRID®" and a "Help" link. Below the header is a progress bar with eight steps: 1. License, 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP, 6. DNS (highlighted in blue), 7. Passwords, and 8. Summary. Below the progress bar is the "Domain Name Service" section. It contains a text box for "Server 1" with the value "10.224.223.130" and a minus sign icon. Below that is a text box for "Server 2" with the value "10.224.223.136" and a plus-minus icon. The text below the progress bar reads: "Domain Name Service. Enter the IP address for at least one Domain Name System (DNS) server, so that server hostnames can be used instead of IP addresses. Specifying at least two DNS servers is recommended. Configuring DNS enables server connectivity, email notifications, and NetApp AutoSupport."

The best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

Specify the StorageGRID system passwords

As part of installing your StorageGRID system, you need to enter the passwords to use to secure your system and perform maintenance tasks.

About this task

Use the Install passwords page to specify the provisioning passphrase and the grid management root user password.

- The provisioning passphrase is used as an encryption key and is not stored by the StorageGRID system.
- You must have the provisioning passphrase for installation, expansion, and maintenance procedures, including downloading the Recovery Package. Therefore, it is important that you store the provisioning passphrase in a secure location.
- You can change the provisioning passphrase from the Grid Manager if you have the current one.
- The grid management root user password can be changed using the Grid Manager.

- Randomly generated command line console and SSH passwords are stored in the `Passwords.txt` file in the Recovery Package.

Steps

1. In **Provisioning Passphrase**, enter the provisioning passphrase that will be required to make changes to the grid topology of your StorageGRID system.

Store the provisioning passphrase in a secure place.



If after the installation completes and you want to change the provisioning passphrase later, you can use the Grid Manager. Select **CONFIGURATION > Access control > Grid passwords**.

2. In **Confirm Provisioning Passphrase**, reenter the provisioning passphrase to confirm it.
3. In **Grid Management Root User Password**, enter the password to use to access the Grid Manager as the "root" user.

Store the password in a secure place.

4. In **Confirm Root User Password**, reenter the Grid Manager password to confirm it.

NetApp® StorageGRID®
Help

Install

1 License
2 Sites
3 Grid Network
4 Grid Nodes
5 NTP
6 DNS
7 Passwords
8 Summary

Passwords

Enter secure passwords that meet your organization's security policies. A text file containing the command line passwords must be downloaded during the final installation step.

Provisioning Passphrase	
Confirm Provisioning Passphrase	
Grid Management Root User Password	
Confirm Root User Password	

☒ Create random command line passwords.

5. If you are installing a grid for proof of concept or demo purposes, optionally clear the **Create random command line passwords** checkbox.

For production deployments, random passwords should always be used for security reasons. Clear **Create random command line passwords** only for demo grids if you want to use default passwords to access grid nodes from the command line using the "root" or "admin" account.



You are prompted to download the Recovery Package file (`sgws-recovery-package-id-revision.zip`) after you click **Install** on the Summary page. You must [download this file](#) to complete the installation. The passwords required to access the system are stored in the `Passwords.txt` file, contained in the Recovery Package file.

6. Click **Next**.

Review your configuration and complete installation

You must carefully review the configuration information you have entered to ensure that the installation completes successfully.

Steps

1. View the **Summary** page.

NetApp® StorageGRID® Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 **Summary**

Summary

Verify that all of the grid configuration information is correct, and then click Install. You can view the status of each grid node as it installs. Click the Modify links to go back and change the associated information.

General Settings

Grid Name	Grid1	Modify License
Passwords	Auto-generated random command line passwords	Modify Passwords

Networking

NTP	10.60.248.183 10.227.204.142 10.235.48.111	Modify NTP
DNS	10.224.223.130 10.224.223.136	Modify DNS
Grid Network	172.16.0.0/21	Modify Grid Network

Topology

Topology	Atlanta	Modify Sites	Modify Grid Nodes
	Raleigh		
	dc1-adm1 dc1-g1 dc1-s1 dc1-s2 dc1-s3 NetApp-SGA		

2. Verify that all of the grid configuration information is correct. Use the Modify links on the Summary page to go back and correct any errors.

3. Click **Install**.



If a node is configured to use the Client Network, the default gateway for that node switches from the Grid Network to the Client Network when you click **Install**. If you lose connectivity, you must ensure that you are accessing the primary Admin Node through an accessible subnet. See [Networking guidelines](#) for details.

4. Click **Download Recovery Package**.

When the installation progresses to the point where the grid topology is defined, you are prompted to download the Recovery Package file (.zip), and confirm that you can successfully access the contents of this file. You must download the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fail. The installation continues in the background, but you can't complete the installation and access the StorageGRID system until you download and verify this file.

5. Verify that you can extract the contents of the .zip file, and then save it in two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

6. Select the **I have successfully downloaded and verified the Recovery Package file** checkbox, and click **Next**.

If the installation is still in progress, the status page appears. This page indicates the progress of the installation for each grid node.

Installation Status

If necessary, you may [Download the Recovery Package file](#) again.

Search

Name	IT	Site	IT	Grid Network IPv4 Address	Progress	IT	Stage	IT
dc1-adm1		Site1		172.16.4.215/21			Starting services	
dc1-g1		Site1		172.16.4.216/21			Complete	
dc1-s1		Site1		172.16.4.217/21			Waiting for Dynamic IP Service peers	
dc1-s2		Site1		172.16.4.218/21			Downloading hotfix from primary Admin if needed	
dc1-s3		Site1		172.16.4.219/21			Downloading hotfix from primary Admin if needed	

When the Complete stage is reached for all grid nodes, the sign-in page for the Grid Manager appears.

7. Sign in to the Grid Manager using the "root" user and the password you specified during the installation.

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You can't set up DHCP during configuration.



Nodes reboot when the Grid Network configuration is changed by DHCP, which can cause outages if a DHCP change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to reestablish these connections.

Installation REST API

StorageGRID provides the StorageGRID Installation API for performing installation tasks.

The API uses the Swagger open source API platform to provide the API documentation. Swagger allows both developers and non-developers to interact with the API in a user interface that illustrates how the API responds to parameters and options. This documentation assumes that you are familiar with standard web technologies and the JSON data format.



Any API operations you perform using the API Documentation webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Each REST API command includes the API's URL, an HTTP action, any required or optional URL parameters, and an expected API response.

StorageGRID Installation API

The StorageGRID Installation API is only available when you are initially configuring your StorageGRID system, and if you need to perform a primary Admin Node recovery. The Installation API can be accessed over HTTPS from the Grid Manager.

To access the API documentation, go to the installation web page on the primary Admin Node and select **Help > API documentation** from the menu bar.

The StorageGRID Installation API includes the following sections:

- **config** — Operations related to the product release and versions of the API. You can list the product release version and the major versions of the API supported by that release.
- **grid** — Grid-level configuration operations. You can get and update grid settings, including grid details, Grid Network subnets, grid passwords, and NTP and DNS server IP addresses.
- **nodes** — Node-level configuration operations. You can retrieve a list of grid nodes, delete a grid node, configure a grid node, view a grid node, and reset a grid node's configuration.
- **provision** — Provisioning operations. You can start the provisioning operation and view the status of the provisioning operation.
- **recovery** — Primary Admin Node recovery operations. You can reset information, upload the Recover Package, start the recovery, and view the status of the recovery operation.
- **recovery-package** — Operations to download the Recovery Package.
- **sites** — Site-level configuration operations. You can create, view, delete, and modify a site.
- **temporary-password** — Operations on the temporary password to secure the mgmt-api during installation.

Where to go next

After completing an installation, perform the required integration and configuration tasks. You can perform the optional tasks as needed.

Required tasks

- Configure VMware vSphere Hypervisor for automatic restart.

You must configure the hypervisor to restart the virtual machines when the server restarts. Without an automatic restart, the virtual machines and grid nodes remain shut down after the server restarts. For details, see the VMware vSphere Hypervisor documentation.

- [Create a tenant account](#) for the S3 client protocol that will be used to store objects on your StorageGRID system.
- [Control system access](#) by configuring groups and user accounts. Optionally, you can [configure a federated identity source](#) (such as Active Directory or OpenLDAP), so you can import administration groups and users. Or, you can [create local groups and users](#).
- Integrate and test the [S3 API](#) client applications you will use to upload objects to your StorageGRID system.
- [Configure the information lifecycle management \(ILM\) rules and ILM policy](#) you want to use to protect object data.
- If your installation includes appliance Storage Nodes, use SANtricity OS to complete the following tasks:
 - Connect to each StorageGRID appliance.
 - Verify receipt of AutoSupport data.

See [Set up hardware](#).

- Review and follow the [StorageGRID system hardening guidelines](#) to eliminate security risks.
- [Configure email notifications for system alerts](#).

Optional tasks

- [Update grid node IP addresses](#) if they have changed since you planned your deployment and generated the Recovery Package.
- [Configure storage encryption](#), if required.
- [Configure storage compression](#) to reduce the size of stored objects, if required.
- [Configure VLAN interfaces](#) to isolate and partition network traffic, if required.
- [Configure high availability groups](#) to improve connection availability for the Grid Manager, Tenant Manager, and S3 clients, if required.
- [Configure load balancer endpoints](#) for S3 client connectivity, if required.

Troubleshoot installation issues

If any problems occur while installing your StorageGRID system, you can access the installation log files.

The following are the main installation log files, which technical support might need to resolve issues.

- `/var/local/log/install.log` (found on all grid nodes)
- `/var/local/log/gdu-server.log` (found on the primary Admin Node)

Related information

To learn how to access the log files, see [Log files reference](#).

If you need additional help, contact [NetApp Support](#).

Virtual machine resource reservation requires adjustment

OVF files include a resource reservation designed to ensure that each grid node has sufficient RAM and CPU to operate efficiently. If you create virtual machines by deploying these OVF files on VMware and the predefined number of resources aren't available, the virtual machines will not start.

About this task

If you are certain that the VM host has sufficient resources for each grid node, manually adjust the resources allocated for each virtual machine, and then try starting the virtual machines.

Steps

1. In the VMware vSphere Hypervisor client tree, select the virtual machine that is not started.
2. Right-click the virtual machine, and select **Edit Settings**.
3. From the Virtual Machines Properties window, select the **Resources** tab.
4. Adjust the resources allocated to the virtual machine:
 - a. Select **CPU**, and then use the Reservation slider to adjust the MHz reserved for this virtual machine.
 - b. Select **Memory**, and then use the Reservation slider to adjust the MB reserved for this virtual machine.
5. Click **OK**.
6. Repeat as required for other virtual machines hosted on the same VM host.

Temporary installation password was disabled

When you deploy a VMware node, you can optionally specify a temporary installation password. You must have this password to access the VM console or use SSH before the new node joins the grid.

If you opted to disable the temporary installation password, you must perform additional steps to debug installation issues.

You can do either of the following:

- Redeploy the VM but specify a temporary installation password so you can access the console or use SSH to debug installation issues.
- Use vCenter to set the password:
 1. Power off the VM.
 2. Go to **VM**, select the **Configure** tab, and select **vApp Options**.
 3. Specify type of temporary installation password to set:
 - Select **CUSTOM_TEMPORARY_PASSWORD** to set a custom temporary password.
 - Select **TEMPORARY_PASSWORD_TYPE** to use the node name as the temporary password.
 4. Select **Set Value**.
 5. Set the temporary password:
 - Change **CUSTOM_TEMPORARY_PASSWORD** to a custom password value.
 - Update the **TEMPORARY_PASSWORD_TYPE** with the **Use node name** value.
 6. Restart the VM to apply the new password.

Upgrade StorageGRID software

Upgrade StorageGRID software

Use these instructions to upgrade a StorageGRID system to a new release.

When you perform the upgrade, all nodes in your StorageGRID system are upgraded.

Before you begin

Review these topics to learn about the new features and enhancements in StorageGRID 11.9, determine whether any features have been deprecated or removed, and find out about changes to StorageGRID APIs.

- [What's new in StorageGRID 11.9](#)
- [Removed or deprecated features](#)
- [Changes to the Grid Management API](#)
- [Changes to the Tenant Management API](#)

What's new in StorageGRID 11.9

This release of StorageGRID introduces the following features and functional changes.

Scalability

Data-only Storage Nodes

To allow for more granular scaling, you can now install [data-only Storage Nodes](#). Where metadata processing isn't critical, you can optimize your infrastructure cost-effectively. This flexibility helps accommodate varying workloads and growth patterns.

Cloud Storage Pool enhancements

IAM Roles Anywhere

StorageGRID now supports short term credentials using [IAM Roles Anywhere in Amazon S3 for Cloud Storage Pools](#).

Using long-term credentials to access S3 buckets poses security risks if these credentials are compromised. Short-term credentials have a limited lifespan, which reduces the risk of unauthorized access.

S3 Object Lock buckets

You can now [configure a Cloud Storage Pool using an Amazon S3 endpoint](#). S3 Object Lock helps prevent accidental or malicious deletion of objects. If you tier data from StorageGRID to Amazon S3, having object lock enabled on both systems enhances data protection across the data's lifecycle.

Multi-tenancy

Bucket limits

By [setting limits on S3 buckets](#), you can prevent tenants from monopolizing capacity. Additionally, uncontrolled growth can result in unexpected costs. By having defined limits, you can better estimate tenant storage

expenses.

5,000 buckets per tenant

To enhance scalability, StorageGRID now supports up to [5,000 S3 buckets per tenant](#). Each grid can have a maximum of 100,000 buckets.

To support 5,000 buckets, each Storage Node in the grid must have a minimum of 64 GB of RAM.

S3 Object Lock improvements

Per-tenant configuration capabilities provide the appropriate balance of flexibility and data security. You can now configure per-tenant retention settings to:

- Allow or disallow compliance mode
- Set a maximum retention period

Refer to:

- [Manage objects with S3 Object Lock](#)
- [How grid administrators control object retention](#)
- [Create tenant account](#)

S3 compatibility

x-amz-checksum-sha256 checksum

- The S3 REST API now provides support for `x-amz-checksum-sha256` [checksum](#).
- StorageGRID now provides SHA-256 checksum support for PUT, GET and HEAD operations. These checksums enhance data integrity.

Changes to S3 protocol support

- Added support for Mountpoint for Amazon S3, which allows applications to connect directly to S3 buckets as if they were local file systems. You can now use StorageGRID with more applications and more use cases.
- As part of adding support for Mountpoint, StorageGRID 11.9 contains [additional changes to S3 protocol support](#).

Maintenance and Supportability

AutoSupport

[AutoSupport](#) now automatically creates hardware failure cases for legacy appliances.

Expanded node clone operations

Node clone usability has been expanded to support larger storage nodes.

Improved ILM handling of expired delete markers

ILM ingest time rules with a time period of Days now also remove expired object delete markers. Delete markers are only removed when a time period of Days has passed and the current delete maker has become

expired (there are no non-current versions).

Refer to [How S3 versioned objects are deleted](#) and [Example of bucket lifecycle taking priority over ILM policy](#).

Improved node decommissioning

To provide a smooth and efficient transition to StorageGRID next-generation hardware, [node decommissioning](#) has been improved.

Syslog for load balancer endpoints

Load balancer endpoint access logs contain troubleshooting information, such as HTTP status codes. StorageGRID now supports [exporting these logs to an external syslog server](#). This enhancement allows for more efficient log management and integration with existing monitoring and alerting systems.

Additional enhancements for maintenance and supportability

- Metrics UI update
- New operating system qualifications
- Support for new third-party components

Security

SSH access keys rotation

Grid administrators can now [update and rotate SSH keys](#). The ability to rotate SSH keys is a security best practice and a proactive defense mechanism.

Alerts for root logins

When an unknown entity signs in to the Grid Manager as root, [an alert is triggered](#). Monitoring root SSH logins is a proactive step toward safeguarding your infrastructure.

Grid Manager enhancements

Erasure-coding profiles page moved

The Erasure-coding profiles page is now located at **CONFIGURATION > System > Erasure coding**. It used to be in the ILM menu.

Search enhancements

The [search field in the Grid Manager](#) now includes better matching logic, allowing you to find pages by searching for common abbreviations and by the names of certain settings within a page. You can also search for more types of items, like nodes, users, and tenant accounts.

Removed or deprecated features and capabilities

Some features and capabilities were removed or deprecated in this release. Review these items to understand whether you need to update client applications or modify your configuration before you upgrade.

Definitions

Deprecated

The feature **should not** be used in new production environments. Existing production environments can continue using the feature.

End of Life

Last shipped version that supports the feature. In some cases, documentation for the feature might be removed at this stage.

Removed

First version that **does not** support the feature.

StorageGRID end of feature support

Deprecated features will be removed in N+2 major versions. For example, if a feature is deprecated in version N (for example, 6.3), the last version where the feature will exist is N+1 (for example, 6.4). Version N+2 (for example, 6.5) is the first release when the feature doesn't exist in the product.

See the [Software Version Support page](#) for additional information.



In certain situations, NetApp might end support for particular features sooner than indicated.

Feature	Deprecated	End of Life	Removed	Links to earlier documentation
Legacy Alarms (<i>not Alerts</i>)	11.7	11.8	11.9	Alarms reference (StorageGRID 11.8)
Archive Node support	11.7	11.8	11.9	Considerations for decommissioning Archive Nodes (StorageGRID 11.8) Note: Before starting your upgrade, you must: - Decommission all Archive Nodes. See Grid node decommissioning (StorageGRID 11.8 doc site). - Remove all Archive Node references from storage pools and ILM policies. See NetApp Knowledge Base: StorageGRID 11.9 software upgrade resolution guide.
Audit export through CIFS/Samba	11.1	11.6	11.7	
CLB service	11.4	11.6	11.7	

Feature	Deprecated	End of Life	Removed	Links to earlier documentation
Docker container engine	11.8	11.9	TBD	Support for Docker as the container engine for software-only deployments is deprecated. Docker will be replaced with another container engine in a future release. Refer to the list of Docker versions currently supported .
NFS audit export	11.8	11.9	12.0	Configure audit client access for NFS (StorageGRID 11.8)
Swift API support	11.7	11.9	12.0	Use Swift REST API (StorageGRID 11.8)
RHEL 8.8	11.9	11.9	12.0	
RHEL 9.0	11.9	11.9	12.0	
RHEL 9.2	11.9	11.9	12.0	
Ubuntu 18.04	11.9	11.9	12.0	
Ubuntu 20.04	11.9	11.9	12.0	
Debian 11	11.9	11.9	12.0	

Also refer to:

- [Changes to the Grid Management API](#)
- [Changes to the Tenant Management API](#)

Changes to the Grid Management API

StorageGRID 11.9 uses version 4 of the Grid Management API. Version 4 deprecates version 3; however, versions 1, 2, and 3 are still supported.



You can continue to use deprecated versions of the management API with StorageGRID 11.9; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.9, you can deactivate the deprecated APIs by using the `PUT /grid/config/management` API.

To learn more, go to [Use the Grid Management API](#).

Review compliance settings after enabling global S3 Object Lock

Review the compliance settings of existing tenants after you enable the global S3 Object Lock setting. When you enable this setting, the S3 Object Lock per-tenant settings depend on the StorageGRID release at the time the tenant was created.

Legacy mgmt-api requests removed

These legacy requests have been removed:

`/grid/server-types`

`/grid/ntp-roles`

Changes to GET `/private/storage-usage` API

- A new property, `usageCacheDuration`, has been added to the response body. This property specifies the duration (in seconds) for which the usage lookup cache remains valid. This value applies when checking the usage against tenant storage quota and bucket capacity limits.
- The GET `/api/v4/private/storage-usage` behavior has been corrected to match nesting from the schema.
- These changes apply only to the private API.

Changes to GET `cross-grid-replication` API

The `/org/containers/:name/cross-grid-replication` GET API no longer requires the Root access (`rootAccess`) permission; however, you must belong to a user group that has the Manage all buckets (`manageAllContainers`) or View all buckets (`viewAllContainers`) permission.

The `/org/containers/:name/cross-grid-replication` PUT API is unchanged and still requires the Root access (`rootAccess`) permission.

Changes to the Tenant Management API

StorageGRID 11.9 uses version 4 of the Tenant Management API. Version 4 deprecates version 3; however, versions 1, 2, and 3 are still supported.



You can continue to use deprecated versions of the Tenant Management API with StorageGRID 11.9; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.9, you can deactivate the deprecated APIs by using the PUT `/grid/config/management` API.

To learn more, go to [Understand the Tenant Management API](#).

New API for bucket capacity limit

You can use the `/org/containers/{bucketName}/quota-object-bytes` API with GET/PUT operations to get and set the storage capacity limit for a bucket.

Plan and prepare for upgrade

Estimate the time to complete an upgrade

Consider when to upgrade, based on how long the upgrade might take. Be aware of which operations you can and can't perform during each stage of the upgrade.

About this task

The time required to complete a StorageGRID upgrade depends on a variety of factors such as client load and hardware performance.

The table summarizes the main upgrade tasks and lists the approximate time required for each task. The steps after the table provide instructions you can use to estimate the upgrade time for your system.

Upgrade task	Description	Approximate time required	During this task
Run prechecks and upgrade primary Admin Node	The upgrade prechecks are run, and the primary Admin Node is stopped, upgraded, and restarted.	30 minutes to 1 hour, with services appliance nodes requiring the most time. Unresolved precheck errors will increase this time.	You can't access the primary Admin Node. Connection errors might be reported, which you can ignore. Running the upgrade prechecks before starting the upgrade lets you resolve any errors before the scheduled upgrade maintenance window.
Start upgrade service	The software file is distributed, and the upgrade service is started.	3 minutes per grid node	
Upgrade other grid nodes	The software on all other grid nodes is upgraded, in the order in which you approve the nodes. Every node in your system will be brought down one at a time.	15 minutes to 1 hour per node, with appliance nodes requiring the most time Note: For appliance nodes, the StorageGRID Appliance Installer is automatically updated to the latest release.	• Don't change the grid configuration. • Don't change the audit level configuration. • Don't update the ILM configuration. • You are prevented from performing other maintenance procedures, such as hotfix, decommission, or expansion. Note: If you need to perform a recovery, contact technical support.
Enable features	The new features for the new version are enabled.	Less than 5 minutes	• Don't change the grid configuration. • Don't change the audit level configuration. • Don't update the ILM configuration. • You can't perform another maintenance procedure.
Upgrade database	The upgrade process checks each node to verify that the Cassandra database does not need to be updated.	10 seconds per node or a few minutes for the entire grid	The upgrade from StorageGRID 11.8 to 11.9 does not require a Cassandra database upgrade; however, the Cassandra service will be stopped and restarted on each Storage Node. For future StorageGRID feature releases, the Cassandra database update step might take several days to complete.

Upgrade task	Description	Approximate time required	During this task
Final upgrade steps	Temporary files are removed and the upgrade to the new release completes.	5 minutes	When the Final upgrade steps task completes, you can perform all maintenance procedures.

Steps

1. Estimate the time required to upgrade all grid nodes.
 - a. Multiply the number of nodes in your StorageGRID system by 1 hour/node.

As a general rule, appliance nodes take longer to upgrade than software-based nodes.

 - b. Add 1 hour to this time to account for the time required to download the `.upgrade` file, run precheck validations, and complete the final upgrade steps.
2. If you have Linux nodes, add 15 minutes for each node to account for the time required to download and install the RPM or DEB package.
3. Calculate the total estimated time for the upgrade by adding the results of steps 1 and 2.

Example: Estimated time to upgrade to StorageGRID 11.9

Suppose your system has 14 grid nodes, of which 8 are Linux nodes.

1. Multiply 14 by 1 hour/node.
2. Add 1 hour to account for the download, precheck, and final steps.

The estimated time to upgrade all nodes is 15 hours.

3. Multiply 8 by 15 minutes/node to account for the time to install the RPM or DEB package on the Linux nodes.

The estimated time for this step is 2 hours.

4. Add the values together.

You should allow up to 17 hours to complete the upgrade of your system to StorageGRID 11.9.0.



As required, you can split the maintenance window into smaller windows by approving subsets of grid nodes to upgrade in multiple sessions. For example, you might prefer to upgrade the nodes at site A in one session and then upgrade the nodes at site B in a later session. If you choose to perform the upgrade in more than one session, be aware that you can't start using the new features until all nodes have been upgraded.

How your system is affected during the upgrade

Learn how your StorageGRID system will be affected during upgrade.

StorageGRID upgrades are non-disruptive

The StorageGRID system can ingest and retrieve data from client applications throughout the upgrade process. If you approve all nodes of the same type to upgrade (for example, Storage Nodes), the nodes are brought down one at a time, so there is no time when all grid nodes or all grid nodes of a certain type are unavailable.

To allow for continued availability, ensure that your ILM policy contains rules that specify storing multiple copies of each object. You must also ensure that all external S3 clients are configured to send requests to one of the following:

- A high availability (HA) group virtual IP address
- A high availability third-party load balancer
- Multiple Gateway Nodes for each client
- Multiple Storage Nodes for each client

Client applications might experience short-term disruptions

The StorageGRID system can ingest and retrieve data from client applications throughout the upgrade process; however, client connections to individual Gateway Nodes or Storage Nodes might be disrupted temporarily if the upgrade needs to restart services on those nodes. Connectivity will be restored after the upgrade process completes and services resume on the individual nodes.

You might need to schedule downtime to apply an upgrade if loss of connectivity for a short period is not acceptable. You can use selective approval to schedule when certain nodes are updated.



You can use multiple gateways and high availability (HA) groups to provide automatic failover during the upgrade process. See the instructions for [configuring high availability groups](#).

Appliance firmware is upgraded

During the StorageGRID 11.9 upgrade:

- All StorageGRID appliance nodes are automatically upgraded to StorageGRID Appliance Installer firmware version 3.9.
- SG6060 and SGF6024 appliances are automatically upgraded to BIOS firmware version 3B08.EX and BMC firmware version 4.00.07.
- SG100 and SG1000 appliances are automatically upgraded to BIOS firmware version 3B13.EC and BMC firmware version 4.74.07.
- SGF6112, SG6160, SG110 and SG1100 appliances are automatically upgraded to BMC firmware version 3.16.07.

ILM policies are handled differently according to their status

- The active policy will remain the same after upgrade.
- Only the latest 10 historical policies are preserved on upgrade.
- If there is a proposed policy, it will be deleted during upgrade.

Alerts might be triggered

Alerts might be triggered when services start and stop and when the StorageGRID system is operating as a

mixed-version environment (some grid nodes running an earlier version, while others have been upgraded to a later version). Other alerts might be triggered after the upgrade completes.

For example, you might see the **Unable to communicate with node** alert when services are stopped, or you might see the **Cassandra communication error** alert when some nodes have been upgraded to StorageGRID 11.9 but other nodes are still running StorageGRID 11.8. In general, these alerts will clear when the upgrade completes.

The **ILM placement unachievable** alert might be triggered when Storage Nodes are stopped during the upgrade to StorageGRID 11.9. This alert might persist for 1 day after the upgrade completes.

After the upgrade completes, you can review any upgrade-related alerts by selecting **Recently resolved alerts** or **Current alerts** from the Grid Manager dashboard.

Many SNMP notifications are generated

Be aware that a large number of SNMP notifications might be generated when grid nodes are stopped and restarted during the upgrade. To avoid excessive notifications, clear the **Enable SNMP Agent Notifications** checkbox (**CONFIGURATION > Monitoring > SNMP agent**) to disable SNMP notifications before you start the upgrade. Then, re-enable notifications after the upgrade is complete.

Configuration changes are restricted



This list applies specifically to upgrades from StorageGRID 11.8 to StorageGRID 11.9. If you're upgrading to another StorageGRID release, refer to the list of restricted changes in the upgrade instructions for that release.

Until the **Enable New Feature** task completes:

- Don't make any grid configuration changes.
- Don't enable or disable any new features.
- Don't update the ILM configuration. Otherwise, you might experience inconsistent and unexpected ILM behavior.
- Don't apply a hotfix or recover a grid node.



Contact technical support if you need to recover a node during upgrade.

- You should not manage HA groups, VLAN interfaces, or load balancer endpoints while you're upgrading to StorageGRID 11.9.
- Don't delete any HA groups until the upgrade to StorageGRID 11.9 is complete. Virtual IP addresses in other HA groups might become inaccessible.

Until the **Final Upgrade Steps** task completes:

- Don't perform an expansion procedure.
- Don't perform a decommission procedure.

You can't view bucket details or manage buckets from the Tenant Manager

During the upgrade to StorageGRID 11.9 (that is, while the system is operating as a mixed-version environment), you can't view bucket details or manage buckets using the Tenant Manager. One of the following errors appears on the Buckets page in Tenant Manager:

- You can't use this API while you're upgrading to 11.9.
- You can't view bucket versioning details in the Tenant Manager while you're upgrading to 11.9.

This error will resolve after the upgrade to 11.9 is complete.

Workaround

While the 11.9 upgrade is in progress, use the following tools to view bucket details or manage buckets, instead of using the Tenant Manager:

- To perform standard S3 operations on a bucket, use either the [S3 REST API](#) or the [Tenant Management API](#).
- To perform StorageGRID custom operations on a bucket (for example, viewing and modifying the bucket consistency, enabling or disabling last access time updates, or configuring search integration), use the Tenant Management API.

Verify the installed version of StorageGRID

Before starting the upgrade, verify that the previous version of StorageGRID is currently installed with the latest available hotfix applied.

About this task

Before you upgrade to StorageGRID 11.9, your grid must have StorageGRID 11.8 installed. If you are currently using a previous version of StorageGRID, you must install all previous upgrade files along with their latest hotfixes (strongly recommended) until your grid's current version is StorageGRID 11.8.x.y.

One possible upgrade path is shown in the [example](#).



NetApp strongly recommends that you apply the latest hotfix for each StorageGRID version before upgrading to the next version and that you also apply the latest hotfix for each new version you install. In some cases, you must apply a hotfix to avoid the risk of data loss. See [NetApp Downloads: StorageGRID](#) and the release notes for each hotfix to learn more.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. From the top of the Grid Manager, select **Help > About**.
3. Verify that **Version** is 11.8.x.y.

In the StorageGRID 11.8.x.y version number:

- The **major release** has an x value of 0 (11.8.0).
 - A **hotfix**, if one has been applied, has a y value (for example, 11.8.0.1).
4. If **Version** is not 11.8.x.y, go to [NetApp Downloads: StorageGRID](#) to download the files for each previous release, including the latest hotfix for each release.
 5. Obtain the the upgrade instructions for each release you downloaded. Then, perform the software upgrade procedure for that release, and apply the latest hotfix for that release (strongly recommended).

See the [StorageGRID hotfix procedure](#).

Example: Upgrade to StorageGRID 11.9 from version 11.6

The following example shows the steps to upgrade from StorageGRID version 11.6 to version 11.8 in preparation for a StorageGRID 11.9 upgrade.

Download and install software in the following sequence to prepare your system for upgrade:

1. Upgrade to the StorageGRID 11.6.0 major release.
2. Apply the latest StorageGRID 11.6.0.y hotfix.
3. Upgrade to the StorageGRID 11.7.0 major release.
4. Apply the latest StorageGRID 11.7.0.y hotfix.
5. Upgrade to the StorageGRID 11.8.0 major release.
6. Apply the latest StorageGRID 11.8.0.y hotfix.

Obtain the required materials for a software upgrade

Before you begin the software upgrade, obtain all required materials.

Item	Notes
Service laptop	The service laptop must have: • Network port • SSH client (for example, PuTTY)
Supported web browser	Browser support typically changes for each StorageGRID release. Make sure your browser is compatible with the new StorageGRID version.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not listed in the <code>Passwords.txt</code> file.
Linux RPM or DEB archive	If any nodes are deployed on Linux hosts, you must download and install the RPM or DEB package on all hosts before you start the upgrade. Ensure your operating system meets StorageGRID's minimum kernel version requirements: • Install StorageGRID on Red Hat Enterprise Linux hosts • Install StorageGRID on Ubuntu or Debian hosts
StorageGRID documentation	• Release notes for StorageGRID 11.9 (sign in required). Be sure to read these carefully before starting the upgrade. • StorageGRID software upgrade resolution guide for the major version you are upgrading to (sign in required) • Other StorageGRID documentation, as required.

Check the system’s condition

Before upgrading a StorageGRID system, verify the system is ready to accommodate the upgrade. Ensure that the system is running normally and that all grid nodes are operational.

Steps

- 1. Sign in to the Grid Manager using a [supported web browser](#).
- 2. Check for and resolve any active alerts.
- 3. Confirm that no conflicting grid tasks are active or pending.
 - a. Select **SUPPORT > Tools > Grid topology**.
 - b. Select **site > primary Admin Node > CMN > Grid Tasks > Configuration**.

Information lifecycle management evaluation (ILME) tasks are the only grid tasks that can run concurrently with the software upgrade.

- c. If any other grid tasks are active or pending, wait for them to finish or release their lock.



Contact technical support if a task does not finish or release its lock.

- 4. Refer to [Internal grid node communications](#) and [External communications](#) to ensure that all required ports for StorageGRID 11.9 are opened before you upgrade.



No additional ports are required when upgrading to StorageGRID 11.9.

The following required port was added in StorageGRID 11.7. Make sure it’s available before you upgrade to StorageGRID 11.9.

Port	Description
18086	TCP port used for S3 requests from the StorageGRID load balancer to LDR and the new LDR service. Before upgrading, confirm that this port is open from all grid nodes to all Storage Nodes. Blocking this port will cause S3 service interruptions after upgrade to StorageGRID 11.9.



If you have opened any custom firewall ports, you are notified during the upgrade precheck. You must contact technical support before proceeding with the upgrade.

Upgrade software

Upgrade quick start

Before starting the upgrade, review the general workflow. The StorageGRID Upgrade page guides you through each upgrade step.

1

Prepare Linux hosts

If any StorageGRID nodes are deployed on Linux hosts, [install the RPM or DEB package on each host](#) before you start the upgrade.

2

Upload upgrade and hotfix files

From the primary Admin Node, access the StorageGRID Upgrade page and upload the upgrade file and the hotfix file, if required.

3

Download Recovery Package

Download the current Recovery Package before you start the upgrade.

4

Run upgrade prechecks

Upgrade prechecks help you detect issues, so you can resolve them before you start the actual upgrade.

5

Start upgrade

When you start the upgrade, the prechecks are run again and the primary Admin Node is upgraded automatically. You can't access the Grid Manager while the primary Admin Node is being upgraded. Audit logs will also be unavailable. This upgrade can take up to 30 minutes.

6

Download Recovery Package

After the primary Admin Node has been upgraded, download a new Recovery Package.

7

Approve nodes

You can approve individual grid nodes, groups of grid nodes, or all grid nodes.



Don't approve the upgrade for a grid node unless you are sure that node is ready to be stopped and rebooted.

8

Resume operations

When all grid nodes have been upgraded, new features are enabled and you can resume operations. You must wait to perform a decommission or expansion procedure until the background **Upgrade database** task and the **Final upgrade steps** task have completed.

Related information

[Estimate the time to complete an upgrade](#)

Linux: Download and install the RPM or DEB package on all hosts

If any StorageGRID nodes are deployed on Linux hosts, download and install an additional RPM or DEB package on each of these hosts before you start the upgrade.

Download upgrade, Linux, and hotfix files

When you perform a StorageGRID upgrade from the Grid Manager, you are prompted to download the upgrade archive and any required hotfix as the first step. However, if you need to download files to upgrade Linux hosts, you can save time by downloading all required files in advance.

Steps

1. Go to [NetApp Downloads: StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.

StorageGRID software versions have this format: 11.x.y. StorageGRID hotfixes have this format: 11.x.y.z.

3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead notice appears, make note of the hotfix number, and select the checkbox.
5. Read the End User License Agreement (EULA), select the checkbox, and then select **Accept & Continue**.

The downloads page for the version you selected appears. The page contains three columns.

6. From the second column (**Upgrade StorageGRID**), download two files:
 - The upgrade archive for the latest release (this is the file in the section labeled **VMware, SG1000, or SG100 Primary Admin Node**). While this file is not needed until you perform the upgrade, downloading it now will save time.
 - An RPM or DEB archive in either .tgz or .zip format. Select the .zip file if you are running Windows on the service laptop.
 - Red Hat Enterprise Linux
`StorageGRID-Webscale-version-RPM-uniqueID.zip`
`StorageGRID-Webscale-version-RPM-uniqueID.tgz`
 - Ubuntu or Debian
`StorageGRID-Webscale-version-DEB-uniqueID.zip`
`StorageGRID-Webscale-version-DEB-uniqueID.tgz`
7. If you needed to agree to a Caution/MustRead notice because of a required hotfix, download the hotfix:
 - a. Go back to [NetApp Downloads: StorageGRID](#).
 - b. Select the hotfix number from the drop-down.
 - c. Agree to the Caution notice and EULA again.
 - d. Download and save the hotfix and its README.

You will be prompted to upload the hotfix file on the StorageGRID Upgrade page when you start the upgrade.

Install archive on all Linux hosts

Perform these steps before upgrading StorageGRID software.

Steps

1. Extract the RPM or DEB packages from the installation file.
2. Install the RPM or DEB packages on all Linux hosts.

See the steps for installing StorageGRID host services in the installation instructions:

- [Red Hat Enterprise Linux: Install StorageGRID host services](#)
- [Ubuntu or Debian: Install StorageGRID host services](#)

The new packages are installed as additional packages.

Remove installation archives for previous versions

To free up space on Linux hosts, you can remove the installation archives for previous versions of StorageGRID that you no longer need.

Steps

1. Remove the old StorageGRID installation archives.

Red Hat

- a. Capture the list of StorageGRID packages installed: `dnf list | grep -i storagegrid`.

Example:

```
[root@rhel-example ~]# dnf list | grep -i storagegrid
StorageGRID-Webscale-Images-11-6-0.x86_64 11.6.0-
20220210.0232.8d56cfe @System
StorageGRID-Webscale-Images-11-7-0.x86_64 11.7.0-
20230424.2238.1a2cf8c @System
StorageGRID-Webscale-Images-11-8-0.x86_64 11.8.0-
20240131.0139.e3e0c87 @System
StorageGRID-Webscale-Images-11-9-0.x86_64 11.9.0-
20240826.1753.4aeeb70 @System
StorageGRID-Webscale-Service-11-6-0.x86_64 11.6.0-
20220210.0232.8d56cfe @System
StorageGRID-Webscale-Service-11-7-0.x86_64 11.7.0-
20230424.2238.1a2cf8c @System
StorageGRID-Webscale-Service-11-8-0.x86_64 11.8.0-
20240131.0139.e3e0c87 @System
StorageGRID-Webscale-Service-11-9-0.x86_64 11.9.0-
20240826.1753.4aeeb70 @System
[root@rhel-example ~]#
```

- b. Remove previous StorageGRID packages: `dnf remove images-package service-package`



Do not remove the installation archives for the version of StorageGRID you are currently running or the versions of StorageGRID you are planning to upgrade to.

You can safely ignore the warnings that appear. They refer to files that have been replaced when you install newer StorageGRID packages.

Example:

```
[root@rhel-example ~]# dnf remove StorageGRID-Webscale-Images-11-6-
0.x86_64 StorageGRID-Webscale-Service-11-6-0.x86_64
Updating Subscription Management repositories.
Unable to read consumer identity

This system is not registered with an entitlement server. You can
use subscription-manager to register.

Dependencies resolved.
=====
=====
```


Package Size	Architecture	Version	Repository
=====			
=====			
Removing:			
StorageGRID-Webscale-Images-11-6-0	x86_64	11.6.0-	
20220210.0232.8d56cfe	@System	2.7	G
StorageGRID-Webscale-Service-11-6-0	x86_64	11.6.0-	
20220210.0232.8d56cfe	@System	7.5	M
Transaction Summary			
=====			
=====			
Remove 2 Packages			
Freed space: 2.8 G			
Is this ok [y/N]: y			
Running transaction check			
Transaction check succeeded.			
Running transaction test			
Transaction test succeeded.			
Running transaction			
Preparing: 1/1			
Running scriptlet: StorageGRID-Webscale-Service-11-6-0-11.6.0-			
20220210.0232.8d56cfe.x86_64 1/2			
Erasing: StorageGRID-Webscale-Service-11-6-0-11.6.0-			
20220210.0232.8d56cfe.x86_64 1/2			
warning: file /usr/lib64/python2.7/site-			
packages/netapp/storagegrid/vendor/latest/netaddr/strategy/ipv6.pyc:			
remove failed: No such file or directory			
warning: file /usr/lib64/python2.7/site-			
packages/netapp/storagegrid/vendor/latest/netaddr/strategy/ipv4.pyc:			
remove failed: No such file or directory			
warning: file /usr/lib64/python2.7/site-			
packages/netapp/storagegrid/vendor/latest/netaddr/strategy/eui64.pyc			
: remove failed: No such file or directory			
warning: file /usr/lib64/python2.7/site-			
packages/netapp/storagegrid/vendor/latest/netaddr/strategy/eui48.pyc			
: remove failed: No such file or directory			
warning: file /usr/lib64/python2.7/site-			
packages/netapp/storagegrid/vendor/latest/netaddr/strategy/__init__.			
pyc: remove failed: No such file or directory			
warning: file /usr/lib64/python2.7/site-			
packages/netapp/storagegrid/vendor/latest/netaddr/ip/sets.pyc:			
remove failed: No such file or directory			
warning: file /usr/lib64/python2.7/site-			

```

packages/netapp/storagegrid/vendor/latest/netaddr/ip/rfc1924.pyc:
remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/ip/nmap.pyc:
remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/ip/iana.pyc:
remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/ip/glob.pyc:
remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/ip/__init__.pyc:
remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/fbsocket.pyc:
remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/eui/ieee.pyc:
remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/eui/__init__.pyc:
remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/core.pyc: remove
failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/contrib/subnet_spl
itter.pyc: remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/contrib/__init__.p
yc: remove failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/compat.pyc: remove
failed: No such file or directory
warning: file /usr/lib64/python2.7/site-
packages/netapp/storagegrid/vendor/latest/netaddr/__init__.pyc:
remove failed: No such file or directory

```

```

Erasing: StorageGRID-Webscale-Images-11-6-0-11.6.0-
20220210.0232.8d56cfe.x86_64 2/2

```

```

Verifying: StorageGRID-Webscale-Images-11-6-0-11.6.0-
20220210.0232.8d56cfe.x86_64 1/2

```

```

Verifying: StorageGRID-Webscale-Service-11-6-0-11.6.0-
20220210.0232.8d56cfe.x86_64 2/2

```

```

Installed products updated.

```

Removed:

```
StorageGRID-Webscale-Images-11-6-0-11.6.0-  
20220210.0232.8d56cfe.x86_64  
StorageGRID-Webscale-Service-11-6-0-11.6.0-  
20220210.0232.8d56cfe.x86_64
```

Complete!

```
[root@rhel-example ~]#
```

Ubuntu and Debian

- a. Capture the list of StorageGRID packages installed: `dpkg -l | grep storagegrid`

Example:

```
root@debian-example:~# dpkg -l | grep storagegrid  
ii storagegrid-webscale-images-11-6-0 11.6.0-20220210.0232.8d56cfe  
amd64 StorageGRID Webscale docker images for 11.6.0  
ii storagegrid-webscale-images-11-7-0 11.7.0-  
20230424.2238.1a2cf8c.dev-signed amd64 StorageGRID Webscale docker  
images for 11.7.0  
ii storagegrid-webscale-images-11-8-0 11.8.0-20240131.0139.e3e0c87  
amd64 StorageGRID Webscale docker images for 11.8.0  
ii storagegrid-webscale-images-11-9-0 11.9.0-20240826.1753.4aeeb70  
amd64 StorageGRID Webscale docker images for 11.9.0  
ii storagegrid-webscale-service-11-6-0 11.6.0-20220210.0232.8d56cfe  
amd64 StorageGRID Webscale host services for 11.6.0  
ii storagegrid-webscale-service-11-7-0 11.7.0-20230424.2238.1a2cf8c  
amd64 StorageGRID Webscale host services for 11.7.0  
ii storagegrid-webscale-service-11-8-0 11.8.0-20240131.0139.e3e0c87  
amd64 StorageGRID Webscale host services for 11.8.0  
ii storagegrid-webscale-service-11-9-0 11.9.0-20240826.1753.4aeeb70  
amd64 StorageGRID Webscale host services for 11.9.0  
root@debian-example:~#
```

- b. Remove previous StorageGRID packages: `dpkg -r images-package service-package`



Do not remove the installation archives for the version of StorageGRID you are currently running or the versions of StorageGRID you are planning to upgrade to.

Example:

```
root@debian-example:~# dpkg -r storagegrid-webscale-service-11-6-0
storagegrid-webscale-images-11-6-0
(Reading database ... 38190 files and directories currently
installed.)
Removing storagegrid-webscale-service-11-6-0 (11.6.0-
20220210.0232.8d56cfe) ...
locale: Cannot set LC_CTYPE to default locale: No such file or
directory
locale: Cannot set LC_MESSAGES to default locale: No such file or
directory
locale: Cannot set LC_ALL to default locale: No such file or
directory
dpkg: warning: while removing storagegrid-webscale-service-11-6-0,
directory '/usr/lib/python2.7/dist-
packages/netapp/storagegrid/vendor/latest' not empty so not removed
Removing storagegrid-webscale-images-11-6-0 (11.6.0-
20220210.0232.8d56cfe) ...
root@debian-example:~#
```

2. Remove StorageGRID container images.

Docker

- a. Capture the list of container images installed: `docker images`

Example:

```
[root@docker-example ~]# docker images
REPOSITORY          TAG                 IMAGE ID            CREATED
SIZE
storagegrid-11.9.0   Admin_Node         610f2595bcb4       2 days ago
2.77GB
storagegrid-11.9.0   Storage_Node       7f73d33eb880       2 days ago
2.65GB
storagegrid-11.9.0   API_Gateway        2f0bb79526e9       2 days ago
1.82GB
storagegrid-11.8.0   Storage_Node       7125480de71b       7 months ago
2.54GB
storagegrid-11.8.0   Admin_Node         404e9f1bd173       7 months ago
2.63GB
storagegrid-11.8.0   Archive_Node       c3294a29697c       7 months ago
2.39GB
storagegrid-11.8.0   API_Gateway        1f88f24b9098       7 months ago
1.74GB
storagegrid-11.7.0   Storage_Node       1655350eff6f       16 months ago
2.51GB
storagegrid-11.7.0   Admin_Node         872258dd0dc8       16 months ago
2.48GB
storagegrid-11.7.0   Archive_Node       121e7c8b6d3b       16 months ago
2.41GB
storagegrid-11.7.0   API_Gateway        5b7a26e382de       16 months ago
1.77GB
storagegrid-11.6.0   Admin_Node         ee39f71a73e1       2 years ago
2.38GB
storagegrid-11.6.0   Storage_Node       f5ef895dcad0       2 years ago
2.08GB
storagegrid-11.6.0   Archive_Node       5782de552db0       2 years ago
1.95GB
storagegrid-11.6.0   API_Gateway        cb480ed37eea       2 years ago
1.35GB
[root@docker-example ~]#
```

- b. Remove the container images for previous StorageGRID versions: `docker rmi image id`



Do not remove the container images for the version of StorageGRID you are currently running or the versions of StorageGRID you are planning to upgrade to.

Example:

```
[root@docker-example ~]# docker rmi cb480ed37eea
Untagged: storagegrid-11.6.0:API_Gateway
Deleted:
sha256:cb480ed37eea0ae9cf3522de1dadfbff0075010d89c1c0a2337a3178051ddf02
Deleted:
sha256:5f269aabf15c32c1fe6f36329c304b6c6ecb563d973794b9b59e8e5ab8ccc
afa
Deleted:
sha256:47c2b2c295a77b312b8db69db58a02d8e09e929e121352bec713fa12dae66
bde
[root@docker-example ~]#
```

Podman

- a. Capture the list of container images installed: `podman images`

Example:

```
[root@podman-example ~]# podman images
REPOSITORY                                TAG          IMAGE ID      CREATED
SIZE
localhost/storagegrid-11.8.0             Storage_Node  7125480de71b  7 months
ago    2.57 GB
localhost/storagegrid-11.8.0             Admin_Node   404e9f1bd173  7 months
ago    2.67 GB
localhost/storagegrid-11.8.0             Archive_Node c3294a29697c  7 months
ago    2.42 GB
localhost/storagegrid-11.8.0             API_Gateway  1f88f24b9098  7 months
ago    1.77 GB
localhost/storagegrid-11.7.0             Storage_Node 1655350eff6f  16 months
ago    2.54 GB
localhost/storagegrid-11.7.0             Admin_Node   872258dd0dc8  16 months
ago    2.51 GB
localhost/storagegrid-11.7.0             Archive_Node 121e7c8b6d3b  16 months
ago    2.44 GB
localhost/storagegrid-11.7.0             API_Gateway  5b7a26e382de  16 months
ago    1.8 GB
localhost/storagegrid-11.6.0             Admin_Node   ee39f71a73e1  2 years
ago    2.42 GB
localhost/storagegrid-11.6.0             Storage_Node f5ef895dcad0  2 years
ago    2.11 GB
localhost/storagegrid-11.6.0             Archive_Node 5782de552db0  2 years
ago    1.98 GB
localhost/storagegrid-11.6.0             API_Gateway  cb480ed37eea  2 years
ago    1.38 GB
[root@podman-example ~]#
```

b. Remove the container images for previous StorageGRID versions: `podman rmi image id`



Do not remove the container images for the version of StorageGRID you are currently running or the versions of StorageGRID you are planning to upgrade to.

Example:

```
[root@podman-example ~]# podman rmi f5ef895dcad0
Untagged: localhost/storagegrid-11.6.0:Storage_Node
Deleted:
f5ef895dcad0d78d0fd21a07dd132d7c7f65f45d80ee7205a4d615494e44cbb7
[root@podman-example ~]#
```

Perform the upgrade

You can upgrade to StorageGRID 11.9 and apply the latest hotfix for that release at the same time. The StorageGRID upgrade page provides the recommended upgrade path and links directly to the correct download pages.

Before you begin

You have reviewed all of the considerations and completed all of the planning and preparation steps.

Access StorageGRID Upgrade page

As a first step, access the StorageGRID Upgrade page in the Grid Manager.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. Select **MAINTENANCE** > **System** > **Software update**.
3. From the StorageGRID upgrade tile, select **Upgrade**.

Select files

The update path on the StorageGRID Upgrade page indicates which major versions (for example, 11.9.0) and hotfixes (for example, 11.9.0.1) you must install to get to the latest StorageGRID release. You should install the recommended versions and hotfixes in the order shown.



If no update path is shown, your browser might not be able to access the NetApp Support Site, or the **Check for software updates** checkbox on the AutoSupport page (**SUPPORT** > **Tools** > **AutoSupport** > **Settings**) might be disabled.

Steps

1. For the **Select files** step, review the update path.
2. From the Download files section, select each **Download** link to download the required files from the NetApp Support Site.

If no update path is shown, go to the [NetApp Downloads: StorageGRID](#) to determine if a new version or hotfix is available and to download the files you need.



If you needed to download and install an RPM or DEB package on all Linux hosts, you might already have the StorageGRID upgrade and hotfix files listed in the update path.

3. Select **Browse** to upload the version upgrade file to StorageGRID:
`NetApp_StorageGRID_11.9.0_Software_uniqueID.upgrade`

When the upload and validation process is done, a green check mark appears next to the file name.

4. If you downloaded a hotfix file, select **Browse** to upload that file. The hotfix will be automatically applied as part of the version upgrade.
5. Select **Continue**.

Run prechecks

Running prechecks allows you to detect and resolve any upgrade issues before you start upgrading your grid.

Steps

1. For the **Run prechecks** step, start by entering the provisioning passphrase for your grid.
2. Select **Download recovery package**.

You should download the current copy of the Recovery Package file before you upgrade the primary Admin Node. The Recovery Package file allows you to restore the system if a failure occurs.

3. When the file is downloaded, confirm that you can access the contents, including the `Passwords.txt` file.
4. Copy the downloaded file (`.zip`) to two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

5. Select **Run prechecks**, and wait for the prechecks to complete.
6. Review the details for each reported precheck and resolve any reported errors. See the [StorageGRID software upgrade resolution guide](#) for the StorageGRID 11.9 release.

You must resolve all precheck *errors* before you can upgrade your system. However, you don't need to address precheck *warnings* before upgrading.



If you have opened any custom firewall ports, you are notified during the precheck validation. You must contact technical support before proceeding with the upgrade.

7. If you made any configuration changes to resolve the reported issues, select **Run prechecks** again to get updated results.

If all errors have been resolved, you are prompted to start the upgrade.

Start upgrade and upgrade primary Admin Node

When you start the upgrade, the upgrade prechecks are run again, and the primary Admin Node is automatically upgraded. This part of the upgrade can take up to 30 minutes.



You will not be able to access any other Grid Manager pages while the primary Admin Node is being upgraded. Audit logs will also be unavailable.

Steps

1. Select **Start upgrade**.

A warning appears to remind you will temporarily lose access to the Grid Manager.

2. Select **OK** to acknowledge the warning and start the upgrade.
3. Wait for the upgrade prechecks to be performed and for the primary Admin Node to be upgraded.



If any precheck errors are reported, resolve them and select **Start upgrade** again.

If the grid has another Admin Node that is online and ready, you can use it to monitor the status of the primary Admin Node. As soon as the primary Admin Node is upgraded, you can approve the other grid nodes.

4. As required, select **Continue** to access the **Upgrade other nodes** step.

Upgrade other nodes

You must upgrade all grid nodes, but you can perform multiple upgrade sessions and customize the upgrade sequence. For example, you might prefer to upgrade the nodes at site A in one session and then upgrade the nodes at site B in a later session. If you choose to perform the upgrade in more than one session, be aware that you can't start using the new features until all nodes have been upgraded.

If the order in which nodes are upgraded is important, approve nodes or groups of nodes one at a time and wait until the upgrade is complete on each node before approving the next node or group of nodes.



When the upgrade starts on a grid node, the services on that node are stopped. Later, the grid node is rebooted. To avoid service interruptions for client applications that are communicating with the node, don't approve the upgrade for a node unless you are sure that node is ready to be stopped and rebooted. As required, schedule a maintenance window or notify customers.

Steps

1. For the **Upgrade other nodes** step, review the Summary, which provides the start time for the upgrade as a whole and the status for each major upgrade task.
 - **Start upgrade service** is the first upgrade task. During this task, the software file is distributed to the grid nodes, and the upgrade service is started on each node.
 - When the **Start upgrade service** task is complete, the **Upgrade other grid nodes** task starts, and you are prompted to download a new copy of the Recovery Package.
2. When prompted, enter your provisioning passphrase and download a new copy of the Recovery Package.



You should download a new copy of the Recovery Package file after the primary Admin Node is upgraded. The Recovery Package file allows you to restore the system if a failure occurs.

3. Review the status tables for each type of node. There are tables for non-primary Admin Nodes, Gateway Nodes, and Storage Nodes.

A grid node can be in one of these stages when the tables first appear:

- Unpacking the upgrade
 - Downloading
 - Waiting to be approved
4. When you are ready to select grid nodes for upgrade (or if you need to unapprove selected nodes), use these instructions:

Task	Instruction
Search for specific nodes to approve, such as all nodes at a particular site	Enter the search string in the Search field

Task	Instruction
Select all nodes for upgrade	Select Approve all nodes
Select all nodes of the same type for upgrade (for example, all Storage Nodes)	Select the Approve all button for the node type If you approve more than one node of the same type, the nodes will be upgraded one at a time.
Select an individual node for upgrade	Select the Approve button for the node
Postpone the upgrade on all selected nodes	Select Unapprove all nodes
Postpone the upgrade on all selected nodes of the same type	Select the Unapprove all button for the node type
Postpone the upgrade on an individual node	Select the Unapprove button for the node

5. Wait for the approved nodes to proceed through these upgrade stages:

- Approved and waiting to be upgraded
- Stopping services



You can't remove a node when its Stage reaches **Stopping services**. The **Unapprove** button is disabled.

- Stopping container
- Cleaning up Docker images
- Upgrading base OS packages



When an appliance node reaches this stage, the StorageGRID Appliance Installer software on the appliance is updated. This automated process ensures that the StorageGRID Appliance Installer version remains in sync with the StorageGRID software version.

- Rebooting



Some appliance models might reboot multiple times to upgrade the firmware and BIOS.

- Performing steps after reboot
- Starting services
- Done

6. Repeat the [approval step](#) as many times as needed until all grid nodes have been upgraded.

Complete upgrade

When all grid nodes have completed the upgrade stages, the **Upgrade other grid nodes** task is shown as Completed. The remaining upgrade tasks are performed automatically in the background.

Steps

1. As soon as the **Enable features** task is complete (which occurs quickly), you can start using the [new features](#) in the upgraded StorageGRID version.
2. During the **Upgrade database** task, the upgrade process checks each node to verify that the Cassandra database does not need to be updated.



The upgrade from StorageGRID 11.8 to 11.9 does not require a Cassandra database upgrade; however, the Cassandra service will be stopped and restarted on each Storage Node. For future StorageGRID feature releases, the Cassandra database update step might take several days to complete.

3. When the **Upgrade database** task has completed, wait a few minutes for the **Final upgrade steps** to complete.
4. When the **Final upgrade steps** have completed, the upgrade is done. The first step, **Select files**, is redisplayed with a green success banner.
5. Verify that grid operations have returned to normal:
 - a. Check that the services are operating normally and that there are no unexpected alerts.
 - b. Confirm that client connections to the StorageGRID system are operating as expected.

Troubleshoot upgrade issues

If something goes wrong when you perform an upgrade, you might be able to resolve the issue yourself. If you can't resolve an issue, gather as much information as you can and then contact technical support.

Upgrade does not complete

The following sections describe how to recover from situations where the upgrade has partially failed.

Upgrade precheck errors

To detect and resolve issues, you can manually run the upgrade prechecks before starting the actual upgrade. Most precheck errors provide information about how to resolve the issue.

Provisioning failures

If the automatic provisioning process fails, contact technical support.

Grid node crashes or fails to start

If a grid node crashes during the upgrade process or fails to start successfully after the upgrade finishes, contact technical support to investigate and to correct any underlying issues.

Ingest or data retrieval is interrupted

If data ingest or retrieval is unexpectedly interrupted when you aren't upgrading a grid node, contact technical support.

Database upgrade errors

If the database upgrade fails with an error, retry the upgrade. If it fails again, contact technical support.

Related information

[Checking the system's condition before upgrading software](#)

User interface issues

You might experience issues with the Grid Manager or the Tenant Manager during or after the upgrade.

Grid Manager displays multiple error messages during upgrade

If you refresh your browser or navigate to another Grid Manager page while the primary Admin Node is being upgraded, you might see multiple "503: Service Unavailable" and "Problem connecting to the server" messages. You can safely ignore these messages—they will stop appearing soon as the node is upgraded.

If these messages appear for more than an hour after you started the upgrade, something might have occurred that prevented the primary Admin Node from being upgraded. If you are unable to resolve the issue on your own, contact technical support.

Web interface does not respond as expected

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

If you experience issues with the web interface:

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



Browser support typically changes for each StorageGRID release.

- Clear your web browser cache.

Clearing the cache removes outdated resources used by the previous version of StorageGRID software, and permits the user interface to operate correctly again. For instructions, see the documentation for your web browser.

"Docker image availability check" error messages

When attempting to start the upgrade process, you might receive an error message that states "The following issues were identified by the Docker image availability check validation suite." All issues must be resolved before you can complete the upgrade.

Contact technical support if you are unsure of the changes required to resolve the identified issues.

Message	Cause	Solution
Unable to determine upgrade version. Upgrade version info file {file_path} did not match the expected format.	The upgrade package is corrupt.	Re-upload the upgrade package, and try again. If the problem persists, contact technical support.

Message	Cause	Solution
Upgrade version info file {file_path} was not found. Unable to determine upgrade version.	The upgrade package is corrupt.	Re-upload the upgrade package, and try again. If the problem persists, contact technical support.
Unable to determine currently installed release version on {node_name}.	A critical file on the node is corrupt.	Contact technical support.
Connection error while attempting to list versions on {node_name}	The node is offline or the connection was interrupted.	Check to make sure that all nodes are online and reachable from the primary Admin Node, and try again.
The host for node {node_name} does not have StorageGRID {upgrade_version} image loaded. Images and services must be installed on the host before the upgrade can proceed.	The RPM or DEB packages for the upgrade have not been installed on the host where the node is running, or the images are still in the process of being imported. Note: This error only applies to nodes that are running as containers on Linux.	Check to make sure that the RPM or DEB packages have been installed on all Linux hosts where nodes are running. Make sure the version is correct for both the service and the images file. Wait a few minutes, and try again. See Linux: Install RPM or DEB package on all hosts.
Error while checking node {node_name}	An unexpected error occurred.	Wait a few minutes, and try again.
Uncaught error while running prechecks. {error_string}	An unexpected error occurred.	Wait a few minutes, and try again.

Apply StorageGRID hotfix

StorageGRID hotfix procedure

You might need to apply a hotfix to your StorageGRID system if issues with the software are detected and resolved between feature releases.

StorageGRID hotfixes contain software changes that are made available outside of a feature or patch release. The same changes are included in a future release. In addition, each hotfix release contains a roll-up of all previous hotfixes within the feature or patch release.

Considerations for applying a hotfix

You can't apply a StorageGRID hotfix when another maintenance procedure is running. For example, you can't apply a hotfix while a decommission, expansion, or recovery procedure is running.



If a node or site decommission procedure is paused, you can safely apply a hotfix. In addition, you might be able to apply a hotfix during the final stages of a StorageGRID upgrade procedure. See the instructions for upgrading StorageGRID software for details.

After you upload the hotfix in the Grid Manager, the hotfix is applied automatically to the primary Admin Node. Then, you can approve the application of the hotfix to the rest of the nodes in your StorageGRID system.

If a hotfix fails to be applied to one or more nodes, the reason for the failure appears in the Details column of the hotfix progress table. You must resolve whatever issues caused the failures and then retry the entire process. Nodes with a previously successful application of the hotfix will be skipped in subsequent applications. You can safely retry the hotfix process as many times as required until all nodes have been updated. The hotfix must be successfully installed on all grid nodes in order for the application to be complete.

While grid nodes are updated with the new hotfix version, the actual changes in a hotfix might only affect specific services on specific types of nodes. For example, a hotfix might only affect the LDR service on Storage Nodes.

How hotfixes are applied for recovery and expansion

After a hotfix has been applied to your grid, the primary Admin Node automatically installs the same hotfix version to any nodes restored by recovery operations or added in an expansion.

However, if you need to recover the primary Admin Node, you must manually install the correct StorageGRID release and then apply the hotfix. The final StorageGRID version of the primary Admin Node must match the version of the other nodes in the grid.

The following example illustrates how to apply a hotfix when recovering the primary Admin Node:

1. Assume the grid is running a StorageGRID 11.A.B version with the latest hotfix. The "grid version" is 11.A.B.y.
2. The primary Admin Node fails.
3. You redeploy the primary Admin Node using StorageGRID 11.A.B, and perform the recovery procedure.



As required to match the grid version, you can use a minor release when deploying the node; you don't need to deploy the major release first.

4. You then apply hotfix 11.A.B.y to the primary Admin Node.

For more information, see [Configure replacement primary Admin Node](#).

How your system is affected when you apply a hotfix

You must understand how your StorageGRID system will be affected when you apply a hotfix.

StorageGRID hotfixes are non-disruptive

The StorageGRID system can ingest and retrieve data from client applications throughout the hotfix process. If you approve all nodes of the same type to hotfix (for example, Storage Nodes), the nodes are brought down one at a time, so there is no time when all grid nodes or all grid nodes of a certain type are unavailable.

To allow for continued availability, ensure that your ILM policy contains rules that specify storing multiple copies

of each object. You must also ensure that all external S3 clients are configured to send requests to one of the following:

- A high availability (HA) group virtual IP address
- A high availability third-party load balancer
- Multiple Gateway Nodes for each client
- Multiple Storage Nodes for each client

Client applications might experience short-term disruptions

The StorageGRID system can ingest and retrieve data from client applications throughout the hotfix process; however, client connections to individual Gateway Nodes or Storage Nodes might be disrupted temporarily if the hotfix needs to restart services on those nodes. Connectivity will be restored after the hotfix process completes and services resume on the individual nodes.

You might need to schedule downtime to apply a hotfix if loss of connectivity for a short period is not acceptable. You can use selective approval to schedule when certain nodes are updated.



You can use multiple gateways and high availability (HA) groups to provide automatic failover during the hotfix process. See the instructions for [configuring high availability groups](#).

Alerts and SNMP notifications might be triggered

Alerts and SNMP notifications might be triggered when services are restarted and when the StorageGRID system is operating as a mixed-version environment (some grid nodes running an earlier version, while others have been upgraded to a later version). In general, these alerts and notifications will clear when the hotfix completes.

Configuration changes are restricted

When applying a hotfix to StorageGRID:

- Don't make any grid configuration changes (for example, specifying Grid Network subnets or approving pending grid nodes) until the hotfix has been applied to all nodes.
- Don't update the ILM configuration until the hotfix has been applied to all nodes.

Obtain required materials for hotfix

Before applying a hotfix, you must obtain all required materials.

Item	Notes
StorageGRID hotfix file	You must download the StorageGRID hotfix file.
• Network port • Supported web browser • SSH client (for example, PuTTY)	

Item	Notes
Recovery Package (.zip) file	Before applying a hotfix, download the most recent Recovery Package file in case any problems occur during the hotfix. Then, after the hotfix has been applied, download a new copy of the Recovery Package file and save it in a safe location. The updated Recovery Package file allows you to restore the system if a failure occurs.
Passwords.txt file	Optional and used only if you are applying a hotfix manually using the SSH client. The Passwords.txt file is part of the Recovery Package .zip file.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not listed in the Passwords.txt file.
Related documentation	readme.txt file for the hotfix. This file is included on the hotfix download page. Be sure to review the readme file carefully before applying the hotfix.

Download hotfix file

You must download the hotfix file before you can apply the hotfix.

Steps

1. Go to [NetApp Downloads: StorageGRID](#).
2. Select the down arrow under **Available Software** to see a list of hotfixes that are available to download.



Hotfix file versions have the form: 11.4.x.y.

3. Review the changes that are included in the update.



If you have just [recovered the primary Admin Node](#) and you need to apply a hotfix, select the same hotfix version that is installed on the other grid nodes.

- a. Select the hotfix version you want to download, and select **Go**.
- b. Sign in using the username and password for your NetApp account.
- c. Read and accept the End User License Agreement.

The download page for the version you selected appears.

- d. Download the hotfix readme.txt file to view a summary of the changes included in the hotfix.
4. Select the download button for the hotfix, and save the file.



Don't change the name of this file.



If you are using a macOS device, the hotfix file might be automatically saved as a `.txt` file. If it is, you must rename the file without the `.txt` extension.

5. Select a location for the download, and select **Save**.

Check system's condition before applying hotfix

You must verify the system is ready to accommodate the hotfix.

1. Sign in to the Grid Manager using a [supported web browser](#).
2. If possible, ensure that the system is running normally and that all grid nodes are connected to the grid.

Connected nodes have green check marks  on the Nodes page.

3. Check for and resolve any current alerts, if possible.
4. Ensure no other maintenance procedures are in progress, such as an upgrade, recovery, expansion, or decommission procedure.

You should wait for any active maintenance procedures to complete before applying a hotfix.

You can't apply a StorageGRID hotfix when another maintenance procedure is running. For example, you can't apply a hotfix while a decommission, expansion, or recovery procedure is running.



If a node or site [decommission procedure is paused](#), you can safely apply a hotfix. In addition, you might be able to apply a hotfix during the final stages of a StorageGRID upgrade procedure. See the instructions for [upgrading StorageGRID software](#).

Apply hotfix

The hotfix is first applied automatically to the primary Admin Node. Then, you must approve the application of the hotfix to other grid nodes until all nodes are running the same software version. You can customize the approval sequence by selecting to approve individual grid nodes, groups of grid nodes, or all grid nodes.

Before you begin

- You have reviewed the [considerations for applying a hotfix](#).
- You have the provisioning passphrase.
- You have Root access or the Maintenance permission.

About this task

- You can delay applying a hotfix to a node, but the hotfix process is not complete until you apply the hotfix to all nodes.
- You can't perform a StorageGRID software upgrade or a SANtricity OS update until you have completed the hotfix process.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).

2. Select **MAINTENANCE > System > Software update**.

The Software Update page appears.

Software update

You can upgrade StorageGRID software, apply a hotfix, or upgrade the SANtricity OS software on StorageGRID storage appliances. NetApp recommends you apply the latest hotfix before and after each software upgrade. Some hotfixes are required to prevent data loss.

StorageGRID upgrade

Upgrade to the next StorageGRID version and apply the latest hotfix for that version.

Upgrade →

StorageGRID hotfix

Apply a hotfix to your current StorageGRID software version.

Apply hotfix →

SANtricity OS update

Update the SANtricity OS software on your StorageGRID storage appliances.

Update →

3. Select **Apply hotfix**.

The StorageGRID Hotfix page appears.

StorageGRID Hotfix

Before starting the hotfix process, you must confirm that there are no active alerts and that all grid nodes are online and available. When the primary Admin Node is updated, services are stopped and restarted. Connectivity might be interrupted until the services are back online.

Hotfix file

Hotfix file ?

Passphrase

Provisioning Passphrase ?

4. Select the hotfix file you downloaded from the NetApp Support Site.

a. Select **Browse**.

b. Locate and select the file.

`hotfix-install-version`

c. Select **Open**.

The file is uploaded. When the upload is finished, the file name is shown in the Details field.



Don't change the file name because it is part of the verification process.

5. Enter the provisioning passphrase in the text box.

The **Start** button becomes enabled.

6. Select **Start**.

A warning appears stating that your browser's connection might be lost temporarily as services on the primary Admin Node are restarted.

7. Select **OK** to start applying the hotfix to the primary Admin Node.

When the hotfix starts:

- a. The hotfix validations are run.



If any errors are reported, resolve them, re-upload the hotfix file, and select **Start** again.

- b. The hotfix installation progress table appears.

This table shows all nodes in your grid and the current stage of the hotfix installation for each node. The nodes in the table are grouped by type (Admin Nodes, Gateway Nodes, and Storage Nodes).

- c. The progress bar reaches completion, and then the primary Admin Node is shown as "Complete."

Hotfix Installation Progress

						Approve All	Remove All
Admin Nodes - 1 out of 1 completed							
						Search	Q
Site	Name	Progress	Stage	Details	Action		
Vancouver	VTC-ADM1-101-191		Complete				

8. Optionally, sort the lists of nodes in each grouping in ascending or descending order by **Site**, **Name**, **Progress**, **Stage**, or **Details**. Or, enter a term in the **Search** box to search for specific nodes.
9. Approve the grid nodes that are ready to be updated. Approved nodes of the same type are upgraded one at a time.



Don't approve the hotfix for a node unless you are sure the node is ready to be updated. When the hotfix is applied to a grid node, some services on that node might be restarted. These operations might cause service interruptions for clients that are communicating with the node.

- Select one or more **Approve** buttons to add one or more individual nodes to the hotfix queue.
- Select the **Approve All** button within each grouping to add all nodes of the same type to the hotfix queue. If you have entered search criteria in the **Search** box, the **Approve All** button applies to all the nodes selected by the search criteria.



The **Approve All** button at the top of the page approves all nodes listed on the page, while the **Approve All** button at the top of a table grouping only approves all nodes in that group. If the order in which nodes are upgraded is important, approve nodes or groups of nodes one at a time and wait until the upgrade is complete on each node before approving the next node(s).

- Select the top-level **Approve All** button at the top of the page to add all nodes in the grid to the hotfix queue.



You must complete the StorageGRID hotfix before you can start a different software update. If you are unable to complete the hotfix, contact technical support.

- Select **Remove** or **Remove All** to remove a node or all nodes from the hotfix queue.

When the Stage progresses beyond "Queued," the **Remove** button is hidden and you can no longer remove the node from the hotfix process.

Storage Nodes - 1 out of 9 completed						Approve All	Remove All
						Search	
Site	Name	Progress	Stage	Details	Action		
Raleigh	RAL-S1-101-196		Queued		Remove		
Raleigh	RAL-S2-101-197		Complete				
Raleigh	RAL-S3-101-198		Queued		Remove		
Sunnyvale	SVL-S1-101-199		Queued		Remove		
Sunnyvale	SVL-S2-101-93		Waiting for you to approve		Approve		
Sunnyvale	SVL-S3-101-94		Waiting for you to approve		Approve		
Vancouver	VTC-S1-101-193		Waiting for you to approve		Approve		
Vancouver	VTC-S2-101-194		Waiting for you to approve		Approve		
Vancouver	VTC-S3-101-195		Waiting for you to approve		Approve		

10. Wait while the hotfix is applied to each approved grid node.

When the hotfix has been successfully installed on all nodes, the Hotfix Installation Progress table closes. A green banner shows the date and time the hotfix was completed.

11. If the hotfix could not be applied to any nodes, review the error for each node, resolve the issue, and repeat these steps.

The procedure is not complete until the hotfix is successfully applied to all nodes. You can safely retry the hotfix process as many times as required until it is complete.

Configure and manage a StorageGRID system

Administer StorageGRID

Administer StorageGRID

Use these instructions to configure and administer a StorageGRID system.

About these instructions

The primary tasks for configuring and administering StorageGRID allow you to:

- Use the Grid Manager to set up groups and users
- Create tenant accounts to allow S3 client applications to store and retrieve objects
- Configure and manage StorageGRID networks
- Configure AutoSupport
- Manage node settings

Before you begin

- You have a general understanding of the StorageGRID system.
- You have fairly detailed knowledge of Linux command shells, networking, and server hardware setup and configuration.

Get started with Grid Manager

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	119
Microsoft Edge	119
Mozilla Firefox	119

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Sign in to the Grid Manager

You access the Grid Manager sign-in page by entering the fully qualified domain name (FQDN) or IP address of an Admin Node into the address bar of a supported web browser.

Each StorageGRID system includes one primary Admin Node and any number of non-primary Admin Nodes. You can sign in to the Grid Manager on any Admin Node to manage the StorageGRID system. However, some maintenance procedures can only be performed from the primary Admin Node.

Connect to HA group

If Admin Nodes are included in a high availability (HA) group, you connect using the virtual IP address of the HA group or a fully qualified domain name that maps to the virtual IP address. The primary Admin Node should be selected as the group's primary interface, so that when you access the Grid Manager, you access it on the primary Admin Node unless the primary Admin Node is not available. See [Manage high availability groups](#).

Use SSO

The sign-in steps are slightly different if [single sign-on \(SSO\) has been configured](#).

Sign in to Grid Manager on first Admin Node

Before you begin

- You have your login credentials.
- You are using a [supported web browser](#).
- Cookies are enabled in your web browser.
- You belong to a user group that has at least one permission.
- You have the URL for the Grid Manager:

```
https://FQDN_or_Admin_Node_IP/
```

You can use the fully qualified domain name, the IP address of an Admin Node, or the virtual IP address of an HA group of Admin Nodes.

To access the Grid Manager on a port other than the default port for HTTPS (443), include the port number in the URL:

```
https://FQDN_or_Admin_Node_IP:port/
```



SSO is not available on the restricted Grid Manager port. You must use port 443.

Steps

1. Launch a supported web browser.
2. In the browser's address bar, enter the URL for the Grid Manager.
3. If you are prompted with a security alert, install the certificate using the browser's installation wizard. See [Manage security certificates](#).
4. Sign in to the Grid Manager.

The sign-in screen that appears depends on whether single sign-on (SSO) has been configured for StorageGRID.

Not using SSO

- a. Enter your username and password for the Grid Manager.
- b. Select **Sign In**.

The image shows a login form for NetApp StorageGRID Grid Manager. At the top, there is a logo consisting of a square icon with a stylized 'N' followed by the text 'NetApp StorageGRID®'. Below the logo, the title 'Grid Manager' is displayed in a large, bold, sans-serif font. Underneath the title, the label 'Username' is positioned above a text input field. The input field contains a single vertical line cursor. Below the username field, the label 'Password' is positioned above another text input field. At the bottom of the form, there is a blue rectangular button with the text 'Sign in' in white. Below the button, there are three links: 'Tenant sign in', 'NetApp support', and 'NetApp.com', all in a smaller, blue, sans-serif font.

Using SSO

- If StorageGRID is using SSO and this is the first time you have accessed the URL on this browser:
 - a. Select **Sign in**. You can leave the 0 in the Account field.



Sign in

Account

Sign in

[NetApp support](#) | [NetApp.com](#)

- b. Enter your standard SSO credentials on your organization's SSO sign-in page. For example:

Sign in with your organizational account

Sign in

- If StorageGRID is using SSO and you have previously accessed the Grid Manager or a tenant account:
 - a. Enter **0** (the account ID for the Grid Manager) or select **Grid Manager** if it appears in the list of recent accounts.

The image shows a web interface for NetApp StorageGRID. At the top, the logo "NetApp StorageGRID®" is displayed. Below it is the heading "Sign in". Under the heading, there is a section labeled "Recent" with a dropdown menu showing "Grid Manager". Below that is a section labeled "Account" with a text input field containing the character "0". A blue "Sign in" button is positioned below the input field. At the bottom of the form, there are links for "NetApp support" and "NetApp.com".

NetApp StorageGRID®

Sign in

Recent

Grid Manager ▼

Account


0

Sign in

[NetApp support](#) | [NetApp.com](#)

- b. Select **Sign in**.
- c. Sign in with your standard SSO credentials on your organization's SSO sign-in page.

When you are signed in, the home page of the Grid Manager appears, which includes the dashboard. To learn what information is provided, see [View and manage the dashboard](#).



Sign into another Admin Node

Follow these steps to sign in to another Admin Node.

Not using SSO

Steps

1. In the browser's address bar, enter the fully qualified domain name or IP address of the other Admin Node. Include the port number as required.
2. Enter your username and password for the Grid Manager.
3. Select **Sign In**.

Using SSO

If StorageGRID is using SSO and you have signed in to one Admin Node, you can access other Admin Nodes without having to sign in again.

Steps


1. Enter the fully qualified domain name or IP address of the other Admin Node in the browser's address bar.
2. If your SSO session has expired, enter your credentials again.

Sign out of the Grid Manager

When you are done working with the Grid Manager, you must sign out to ensure that unauthorized users can’t access the StorageGRID system. Closing your browser might not sign you out of the system, based on browser cookie settings.

Steps

- 1. Select your user name in the top-right corner.



- 2. Select **Sign out**.

Option	Description
SSO not in use	You are signed out of the Admin Node. The Grid Manager sign in page is displayed. Note: If you signed into more than one Admin Node, you must sign out of each node.
SSO enabled	You are signed out of all Admin Nodes you were accessing. The StorageGRID sign in page is displayed. Grid Manager is listed as the default in the Recent Accounts drop-down, and the Account ID field shows 0. Note: If SSO is enabled and you are also signed in to the Tenant Manager, you must also sign out of the tenant account to sign out of SSO.

Change your password

If you are a local user of the Grid Manager, you can change your own password.

Before you begin

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

About this task

If you sign in to StorageGRID as a federated user or if single sign-on (SSO) is enabled, you can’t change your password in Grid Manager. Instead, you must change your password in the external identity source, for

example, Active Directory or OpenLDAP.

Steps

1. From the Grid Manager header, select ***your name*** > **Change password**.
2. Enter your current password.
3. Type a new password.

Your password must contain at least 8 and no more than 32 characters. Passwords are case-sensitive.

4. Re-enter the new password.
5. Select **Save**.

View StorageGRID license information

You can view the license information for your StorageGRID system, such as the maximum storage capacity of your grid, whenever necessary.

Before you begin

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

About this task

If there is an issue with the software license for this StorageGRID system, the Health status card on the dashboard includes a License status icon and a **License** link. The number indicates the number of license-related issues.



Steps

1. Access the License page by doing one of the following:
 - Select **MAINTENANCE** > **System** > **License**.
 - From the Health status card on the dashboard, select the License status icon or the **License** link.

This link appears only if there is an issue with the license.

2. View the read-only details for the current license:
 - StorageGRID system ID, which is the unique identification number for this StorageGRID installation

- License serial number
- License type, either **Perpetual** or **Subscription**
- Licensed storage capacity of the grid
- Supported storage capacity
- License end date. **N/A** appears for a perpetual license.
- Support end date

This date is read from the current license file and might be out of date if you extended or renewed the support service contract after obtaining the license file. To update this value, see [Update StorageGRID license information](#). You can also view the actual contract end date using Active IQ.

- Contents of the license text file

Update StorageGRID license information

You must update the license information for your StorageGRID system any time the terms of your license change. For example, you must update the license information if you purchase additional storage capacity for your grid.

Before you begin

- You have a new license file to apply to your StorageGRID system.
- You have [specific access permissions](#).
- You have the provisioning passphrase.

Steps

1. Select **MAINTENANCE > System > License**.
2. In the Update license section, select **Browse**.
3. Locate and select the new license file (.txt).

The new license file is validated and displayed.

4. Enter the provisioning passphrase.
5. Select **Save**.

Use the API

Use the Grid Management API

You can perform system management tasks using the Grid Management REST API instead of the Grid Manager user interface. For example, you might want to use the API to automate operations or to create multiple entities, such as users, more quickly.

Top-level resources

The Grid Management API provides the following top-level resources:

- /grid: Access is restricted to Grid Manager users and is based on the configured group permissions.

- `/org`: Access is restricted to users who belong to a local or federated LDAP group for a tenant account. For details, see [Use a tenant account](#).
- `/private`: Access is restricted to Grid Manager users and is based on the configured group permissions. The private APIs are subject to change without notice. StorageGRID private endpoints also ignore the API version of the request.

Issue API requests

The Grid Management API uses the Swagger open source API platform. Swagger provides an intuitive user interface that allows developers and non-developers to perform real-time operations in StorageGRID with the API.

The Swagger user interface provides complete details and documentation for each API operation.

Before you begin


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



Any API operations you perform using the API Documentation webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Steps

1. From the Grid Manager header, select the help icon and select **API documentation**.



2. To perform an operation with the private API, select **Go to private API documentation** on the StorageGRID Management API page.

The private APIs are subject to change without notice. StorageGRID private endpoints also ignore the API version of the request.

3. Select the desired operation.

When you expand an API operation, you can see the available HTTP actions, such as GET, PUT, UPDATE, and DELETE.

4. Select an HTTP action to see the request details, including the endpoint URL, a list of any required or optional parameters, an example of the request body (when required), and the possible responses.

groups Operations on groups

GET `/grid/groups` Lists Grid Administrator Groups

Parameters Try it out

Name	Description
type string (query)	filter by group type Available values : local, federated
limit integer (query)	maximum number of results Default value : 25
marker string (query)	marker-style pagination offset (value is Group's URN)
includeMarker boolean (query)	if set, the marker element is also returned
order string (query)	pagination order (desc requires marker) Available values : asc, desc

Responses Response content type: application/json

Code	Description
200	successfully retrieved

Example Value | Model

```
{
  "responseTime": "2021-03-29T14:22:19.673Z",
  "status": "success",
  "apiVersion": "3.3",
  "deprecated": false,
  "data": [
    {
      "displayName": "Developers",

```

5. Determine if the request requires additional parameters, such as a group or user ID. Then, obtain these values. You might need to issue a different API request first to get the information you need.
6. Determine if you need to modify the example request body. If so, you can select **Model** to learn the requirements for each field.
7. Select **Try it out**.
8. Provide any required parameters, or modify the request body as required.
9. Select **Execute**.

10. Review the response code to determine if the request was successful.

Grid Management API operations

The Grid Management API organizes the available operations into the following sections.



This list only includes operations available in the public API.

- **accounts:** Operations to manage storage tenant accounts, including creating new accounts and retrieving storage usage for a given account.
- **alert-history:** Operations on resolved alerts.
- **alert-receivers:** Operations on alert notification receivers (email).
- **alert-rules:** Operations on alert rules.
- **alert-silences:** Operations on alert silences.
- **alerts:** Operations on alerts.
- **audit:** Operations to list and update the audit configuration.
- **auth:** Operations to perform user session authentication.

The Grid Management API supports the Bearer Token Authentication Scheme. To sign in, you provide a username and password in the JSON body of the authentication request (that is, `POST /api/v3/authorize`). If the user is successfully authenticated, a security token is returned. This token must be provided in the header of subsequent API requests ("Authorization: Bearer *token*"). The token expires after 16 hours.



If single sign-on is enabled for the StorageGRID system, you must perform different steps to authenticate. See "Authenticating in to the API if single sign-on is enabled."

See "Protecting against Cross-Site Request Forgery" for information about improving authentication security.

- **client-certificates:** Operations to configure client certificates so that StorageGRID can be accessed securely using external monitoring tools.
- **config:** Operations related to the product release and versions of the Grid Management API. You can list the product release version and the major versions of the Grid Management API supported by that release, and you can disable deprecated versions of the API.
- **deactivated-features:** Operations to view features that might have been deactivated.
- **dns-servers:** Operations to list and change configured external DNS servers.
- **drive-details:** Operations on drives for specific storage appliance models.
- **endpoint-domain-names:** Operations to list and change S3 endpoint domain names.
- **erasure-coding:** Operations on erasure-coding profiles.
- **expansion:** Operations on expansion (procedure-level).
- **expansion-nodes:** Operations on expansion (node-level).
- **expansion-sites:** Operations on expansion (site-level).
- **grid-networks:** Operations to list and change the Grid Network List.
- **grid-passwords:** Operations for grid password management.

- **groups:** Operations to manage local Grid Administrator Groups and to retrieve federated Grid Administrator Groups from an external LDAP server.
- **identity-source:** Operations to configure an external identity source and to manually synchronize federated group and user information.
- **ilm:** Operations on information lifecycle management (ILM).
- **in-progress-procedures:** Retrieves the maintenance procedures that are currently in progress.
- **license:** Operations to retrieve and update the StorageGRID license.
- **logs:** Operations for collecting and downloading log files.
- **metrics:** Operations on StorageGRID metrics including instant metric queries at a single point in time and range metric queries over a range of time. The Grid Management API uses the Prometheus systems monitoring tool as the backend data source. For information about constructing Prometheus queries, see the Prometheus web site.



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

- **node-details:** Operations on node details.
- **node-health:** Operations on node health status.
- **node-storage-state:** Operations on node storage status.
- **ntp-servers:** Operations to list or update external Network Time Protocol (NTP) servers.
- **objects:** Operations on objects and object metadata.
- **recovery:** Operations for the recovery procedure.
- **recovery-package:** Operations to download the Recovery Package.
- **regions:** Operations to view and create regions.
- **s3-object-lock:** Operations on global S3 Object Lock settings.
- **server-certificate:** Operations to view and update Grid Manager server certificates.
- **snmp:** Operations on the current SNMP configuration.
- **storage-watermarks:** Storage node watermarks.
- **traffic-classes:** Operations for traffic classification policies.
- **untrusted-client-network:** Operations on the untrusted Client Network configuration.
- **users:** Operations to view and manage Grid Manager users.

Grid Management API versioning

The Grid Management API uses versioning to support non-disruptive upgrades.

For example, this Request URL specifies version 4 of the API.

```
https://hostname_or_ip_address/api/v4/authorize
```

The major version of the API is bumped when changes are made that are *not compatible* with older versions. The minor version of the API is bumped when changes are made that *are compatible* with older versions. Compatible changes include the addition of new endpoints or new properties.

The following example illustrates how the API version is bumped based on the type of changes made.

Type of change to API	Old version	New version
Compatible with older versions	2.1	2.2
Not compatible with older versions	2.1	3.0
	3.0	4.0

When you install StorageGRID software for the first time, only the most recent version of the API is enabled. However, when you upgrade to a new feature release of StorageGRID, you continue to have access to the older API version for at least one StorageGRID feature release.



You can configure the supported versions. See the **config** section of the Swagger API documentation for the [Grid Management API](#) for more information. You should deactivate support for the older version after updating all API clients to use the newer version.

Outdated requests are marked as deprecated in the following ways:

- The response header is "Deprecated: true"
- The JSON response body includes "deprecated": true
- A deprecated warning is added to nms.log. For example:

```
Received call to deprecated v2 API at POST "/api/v2/authorize"
```

Determine which API versions are supported in the current release

Use the `GET /versions` API request to return a list of the supported API major versions. This request is located in the **config** section of the Swagger API documentation.

```
GET https://{IP-Address}/api/versions
{
  "responseTime": "2023-06-27T22:13:50.750Z",
  "status": "success",
  "apiVersion": "4.0",
  "data": [
    2,
    3,
    4
  ]
}
```

Specify an API version for a request

You can specify the API version using a path parameter (`/api/v4`) or a header (`Api-Version: 4`). If you provide both values, the header value overrides the path value.

```
curl https://[IP-Address]/api/v4/grid/accounts

curl -H "Api-Version: 4" https://[IP-Address]/api/grid/accounts
```

Protect against Cross-Site Request Forgery (CSRF)

You can help protect against Cross-Site Request Forgery (CSRF) attacks against StorageGRID by using CSRF tokens to enhance authentication that uses cookies. The Grid Manager and Tenant Manager automatically enable this security feature; other API clients can choose whether to enable it when they sign in.

An attacker that can trigger a request to a different site (such as with an HTTP form POST) can cause certain requests to be made using the signed-in user's cookies.

StorageGRID helps protect against CSRF attacks by using CSRF tokens. When enabled, the contents of a specific cookie must match the contents of either a specific header or a specific POST body parameter.

To enable the feature, set the `csrfToken` parameter to `true` during authentication. The default is `false`.

```
curl -X POST --header "Content-Type: application/json" --header "Accept: application/json" -d "{
  \"username\": \"MyUserName\",
  \"password\": \"MyPassword\",
  \"cookie\": true,
  \"csrfToken\": true
}" "https://example.com/api/v3/authorize"
```

When `true`, a `GridCsrfToken` cookie is set with a random value for sign-ins to the Grid Manager, and the `AccountCsrfToken` cookie is set with a random value for sign-ins to the Tenant Manager.

If the cookie is present, all requests that can modify the state of the system (POST, PUT, PATCH, DELETE) must include one of the following:

- The `X-Csrf-Token` header, with the value of the header set to the value of the CSRF token cookie.
- For endpoints that accept a form-encoded body: A `csrfToken` form-encoded request body parameter.

See the online API documentation for additional examples and details.



Requests that have a CSRF token cookie set will also enforce the "Content-Type: application/json" header for any request that expects a JSON request body as an additional protection against CSRF attacks.

Use the API if single sign-on is enabled

Use the API if single sign-on is enabled (Active Directory)

If you have [configured and enabled single sign-on \(SSO\)](#) and you use Active Directory as

the SSO provider, you must issue a series of API requests to obtain an authentication token that is valid for the Grid Management API or the Tenant Management API.

Sign in to the API if single sign-on is enabled

These instructions apply if you are using Active Directory as the SSO identity provider.

Before you begin

- You know the SSO username and password for a federated user who belongs to a StorageGRID user group.
- If you want to access the Tenant Management API, you know the tenant account ID.

About this task

To obtain an authentication token, you can use one of the following examples:

- The `storagegrid-ssoauth.py` Python script, which is located in the StorageGRID installation files directory (`./rpms` for Red Hat Enterprise Linux, `./debs` for Ubuntu or Debian, and `./vsphere` for VMware).
- An example workflow of curl requests.

The curl workflow might time out if you perform it too slowly. You might see the error: A valid SubjectConfirmation was not found on this Response.



The example curl workflow does not protect the password from being seen by other users.

If you have a URL-encoding issue, you might see the error: Unsupported SAML version.

Steps

1. Select one of the following methods to obtain an authentication token:
 - Use the `storagegrid-ssoauth.py` Python script. Go to step 2.
 - Use curl requests. Go to step 3.
2. If you want to use the `storagegrid-ssoauth.py` script, pass the script to the Python interpreter and run the script.

When prompted, enter values for the following arguments:

- The SSO method. Enter ADFS or adfs.
- The SSO username
- The domain where StorageGRID is installed
- The address for StorageGRID
- The tenant account ID, if you want to access the Tenant Management API.

```
python3 storagegrid-ssoauth.py
sso_method: adfs
saml_user: my-sso-username
saml_domain: my-domain
sg_address: storagegrid.example.com
tenant_account_id: 12345
Enter the user's SAML password:
*****

*****

StorageGRID Auth Token: 56eb07bf-21f6-40b7-afob-5c6cacfb25e7
```

The StorageGRID authorization token is provided in the output. You can now use the token for other requests, similar to how you would use the API if SSO was not being used.

3. If you want to use curl requests, use the following procedure.

a. Declare the variables needed to sign in.

```
export SAMLUSER='my-sso-username'
export SAMLPASSWORD='my-password'
export SAMLDOMAIN='my-domain'
export TENANTACCOUNTID='12345'
export STORAGEGRID_ADDRESS='storagegrid.example.com'
export AD_FS_ADDRESS='adfs.example.com'
```



To access the Grid Management API, use 0 as TENANTACCOUNTID.

b. To receive a signed authentication URL, issue a POST request to `/api/v3/authorize-saml`, and remove the additional JSON encoding from the response.

This example shows a POST request for a signed authentication URL for TENANTACCOUNTID. The results will be passed to `python -m json.tool` to remove the JSON encoding.

```
curl -X POST "https://$STORAGEGRID_ADDRESS/api/v3/authorize-saml" \
-H "accept: application/json" -H "Content-Type: application/json" \
--data "{\"accountId\": \"$TENANTACCOUNTID\"}" | python -m
json.tool
```

The response for this example includes a signed URL that is URL-encoded, but it does not include the additional JSON-encoding layer.

```
{
  "apiVersion": "3.0",
  "data":
  "https://adfs.example.com/adfs/ls/?SAMLRequest=fZHLbsIwEEV%2FJTuv7...
sSl%2BfQ33cvfwA%3D&RelayState=12345",
  "responseTime": "2018-11-06T16:30:23.355Z",
  "status": "success"
}
```

- c. Save the SAMLRequest from the response for use in subsequent commands.

```
export SAMLREQUEST='fZHLbsIwEEV%2FJTuv7...sSl%2BfQ33cvfwA%3D'
```

- d. Get a full URL that includes the client request ID from AD FS.

One option is to request the login form using the URL from the previous response.

```
curl "https://$AD_FS_ADDRESS/adfs/ls/?SAMLRequest=
$SAMLREQUEST&RelayState=$TENANTACCOUNTID" | grep 'form method="post"
id="loginForm"'
```

The response includes the client request ID:

```
<form method="post" id="loginForm" autocomplete="off"
novalidate="novalidate" onKeyPress="if (event && event.keyCode == 13)
Login.submitLoginRequest();" action="/adfs/ls/?
SAMLRequest=fZHRToMwFIZfhh...UJikvo77sXPw%3D%3D&RelayState=12345&clie
nt-request-id=00000000-0000-0000-ee02-0080000000de" >
```

- e. Save the client request ID from the response.

```
export SAMLREQUESTID='00000000-0000-0000-ee02-0080000000de'
```

- f. Send your credentials to the form action from the previous response.

```
curl -X POST "https://$AD_FS_ADDRESS
/adfs/ls/?SAMLRequest=$SAMLREQUEST&RelayState=$TENANTACCOUNTID&client
-request-id=$SAMLREQUESTID" \
--data "UserName=$SAMLUSER@$SAMLDOMAIN&Password=
$SAMLPASSWORD&AuthMethod=FormsAuthentication" --include
```


AD FS returns a 302 redirect, with additional information in the headers.



If multi-factor authentication (MFA) is enabled for your SSO system, the form post will also contain the second password or other credentials.

```
HTTP/1.1 302 Found
Content-Length: 0
Content-Type: text/html; charset=utf-8
Location:
https://adfs.example.com/adfs/ls/?SAMLRequest=fZHRT0MwFIZfhb...UJikvo
77sXPw%3D%3D&RelayState=12345&client-request-id=00000000-0000-0000-
ee02-0080000000de
Set-Cookie: MSISAuth=AAEAADAvsHpXk6ApV...pmP0aEiNtJvWY=; path=/adfs;
HttpOnly; Secure
Date: Tue, 06 Nov 2018 16:55:05 GMT
```

g. Save the MSISAuth cookie from the response.

```
export MSISAuth='AAEAADAvsHpXk6ApV...pmP0aEiNtJvWY='
```

h. Send a GET request to the specified location with the cookies from the authentication POST.

```
curl "https://$AD_FS_ADDRESS/adfs/ls/?SAMLRequest=
$SAMLREQUEST&RelayState=$TENANTACCOUNTID&client-request-
id=$SAMLREQUESTID" \
--cookie "MSISAuth=$MSISAuth" --include
```

The response headers will contain AD FS session information for later logout usage, and the response body contains the SAMLResponse in a hidden form field.

```

HTTP/1.1 200 OK
Cache-Control: no-cache,no-store
Pragma: no-cache
Content-Length: 5665
Content-Type: text/html; charset=utf-8
Expires: -1
Server: Microsoft-HTTPAPI/2.0
P3P: ADFS doesn't have P3P policy, please contact your site's admin
for more details
Set-Cookie:
SamlSession=a3dpbnRlcnMtUHJpbWFyeS1BZG1pbi0xNzgmRmFsc2Umcng4NnJDZmFKV
XFxVWx3bk1lMnFuUSUzZCUzZCYmJiYmXzE3MjAyZTA5LTNmMDgtNDRkZC04Yzg5LTQ3ND
UxYzA3ZjkzYw==; path=/adfs; HttpOnly; Secure
Set-Cookie: MSISAuthenticated=MTEvNy8yMDE4IDQ6MzI6NTkgUE0=;
path=/adfs; HttpOnly; Secure
Set-Cookie: MSISLoopDetectionCookie=MjAxOC0xMS0wNzoxNjozMj01OVpcMQ==;
path=/adfs; HttpOnly; Secure
Date: Wed, 07 Nov 2018 16:32:59 GMT

<form method="POST" name="hiddenform"
action="https://storagegrid.example.com:443/api/saml-response">
  <input type="hidden" name="SAMLResponse"
value="PHNhbWxwOlJlc3Bvb3N...1scDpSZXNwb25zZT4=" /><input
type="hidden" name="RelayState" value="12345" />

```

- i. Save the SAMLResponse from the hidden field:

```
export SAMLResponse='PHNhbWxwOlJlc3Bvb3N...1scDpSZXNwb25zZT4='
```

- j. Using the saved SAMLResponse, make a StorageGRID/api/saml-response request to generate a StorageGRID authentication token.

For RelayState, use the tenant account ID or use 0 if you want to sign in to the Grid Management API.

```

curl -X POST "https://$STORAGEGRID_ADDRESS:443/api/saml-response" \
-H "accept: application/json" \
--data-urlencode "SAMLResponse=$SAMLResponse" \
--data-urlencode "RelayState=$TENANTACCOUNTID" \
| python -m json.tool

```

The response includes the authentication token.

```
{
  "apiVersion": "3.0",
  "data": "56eb07bf-21f6-40b7-af0b-5c6cacfb25e7",
  "responseTime": "2018-11-07T21:32:53.486Z",
  "status": "success"
}
```

k. Save the authentication token in the response as MYTOKEN.

```
export MYTOKEN="56eb07bf-21f6-40b7-af0b-5c6cacfb25e7"
```

You can now use MYTOKEN for other requests, similar to how you would use the API if SSO was not being used.

Sign out of the API if single sign-on is enabled

If single sign-on (SSO) has been enabled, you must issue a series of API requests to sign out of the Grid Management API or the Tenant Management API.

These instructions apply if you are using Active Directory as the SSO identity provider

About this task

If required, you can sign out of the StorageGRID API by logging out from your organization's single logout page. Or, you can trigger single logout (SLO) from StorageGRID, which requires a valid StorageGRID bearer token.

Steps

1. To generate a signed logout request, pass `cookie "sso=true" to the SLO API:

```
curl -k -X DELETE "https://$STORAGEGRID_ADDRESS/api/v3/authorize" \
-H "accept: application/json" \
-H "Authorization: Bearer $MYTOKEN" \
--cookie "sso=true" \
| python -m json.tool
```

A logout URL is returned:

```
{
  "apiVersion": "3.0",
  "data":
  "https://adfs.example.com/adfs/ls/?SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3D",
  "responseTime": "2018-11-20T22:20:30.839Z",
  "status": "success"
}
```

2. Save the logout URL.

```
export LOGOUT_REQUEST
='https://adfs.example.com/adfs/ls/?SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3D'
```

3. Send a request to the logout URL to trigger SLO and to redirect back to StorageGRID.

```
curl --include "$LOGOUT_REQUEST"
```

The 302 response is returned. The redirect location is not applicable to API-only logout.

```
HTTP/1.1 302 Found
Location: https://$STORAGEGRID_ADDRESS:443/api/saml-logout?SAMLResponse=fVLLasMwEPwVo7ss%...%23rsa-sha256
Set-Cookie: MSISSignoutProtocol=U2FtbA==; expires=Tue, 20 Nov 2018 22:35:03 GMT; path=/adfs; HttpOnly; Secure
```

4. Delete the StorageGRID bearer token.

Deleting the StorageGRID bearer token works the same way as without SSO. If `cookie "sso=true" is not provided, the user is logged out of StorageGRID without affecting the SSO state.

```
curl -X DELETE "https://$STORAGEGRID_ADDRESS/api/v3/authorize" \
-H "accept: application/json" \
-H "Authorization: Bearer $MYTOKEN" \
--include
```

A 204 No Content response indicates the user is now signed out.

```
HTTP/1.1 204 No Content
```

Use the API if single sign-on is enabled (Azure)

If you have [configured and enabled single sign-on \(SSO\)](#) and you use Azure as the SSO provider, you can use two example scripts to obtain an authentication token that is valid for the Grid Management API or the Tenant Management API.

Sign in to the API if Azure single sign-on is enabled

These instructions apply if you are using Azure as the SSO identity provider

Before you begin

- You know the SSO email address and password for a federated user who belongs to a StorageGRID user group.
- If you want to access the Tenant Management API, you know the tenant account ID.

About this task

To obtain an authentication token, you can use the following example scripts:

- The `storagegrid-ssoauth-azure.py` Python script
- The `storagegrid-ssoauth-azure.js` Node.js script

Both scripts are located in the StorageGRID installation files directory (`./rpms` for Red Hat Enterprise Linux, `./debs` for Ubuntu or Debian, and `./vsphere` for VMware).

To write your own API integration with Azure, see the `storagegrid-ssoauth-azure.py` script. The Python script makes two requests to StorageGRID directly (first to get the SAMLRequest, and later to get the authorization token), and also calls the Node.js script to interact with Azure to perform the SSO operations.

SSO operations can be executed using a series of API requests, but doing so is not straightforward. The Puppeteer Node.js module is used to scrape the Azure SSO interface.

If you have a URL-encoding issue, you might see the error: `Unsupported SAML version`.

Steps

1. Install the required dependencies, as follows:
 - a. Install Node.js (see <https://nodejs.org/en/download/>).
 - b. Install the required Node.js modules (puppeteer and jsdom):

```
npm install -g <module>
```

2. Pass the Python script to the Python interpreter to run the script.

The Python script will then call the corresponding Node.js script to perform the Azure SSO interactions.

3. When prompted, enter values for the following arguments (or pass them in using parameters):
 - The SSO email address used to sign in to Azure
 - The address for StorageGRID
 - The tenant account ID, if you want to access the Tenant Management API
4. When prompted, enter the password and be prepared to provide an MFA authorization to Azure if

requested.

```
c:\Users\user\Documents\azure_sso>py storagegrid-azure-ssoauth.py --sso-email-address user@my-domain.com
--sg-address storagegrid.examp.e.com --tenant-account-id 0
Enter the user's SSO password:
*****

Watch for and approve a 2FA authorization request
*****

StorageGRID Auth Token: {'responseTime': '2021-10-04T21:30:48.807Z', 'status': 'success', 'apiVersion':
'3.4', 'data': '4807d93e-a3df-48f2-9680-906cd255979e'}
```



The script assumes MFA is done using Microsoft Authenticator. You might need to modify the script to support other forms of MFA (such as entering a code received in a text message).

The StorageGRID authorization token is provided in the output. You can now use the token for other requests, similar to how you would use the API if SSO was not being used.

Use the API if single sign-on is enabled (PingFederate)

If you have [configured and enabled single sign-on \(SSO\)](#) and you use PingFederate as the SSO provider, you must issue a series of API requests to obtain an authentication token that is valid for the Grid Management API or the Tenant Management API.

Sign in to the API if single sign-on is enabled

These instructions apply if you are using PingFederate as the SSO identity provider

Before you begin

- You know the SSO username and password for a federated user who belongs to a StorageGRID user group.
- If you want to access the Tenant Management API, you know the tenant account ID.

About this task

To obtain an authentication token, you can use one of the following examples:

- The `storagegrid-ssoauth.py` Python script, which is located in the StorageGRID installation files directory (`./rpms` for Red Hat Enterprise Linux, `./debs` for Ubuntu or Debian, and `./vsphere` for VMware).
- An example workflow of curl requests.

The curl workflow might time out if you perform it too slowly. You might see the error: A valid SubjectConfirmation was not found on this Response.



The example curl workflow does not protect the password from being seen by other users.

If you have a URL-encoding issue, you might see the error: Unsupported SAML version.

Steps

1. Select one of the following methods to obtain an authentication token:
 - Use the `storagegrid-ssoauth.py` Python script. Go to step 2.

- Use curl requests. Go to step 3.
2. If you want to use the `storagegrid-ssoauth.py` script, pass the script to the Python interpreter and run the script.

When prompted, enter values for the following arguments:

- The SSO method. You can enter any variation of "pingfederate" (PINGFEDERATE, pingfederate, and so on).
- The SSO username
- The domain where StorageGRID is installed. This field is not used for PingFederate. You can leave it blank or enter any value.
- The address for StorageGRID
- The tenant account ID, if you want to access the Tenant Management API.

```
python3 storagegrid-ssoauth.py
sso_method: pingfederate
saml_user: my-sso-username
saml_domain:
sg_address: storagegrid.example.com
tenant_account_id: 12345
Enter the user's SAML password:
*****

*****
StorageGRID Auth Token: 56eb07bf-21f6-40b7-afob-5c6cacfb25e7
```

The StorageGRID authorization token is provided in the output. You can now use the token for other requests, similar to how you would use the API if SSO was not being used.

3. If you want to use curl requests, use the following procedure.
- a. Declare the variables needed to sign in.

```
export SAMLUSER='my-sso-username'
export SAMLPASSWORD='my-password'
export TENANTACCOUNTID='12345'
export STORAGEGRID_ADDRESS='storagegrid.example.com'
```



To access the Grid Management API, use 0 as TENANTACCOUNTID.

- b. To receive a signed authentication URL, issue a POST request to `/api/v3/authorize-saml`, and remove the additional JSON encoding from the response.

This example shows a POST request for a signed authentication URL for TENANTACCOUNTID. The results will be passed to `python -m json.tool` to remove the JSON encoding.

```
curl -X POST "https://$STORAGEGRID_ADDRESS/api/v3/authorize-saml" \
  -H "accept: application/json" -H "Content-Type: application/json" \
  --data "{\"accountId\": \"$TENANTACCOUNTID\"}" | python -m
json.tool
```

The response for this example includes a signed URL that is URL-encoded, but it does not include the additional JSON-encoding layer.

```
{
  "apiVersion": "3.0",
  "data": "https://my-pf-baseurl/idp/SSO.saml2?...",
  "responseTime": "2018-11-06T16:30:23.355Z",
  "status": "success"
}
```

- c. Save the SAMLRequest from the response for use in subsequent commands.

```
export SAMLREQUEST="https://my-pf-baseurl/idp/SSO.saml2?..."
```

- d. Export the response and cookie, and echo the response:

```
RESPONSE=$(curl -c - "$SAMLREQUEST")
```

```
echo "$RESPONSE" | grep 'input type="hidden" name="pf.adapterId"
id="pf.adapterId"'
```

- e. Export the 'pf.adapterId' value, and echo the response:

```
export ADAPTER='myAdapter'
```

```
echo "$RESPONSE" | grep 'base'
```

- f. Export the 'href' value (remove the trailing slash /), and echo the response:

```
export BASEURL='https://my-pf-baseurl'
```



```
echo "$RESPONSE" | grep 'form method="POST"'
```

g. Export the 'action' value:

```
export SSOPING='/idp/.../resumeSAML20/idp/SSO.ping'
```

h. Send cookies along with credentials:

```
curl -b <(echo "$RESPONSE") -X POST "$BASEURL$SSOPING" \
--data "pf.username=$SAMLUSER&pf.pass=
$SAMPLPASSWORD&pf.ok=clicked&pf.cancel=&pf.adapterId=$ADAPTER"
--include
```

i. Save the SAMLResponse from the hidden field:

```
export SAMLResponse='PHNhbWxwOlJlc3BvbnN...1scDpSZXNwb25zZT4='
```

j. Using the saved SAMLResponse, make a StorageGRID/api/saml-response request to generate a StorageGRID authentication token.

For RelayState, use the tenant account ID or use 0 if you want to sign in to the Grid Management API.

```
curl -X POST "https://$STORAGEGRID_ADDRESS:443/api/saml-response" \
-H "accept: application/json" \
--data-urlencode "SAMLResponse=$SAMLResponse" \
--data-urlencode "RelayState=$TENANTACCOUNTID" \
| python -m json.tool
```

The response includes the authentication token.

```
{
  "apiVersion": "3.0",
  "data": "56eb07bf-21f6-40b7-af0b-5c6cacfb25e7",
  "responseTime": "2018-11-07T21:32:53.486Z",
  "status": "success"
}
```

k. Save the authentication token in the response as MYTOKEN.

```
export MYTOKEN="56eb07bf-21f6-40b7-af0b-5c6cacfb25e7"
```

You can now use MYTOKEN for other requests, similar to how you would use the API if SSO was not being used.

Sign out of the API if single sign-on is enabled

If single sign-on (SSO) has been enabled, you must issue a series of API requests to sign out of the Grid Management API or the Tenant Management API.

These instructions apply if you are using PingFederate as the SSO identity provider

About this task

If required, you can sign out of the StorageGRID API by logging out from your organization's single logout page. Or, you can trigger single logout (SLO) from StorageGRID, which requires a valid StorageGRID bearer token.

Steps

1. To generate a signed logout request, pass `cookie "sso=true"` to the SLO API:

```
curl -k -X DELETE "https://$STORAGEGRID_ADDRESS/api/v3/authorize" \
-H "accept: application/json" \
-H "Authorization: Bearer $MYTOKEN" \
--cookie "sso=true" \
| python -m json.tool
```

A logout URL is returned:

```
{
  "apiVersion": "3.0",
  "data": "https://my-ping-
url/idp/SLO.saml2?SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3D",
  "responseTime": "2021-10-12T22:20:30.839Z",
  "status": "success"
}
```

2. Save the logout URL.

```
export LOGOUT_REQUEST='https://my-ping-
url/idp/SLO.saml2?SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3D'
```

3. Send a request to the logout URL to trigger SLO and to redirect back to StorageGRID.

```
curl --include "$LOGOUT_REQUEST"
```

The 302 response is returned. The redirect location is not applicable to API-only logout.

```
HTTP/1.1 302 Found
Location: https://$STORAGEGRID_ADDRESS:443/api/saml-logout?SAMLResponse=fVLLasMwEPwVo7ss%...%23rsa-sha256
Set-Cookie: PF=QoKs...SgCC; Path=/; Secure; HttpOnly; SameSite=None
```

4. Delete the StorageGRID bearer token.

Deleting the StorageGRID bearer token works the same way as without SSO. If `cookie "sso=true" is not provided, the user is logged out of StorageGRID without affecting the SSO state.

```
curl -X DELETE "https://$STORAGEGRID_ADDRESS/api/v3/authorize" \
-H "accept: application/json" \
-H "Authorization: Bearer $MYTOKEN" \
--include
```

A 204 No Content response indicates the user is now signed out.

```
HTTP/1.1 204 No Content
```

Deactivate features with the API

You can use the Grid Management API to completely deactivate certain features in the StorageGRID system. When a feature is deactivated, no one can be assigned permissions to perform the tasks related to that feature.

About this task

The Deactivated Features system allows you to prevent access to certain features in the StorageGRID system. Deactivating a feature is the only way to prevent the root user or users who belong to admin groups with **Root access** permission from being able to use that feature.

To understand how this functionality might be useful, consider the following scenario:

Company A is a service provider who leases the storage capacity of their StorageGRID system by creating tenant accounts. To protect the security of their leaseholders' objects, Company A wants to ensure that its own employees can never access any tenant account after the account has been deployed.

*Company A can accomplish this goal by using the Deactivate Features system in the Grid Management API. By completely deactivating the **Change tenant root password** feature in the Grid Manager (both the UI and the API), Company A ensures that Admin users—including the root user and users belonging to groups with the **Root access** permission—can't change the password for any tenant account's root user.*

Steps

1. Access the Swagger documentation for the Grid Management API. See [Use the Grid Management API](#).
2. Locate the Deactivate Features endpoint.

3. To deactivate a feature, such as Change tenant root password, send a body to the API like this:

```
{ "grid": {"changeTenantRootPassword": true} }
```

When the request is complete, the Change tenant root password feature is disabled. The **Change tenant root password** management permission no longer appears in the user interface, and any API request that attempts to change the root password for a tenant will fail with "403 Forbidden."

Reactivate deactivated features

By default, you can use the Grid Management API to reactivate a feature that has been deactivated. However, if you want to prevent deactivated features from ever being reactivated, you can deactivate the **activateFeatures** feature itself.



The **activateFeatures** feature can't be reactivated. If you decide to deactivate this feature, be aware that you will permanently lose the ability to reactivate any other deactivated features. You must contact technical support to restore any lost functionality.

Steps

1. Access the Swagger documentation for the Grid Management API.
2. Locate the Deactivate Features endpoint.
3. To reactivate all features, send a body to the API like this:

```
{ "grid": null }
```

When this request is complete, all features, including the Change tenant root password feature, are reactivated. The **Change tenant root password** management permission now appears in the user interface, and any API request that attempts to change the root password for a tenant will succeed, assuming the user has the **Root access** or **Change tenant root password** management permission.



The previous example causes *all* deactivated features to be reactivated. If other features have been deactivated that should remain deactivated, you must explicitly specify them in the PUT request. For example, to reactivate the Change tenant root password feature and continue to deactivate the storageAdmin management permission, send this PUT request:

```
{ "grid": {"storageAdmin": true} }
```

Control access to StorageGRID

Control StorageGRID access

You control who can access StorageGRID and which tasks users can perform by creating or importing groups and users and assigning permissions to each group. Optionally, you can enable single sign-on (SSO), create client certificates, and change grid passwords.

Control access to the Grid Manager

You determine who can access the Grid Manager and the Grid Management API by importing groups and users from an identity federation service or by setting up local groups and local users.

Using [identity federation](#) makes setting up [groups](#) and [users](#) faster, and it allows users to sign in to

StorageGRID using familiar credentials. You can configure identity federation if you use Active Directory, OpenLDAP, or Oracle Directory Server.



Contact technical support if you want to use another LDAP v3 service.

You determine which tasks each user can perform by assigning different [permissions](#) to each group. For example, you might want users in one group to be able to manage ILM rules and users in another group to perform maintenance tasks. A user must belong to at least one group to access the system.

Optionally, you can configure a group to be read-only. Users in a read-only group can only view settings and features. They can't make any changes or perform any operations in the Grid Manager or Grid Management API.

Enable single sign-on

The StorageGRID system supports single sign-on (SSO) using the Security Assertion Markup Language 2.0 (SAML 2.0) standard. After you [configure and enable SSO](#), all users must be authenticated by an external identity provider before they can access the Grid Manager, the Tenant Manager, the Grid Management API, or the Tenant Management API. Local users can't sign in to StorageGRID.

Change provisioning passphrase

The provisioning passphrase is required for many installation and maintenance procedures, and for downloading the StorageGRID Recovery Package. The passphrase is also required to download backups of the grid topology information and encryption keys for the StorageGRID system. You can [change the passphrase](#) as required.

Change node console passwords

Each node in your grid has a unique node console password, which you need to log in to the node as "admin" using SSH, or to the root user on a VM/physical console connection. As needed, you can [change the node console password](#) for each node.

Change the provisioning passphrase

Use this procedure to change the StorageGRID provisioning passphrase. The passphrase is required for recovery, expansion, and maintenance procedures. The passphrase is also required to download Recovery Package backups that include the grid topology information, grid node console passwords, and encryption keys for the StorageGRID system.

Before you begin

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

About this task

The provisioning passphrase is required for many installation and maintenance procedures, and for [downloading the Recovery Package](#). The provisioning passphrase is not listed in the `Passwords.txt` file. Make sure to document the provisioning passphrase and keep it in a safe and secure location.

Steps

1. Select **CONFIGURATION > Access control> Grid passwords**.
2. Under **Change provisioning passphrase**, select **Make a change**
3. Enter your current provisioning passphrase.
4. Enter the new passphrase. The passphrase must contain at least 8 and no more than 32 characters. Passphrases are case-sensitive.
5. Store the new provisioning passphrase in a secure location. It is required for installation, expansion, and maintenance procedures.
6. Re-enter the new passphrase, and select **Save**.

The system displays a green success banner when the provisioning passphrase change is complete.

 Provisioning passphrase successfully changed. Go to the [Recovery Package](#) to download a new Recovery Package.

7. Select **Recovery Package**.
8. Enter the new provisioning passphrase to download the new Recovery Package.



After changing the provisioning passphrase, you must immediately download a new Recovery Package. The Recovery Package file allows you to restore the system if a failure occurs.

Change node console passwords

Each node in your grid has a unique node console password, which you need to log in to the node. Use these steps to change each unique node console password for each node in your grid.

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 the current provisioning passphrase.

About this task

Use the node console password to log in to a node as "admin" using SSH, or to the root user on a VM/physical console connection. The change node console password process creates new passwords for each node in your grid and stores the passwords in an updated `Passwords.txt` file in the Recovery Package. The passwords are listed in the Password column in the Passwords.txt file.



There are separate SSH access passwords for the SSH keys used for communication between nodes. The SSH access passwords aren't changed by this procedure.

Access the wizard

Steps

1. Select **CONFIGURATION > Access control > Grid passwords**.
2. Under **Change node console passwords**, select **Make a change**.

Enter the provisioning passphrase

Steps

1. Enter the provisioning passphrase for your grid.
2. Select **Continue**.

Download the current recovery package

Before changing node console passwords, download the current Recovery Package. You can use the passwords in this file if the password change process fails for any node.

Steps

1. Select **Download recovery package**.
2. Copy the Recovery Package file (.zip) to two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

3. Select **Continue**.
4. When the confirmation dialog appears, select **Yes** if you are ready to start changing the node console passwords.

You can't cancel this process after it starts.

Change node console passwords

When the node console password process starts, a new Recovery Package is generated that includes the new passwords. Then, the passwords are updated on each node.

Steps

1. Wait for the new Recovery Package to be generated, which might take a few minutes.
2. Select **Download new recovery package**.
3. When the download completes:
 - a. Open the .zip file.
 - b. Confirm that you can access the contents, including the `Passwords.txt` file, which contains the new node console passwords.
 - c. Copy the new Recovery Package file (.zip) to two safe, secure, and separate locations.



Don't overwrite the old Recovery Package.

The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

4. Select the checkbox to indicate you have downloaded the new Recovery Package and verified the content.
5. Select **Change node console passwords** and wait for all nodes to be updated with the new passwords. This might take a few minutes.

If passwords are changed for all nodes, a green success banner appears. Go to the next step.

If there is an error during the update process, a banner message lists the number of nodes that failed to have their passwords changed. The system will automatically retry the process on any node that failed to have its password changed. If the process ends with some nodes still not having a changed password, the **Retry** button appears.

If the password update failed for one or more nodes:

- a. Review the error messages listed in the table.
- b. Resolve the issues.
- c. Select **Retry**.



Retrying only changes the node console passwords on the nodes that failed during previous password change attempts.

6. After node console passwords have been changed for all nodes, delete the [first Recovery Package you downloaded](#).
7. Optionally, use the **Recovery package** link to download an additional copy of the new Recovery Package.

Change SSH access passwords for Admin Nodes

Changing the SSH access passwords for Admin Nodes also updates the unique sets of internal SSH keys for each node in the grid. The primary Admin Node uses these SSH keys to access nodes using secure, passwordless authentication.

Use an SSH key to log in to a node as `admin` or to the root user on a VM or physical console connection.

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 the current provisioning passphrase.

About this task

The new access passwords for Admin Nodes and the new internal keys for each node are stored in the `Passwords.txt` file in the Recovery Package. The keys are listed in the Password column in that file.

There are separate SSH access passwords for the SSH keys used for communication between nodes. Those aren't changed by this procedure.

Access the wizard

Steps

1. Select **CONFIGURATION > Access control > Grid passwords**.
2. Under **Change SSH keys**, select **Make a change**.

Download the current recovery package

Before changing SSH access keys, download the current Recovery Package. You can use the keys in this file if the key change process fails for any node.

Steps

1. Enter the provisioning passphrase for your grid.
2. Select **Download recovery package**.
3. Copy the Recovery Package file (.zip) to two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

4. Select **Continue**.
5. When the confirmation dialog appears, select **Yes** if you are ready to start changing the SSH access keys.



You can't cancel this process after it starts.

Change SSH access keys

When the change SSH access keys process starts, a new Recovery Package is generated that includes the new keys. Then, the keys are updated on each node.

Steps

1. Wait for the new Recovery Package to be generated, which might take a few minutes.
2. When the Download new Recovery Package button is enabled, select **Download new Recovery Package** and save the new Recovery Package file (.zip) to two safe, secure, and separate locations.
3. When the download completes:
 - a. Open the .zip file.
 - b. Confirm that you can access the contents, including the `Passwords.txt` file, which contains the new SSH access keys.
 - c. Copy the new Recovery Package file (.zip) to two safe, secure, and separate locations.



Don't overwrite the old Recovery Package.

The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

4. Wait for the keys to update on each node, which might take a few minutes.

If keys are changed for all nodes, a green success banner appears.

If there is an error during the update process, a banner message lists the number of nodes that failed to have their keys changed. The system will automatically retry the process on any node that failed to have its key changed. If the process ends with some nodes still not having a changed key, the **Retry** button appears.

If the key update failed for one or more nodes:

- a. Review the error messages listed in the table.
- b. Resolve the issues.
- c. Select **Retry**.

Retrying only changes the SSH access keys on the nodes that failed during previous key change attempts.

5. After SSH access keys have been changed for all nodes, delete the [first Recovery Package you downloaded](#).
6. Optionally, select **MAINTENANCE > System > Recovery package** to download an additional copy of the new Recovery Package.

Use identity federation

Using identity federation makes setting up groups and users faster, and it allows users to sign in to StorageGRID using familiar credentials.

Configure identity federation for Grid Manager

You can configure identity federation in the Grid Manager if you want admin groups and users to be managed in another system such as Active Directory, Azure Active Directory (Azure AD), OpenLDAP, or Oracle Directory Server.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- You are using Active Directory, Azure AD, OpenLDAP, or Oracle Directory Server as the identity provider.



If you want to use an LDAP v3 service that is not listed, contact technical support.

- If you plan to use OpenLDAP, you must configure the OpenLDAP server. See [Guidelines for configuring an OpenLDAP server](#).
- If you plan to enable single sign-on (SSO), you have reviewed the [requirements and considerations for single sign-on](#).
- If you plan to use Transport Layer Security (TLS) for communications with the LDAP server, the identity provider is using TLS 1.2 or 1.3. See [Supported ciphers for outgoing TLS connections](#).

About this task

You can configure an identity source for the Grid Manager if you want to import groups from another system such as Active Directory, Azure AD, OpenLDAP, or Oracle Directory Server. You can import the following types of groups:

- Admin groups. The users in admin groups can sign in to the Grid Manager and perform tasks, based on the management permissions assigned to the group.
- Tenant user groups for tenants that don't use their own identity source. Users in tenant groups can sign in to the Tenant Manager and perform tasks, based on the permissions assigned to the group in the Tenant Manager. See [Create tenant account](#) and [Use a tenant account](#) for details.

Enter the configuration

Steps

1. Select **CONFIGURATION > Access control > Identity federation**.
2. Select **Enable identity federation**.

3. In the LDAP service type section, select the type of LDAP service you want to configure.

LDAP service type

Select the type of LDAP service you want to configure.

Active Directory	Azure	OpenLDAP	Other

Select **Other** to configure values for an LDAP server that uses Oracle Directory Server.

4. If you selected **Other**, complete the fields in the LDAP Attributes section. Otherwise, go to the next step.
- **User Unique Name:** The name of the attribute that contains the unique identifier of an LDAP user. This attribute is equivalent to `sAMAccountName` for Active Directory and `uid` for OpenLDAP. If you are configuring Oracle Directory Server, enter `uid`.
 - **User UUID:** The name of the attribute that contains the permanent unique identifier of an LDAP user. This attribute is equivalent to `objectGUID` for Active Directory and `entryUUID` for OpenLDAP. If you are configuring Oracle Directory Server, enter `nsuniqueid`. Each user's value for the specified attribute must be a 32-digit hexadecimal number in either 16-byte or string format, where hyphens are ignored.
 - **Group Unique Name:** The name of the attribute that contains the unique identifier of an LDAP group. This attribute is equivalent to `sAMAccountName` for Active Directory and `cn` for OpenLDAP. If you are configuring Oracle Directory Server, enter `cn`.
 - **Group UUID:** The name of the attribute that contains the permanent unique identifier of an LDAP group. This attribute is equivalent to `objectGUID` for Active Directory and `entryUUID` for OpenLDAP. If you are configuring Oracle Directory Server, enter `nsuniqueid`. Each group's value for the specified attribute must be a 32-digit hexadecimal number in either 16-byte or string format, where hyphens are ignored.
5. For all LDAP service types, enter the required LDAP server and network connection information in the Configure LDAP server section.
- **Hostname:** The fully qualified domain name (FQDN) or IP address of the LDAP server.
 - **Port:** The port used to connect to the LDAP server.



The default port for STARTTLS is 389, and the default port for LDAPS is 636. However, you can use any port as long as your firewall is configured correctly.

- **Username:** The full path of the distinguished name (DN) for the user that will connect to the LDAP server.

For Active Directory, you can also specify the Down-Level Logon Name or the User Principal Name.

The specified user must have permission to list groups and users and to access the following attributes:

- `sAMAccountName` or `uid`
- `objectGUID`, `entryUUID`, or `nsuniqueid`

- `cn`
 - `memberOf` or `isMemberOf`
 - **Active Directory:** `objectSid`, `primaryGroupID`, `userAccountControl`, and `userPrincipalName`
 - **Azure:** `accountEnabled` and `userPrincipalName`
- **Password:** The password associated with the username.



If you change the password in the future, you must update it on this page.

- **Group Base DN:** The full path of the distinguished name (DN) for an LDAP subtree you want to search for groups. In the Active Directory example (below), all groups whose Distinguished Name is relative to the base DN (`DC=storagegrid,DC=example,DC=com`) can be used as federated groups.



The **Group unique name** values must be unique within the **Group Base DN** they belong to.

- **User Base DN:** The full path of the distinguished name (DN) of an LDAP subtree you want to search for users.



The **User unique name** values must be unique within the **User Base DN** they belong to.

- **Bind username format** (optional): The default username pattern StorageGRID should use if the pattern can't be determined automatically.

Providing **Bind username format** is recommended because it can allow users to sign in if StorageGRID is unable to bind with the service account.

Enter one of these patterns:

- **UserPrincipalName pattern (Active Directory and Azure):** `[USERNAME]@example.com`
- **Down-level logon name pattern (Active Directory and Azure):** `example\[USERNAME]`
- **Distinguished name pattern:** `CN=[USERNAME],CN=Users,DC=example,DC=com`

Include **[USERNAME]** exactly as written.

6. In the Transport Layer Security (TLS) section, select a security setting.

- **Use STARTTLS:** Use STARTTLS to secure communications with the LDAP server. This is the recommended option for Active Directory, OpenLDAP, or Other, but this option is not supported for Azure.
- **Use LDAPS:** The LDAPS (LDAP over SSL) option uses TLS to establish a connection to the LDAP server. You must select this option for Azure.
- **Do not use TLS:** The network traffic between the StorageGRID system and the LDAP server will not be secured. This option is not supported for Azure.



Using the **Do not use TLS** option is not supported if your Active Directory server enforces LDAP signing. You must use STARTTLS or LDAPS.

7. If you selected STARTTLS or LDAPS, choose the certificate used to secure the connection.

- **Use operating system CA certificate:** Use the default Grid CA certificate installed on the operating system to secure connections.
- **Use custom CA certificate:** Use a custom security certificate.

If you select this setting, copy and paste the custom security certificate into the CA certificate text box.

Test the connection and save the configuration

After entering all values, you must test the connection before you can save the configuration. StorageGRID verifies the connection settings for the LDAP server and the bind username format, if you provided one.

Steps

1. Select **Test connection**.
2. If you did not provide a bind username format:
 - A "Test connection successful" message appears if the connection settings are valid. Select **Save** to save the configuration.
 - A "test connection could not be established" message appears if the connection settings are invalid. Select **Close**. Then, resolve any issues and test the connection again.
3. If you provided a bind username format, enter the username and password of a valid federated user.

For example, enter your own username and password. Don't include any special characters in the username, such as @ or /.

Test Connection

To test the connection and the bind username format, enter the username and password of a federated user. For example, enter your own federated username and password. The test values are not saved.

Test username

The username of a federated user.

Test password

[Cancel](#) [Test Connection](#)

- A "Test connection successful" message appears if the connection settings are valid. Select **Save** to save the configuration.
- An error message appears if the connection settings, bind username format, or test username and password are invalid. Resolve any issues and test the connection again.

Force synchronization with the identity source

The StorageGRID system periodically synchronizes federated groups and users from the identity source. You can force synchronization to start if you want to enable or restrict user permissions as quickly as possible.

Steps

1. Go to the Identity federation page.
2. Select **Sync server** at the top of the page.

The synchronization process might take some time depending on your environment.



The **Identity federation synchronization failure** alert is triggered if there is an issue synchronizing federated groups and users from the identity source.

Disable identity federation

You can temporarily or permanently disable identity federation for groups and users. When identity federation is disabled, there is no communication between StorageGRID and the identity source. However, any settings you have configured are retained, allowing you to easily reenable identity federation in the future.

About this task

Before you disable identity federation, you should be aware of the following:

- Federated users will be unable to sign in.
- Federated users who are currently signed in will retain access to the StorageGRID system until their session expires, but they will be unable to sign in after their session expires.
- Synchronization between the StorageGRID system and the identity source will not occur, and alerts will not be raised for accounts that have not been synchronized.
- The **Enable identity federation** checkbox is disabled if single sign-on (SSO) is set to **Enabled** or **Sandbox Mode**. The SSO Status on the Single Sign-on page must be **Disabled** before you can disable identity federation. See [Disable single sign-on](#).

Steps

1. Go to the Identity federation page.
2. Uncheck the **Enable identity federation** checkbox.

Guidelines for configuring an OpenLDAP server

If you want to use an OpenLDAP server for identity federation, you must configure specific settings on the OpenLDAP server.



For identity sources that aren't ActiveDirectory or Azure, StorageGRID will not automatically block S3 access to users who are disabled externally. To block S3 access, delete any S3 keys for the user or remove the user from all groups.

Memberof and refint overlays

The memberof and refint overlays should be enabled. For more information, see the instructions for reverse group membership maintenance in the [OpenLDAP documentation: Version 2.4 Administrator's Guide](#).

Indexing

You must configure the following OpenLDAP attributes with the specified index keywords:

- `olcDbIndex: objectClass eq`
- `olcDbIndex: uid eq,pres,sub`
- `olcDbIndex: cn eq,pres,sub`
- `olcDbIndex: entryUUID eq`

In addition, ensure the fields mentioned in the help for Username are indexed for optimal performance.

See the information about reverse group membership maintenance in the [OpenLDAP documentation: Version 2.4 Administrator's Guide](#).

Manage admin groups

You can create admin groups to manage the security permissions for one or more admin users. Users must belong to a group to be granted access to the StorageGRID system.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- If you plan to import a federated group, you have configured identity federation and the federated group already exists in the configured identity source.

Create an admin group

Admin groups allow you to determine which users can access which features and operations in the Grid Manager and the Grid Management API.

Access the wizard

Steps

1. Select **CONFIGURATION > Access control > Admin groups**.
2. Select **Create group**.

Choose a group type

You can create a local group or import a federated group.

- Create a local group if you want to assign permissions to local users.
- Create a federated group to import users from the identity source.

Local group

Steps

1. Select **Local group**.
2. Enter a display name for the group, which you can update later as required. For example, "Maintenance Users" or "ILM Administrators."
3. Enter a unique name for the group, which you can't update later.
4. Select **Continue**.

Federated group

Steps

1. Select **Federated group**.
2. Enter the name of the group you want to import, exactly as it appears in the configured identity source.
 - For Active Directory and Azure, use the sAMAccountName.
 - For OpenLDAP, use the CN (Common Name).
 - For another LDAP, use the appropriate unique name for the LDAP server.
3. Select **Continue**.

Manage group permissions

Steps

1. For **Access mode**, select whether users in the group can change settings and perform operations in the Grid Manager and the Grid Management API or whether they can only view settings and features.
 - **Read-write** (default): Users can change settings and perform the operations allowed by their management permissions.
 - **Read-only**: Users can only view settings and features. They can't make any changes or perform any operations in the Grid Manager or Grid Management API. Local read-only users can change their own passwords.



If a user belongs to multiple groups and any group is set to **Read-only**, the user will have read-only access to all selected settings and features.

2. Select one or more [admin group permissions](#).

You must assign at least one permission to each group; otherwise, users belonging to the group will not be able to sign in to StorageGRID.

3. If you are creating a local group, select **Continue**. If you are creating a federated group, select **Create group** and **Finish**.

Add users (local groups only)

Steps

1. Optionally, select one or more local users for this group.

If you have not yet created local users, you can save the group without adding users. You can add this

group to the user on the Users page. See [Manage users](#) for details.

2. Select **Create group** and **Finish**.

View and edit admin groups

You can view details for existing groups, modify a group, or duplicate a group.

- To view basic information for all groups, review the table on the Groups page.
- To view all details for a specific group or to edit a group, use the **Actions** menu or the details page.

Task	Actions menu	Details page
View group details	a. Select the checkbox for the group. b. Select Actions > View group details.	Select the group name in the table.
Edit display name (local groups only)	a. Select the checkbox for the group. b. Select Actions > Edit group name. c. Enter the new name. d. Select Save changes.	a. Select the group name to display the details. b. Select the edit icon . c. Enter the new name. d. Select Save changes.
Edit access mode or permissions	a. Select the checkbox for the group. b. Select Actions > View group details. c. Optionally, change the group's Access mode. d. Optionally, select or clear admin group permissions. e. Select Save changes.	a. Select the group name to display the details. b. Optionally, change the group's Access mode. c. Optionally, select or clear admin group permissions. d. Select Save changes.

Duplicate a group

Steps

1. Select the checkbox for the group.
2. Select **Actions > Duplicate group**.
3. Complete the Duplicate group wizard.

Delete a group

You can delete an admin group when you want to remove the group from the system, and remove all permissions associated with the group. Deleting an admin group removes any users from the group, but does not delete the users.

Steps

1. From the Groups page, select the checkbox for each group you want to remove.
2. Select **Actions > Delete group**.
3. Select **Delete groups**.

Admin group permissions

When creating admin user groups, you select one or more permissions to control access to specific features of the Grid Manager. You can then assign each user to one or more of these admin groups to determine which tasks that user can perform.

You must assign at least one permission to each group; otherwise, users belonging to that group will not be able to sign in to the Grid Manager or the Grid Management API.

By default, any user who belongs to a group that has at least one permission can perform the following tasks:

- Sign in to the Grid Manager
- View the dashboard
- View the Nodes pages
- View current and resolved alerts
- Change their own password (local users only)
- View certain information provided on the Configuration and Maintenance pages

Interaction between permissions and Access mode

For all permissions, the group's **Access mode** setting determines whether users can change settings and perform operations or whether they can only view the related settings and features. If a user belongs to multiple groups and any group is set to **Read-only**, the user will have read-only access to all selected settings and features.

The following sections describe the permissions you can assign when creating or editing an admin group. Any functionality not explicitly mentioned requires the **Root access** permission.

Root access

This permission provides access to all grid administration features.

Change tenant root password

This permission provides access to the **Change root password** option on the Tenants page, allowing you to control who can change the password for the tenant's local root user. This permission is also used for migrating S3 keys when the S3 key import feature is enabled. Users who don't have this permission can't see the **Change root password** option.



To grant access to the Tenants page, which contains the **Change root password** option, also assign the **Tenant accounts** permission.

Grid topology page configuration

This permission provides access to the Configuration tabs on the **SUPPORT > Tools > Grid topology** page.



The Grid topology page has been deprecated and will be removed in a future release.

ILM

This permission provides access to the following **ILM** menu options:

- Rules
- Policies
- Policy tags
- Storage pools
- Storage grades
- Regions
- Object metadata lookup



Users must have the **Other grid configuration** and **Grid topology page configuration** permissions to manage storage grades.

Maintenance

Users must have the Maintenance permission to use these options:

- **CONFIGURATION > Access control:**
 - Grid passwords
- **CONFIGURATION > Network:**
 - S3 endpoint domain names
- **MAINTENANCE > Tasks:**
 - Decommission
 - Expansion
 - Object existence check
 - Recovery
- **MAINTENANCE > System:**
 - Recovery package
 - Software update
- **SUPPORT > Tools:**
 - Logs

Users who don't have the Maintenance permission can view, but not edit, these pages:

- **MAINTENANCE > Network:**
 - DNS servers
 - Grid Network
 - NTP servers
- **MAINTENANCE > System:**

- License
- **CONFIGURATION > Network:**
 - S3 endpoint domain names
- **CONFIGURATION > Security:**
 - Certificates
- **CONFIGURATION > Monitoring:**
 - Audit and syslog server

Manage alerts

This permission provides access to options for managing alerts. Users must have this permission to manage silences, alert notifications, and alert rules.

Metrics query

This permission provides access to:

- **SUPPORT > Tools > Metrics** page
- Custom Prometheus metrics queries using the **Metrics** section of the Grid Management API
- Grid Manager dashboard cards that contain metrics

Object metadata lookup

This permission provides access to the **ILM > Object metadata lookup** page.

Other grid configuration

This permission provides access to additional grid configuration options.



To see these additional options, users must also have the **Grid topology page configuration** permission.

- **ILM:**
 - Storage grades
- **CONFIGURATION > System:**
- **SUPPORT > Other:**
 - Link cost

Storage appliance administrator

This permission provides:

- Access to the E-Series SANtricity System Manager on storage appliances through the Grid Manager.
- The ability to perform troubleshooting and maintenance tasks on the Manage drives tab for appliances that support these operations.

Tenant accounts

This permission provides the ability to:

- Access the Tenants page, where you can create, edit, and remove tenant accounts
- View existing traffic classification policies
- View Grid Manager dashboard cards that contain tenant details

Manage users

You can view local and federated users. You can also create local users and assign them to local admin groups to determine which Grid Manager features these users can access.

Before you begin

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

Create a local user

You can create one or more local users and assign each user to one or more local groups. The group's permissions control which Grid Manager and Grid Management API features the user can access.

You can create local users only. Use the external identity source to manage federated users and groups.

The Grid Manager includes one predefined local user, named "root." You can't remove the root user.



If single sign-on (SSO) is enabled, local users can't sign in to StorageGRID.

Access the wizard

Steps

1. Select **CONFIGURATION** > **Access control** > **Admin users**.
2. Select **Create user**.

Enter user credentials

Steps

1. Enter the user's full name, a unique username, and a password.
2. Optionally, select **Yes** if this user should not have access to the Grid Manager or Grid Management API.
3. Select **Continue**.

Assign to groups

Steps

1. Optionally, assign the user to one or more groups to determine the user's permissions.

If you have not yet created groups, you can save the user without selecting groups. You can add this user to a group on the Groups page.

If a user belongs to multiple groups, the permissions are cumulative. See

[Manage admin groups](#) for details.

2. Select **Create user** and select **Finish**.

View and edit local users

You can view details for existing local and federated users. You can modify a local user to change the user's full name, password, or group membership. You can also temporarily prevent a user from accessing the Grid Manager and the Grid Management API.

You can edit local users only. Use the external identity source to manage federated users.

- To view basic information for all local and federated users, review the table on the Users page.
- To view all details for a specific user, edit a local user, or change a local user's password, use the **Actions** menu or the details page.

Any edits are applied the next time the user signs out and then signs back in to the Grid Manager.



Local users can change their own passwords using the **Change password** option in the Grid Manager banner.

Task	Actions menu	Details page
View user details	- Select the checkbox for the user. - Select Actions > View user details.	Select the user's name in the table.
Edit full name (local users only)	- Select the checkbox for the user. - Select Actions > Edit full name. - Enter the new name. - Select Save changes.	- Select the user's name to display the details. - Select the edit icon . - Enter the new name. - Select Save changes.
Deny or allow StorageGRID access	- Select the checkbox for the user. - Select Actions > View user details. - Select the Access tab. - Select Yes to prevent the user from signing in to the Grid Manager or the Grid Management API, or select No to allow the user to sign in. - Select Save changes.	- Select the user's name to display the details. - Select the Access tab. - Select Yes to prevent the user from signing in to the Grid Manager or the Grid Management API, or select No to allow the user to sign in. - Select Save changes.
Change password (local users only)	- Select the checkbox for the user. - Select Actions > View user details. - Select the Password tab. - Enter a new password. - Select Change password.	- Select the user's name to display the details. - Select the Password tab. - Enter a new password. - Select Change password.

Task	Actions menu	Details page
Change groups (local users only)	- Select the checkbox for the user. - Select Actions > View user details. - Select the Groups tab. - Optionally, select the link after a group name to view the group's details in a new browser tab. - Select Edit groups to select different groups. - Select Save changes.	- Select the user's name to display the details. - Select the Groups tab. - Optionally, select the link after a group name to view the group's details in a new browser tab. - Select Edit groups to select different groups. - Select Save changes.

Duplicate a user

You can duplicate an existing user to create a new user with the same permissions.

Steps

1. Select the checkbox for the user.
2. Select **Actions > Duplicate user**.
3. Complete the Duplicate user wizard.

Delete a user

You can delete a local user to permanently remove that user from the system.



You can't delete the root user.

Steps

1. From the Users page, select the checkbox for each user you want to remove.
2. Select **Actions > Delete user**.
3. Select **Delete user**.

Use single sign-on (SSO)

Configure single sign-on

When single sign-on (SSO) is enabled, users can only access the Grid Manager, the Tenant Manager, the Grid Management API, or the Tenant Management API if their credentials are authorized using the SSO sign-in process implemented by your organization. Local users can't sign in to StorageGRID.

How single sign-on works

The StorageGRID system supports single sign-on (SSO) using the Security Assertion Markup Language 2.0 (SAML 2.0) standard.

Before enabling single sign-on (SSO), review how the StorageGRID sign-in and sign-out processes are

affected when SSO is enabled.

Sign in when SSO is enabled


When SSO is enabled and you sign in to StorageGRID, you are redirected to your organization's SSO page to validate your credentials.

Steps

1. Enter the fully qualified domain name or IP address of any StorageGRID Admin Node in a web browser.

The StorageGRID Sign in page appears.

- If this is the first time you have accessed the URL on this browser, you are prompted for an account ID:

The image shows a screenshot of the NetApp StorageGRID Sign in page. At the top, the NetApp StorageGRID logo is displayed. Below the logo, the text "Sign in" is prominently shown. Underneath, the word "Account" is followed by a text input field containing the number "0". A blue "Sign in" button is positioned below the input field. At the bottom of the page, there are links for "NetApp support" and "NetApp.com" separated by a vertical bar.

- If you have previously accessed either the Grid Manager or the Tenant Manager, you are prompted to select a recent account or to enter an account ID:



The StorageGRID Sign in page is not shown when you enter the complete URL for a tenant account (that is, a fully qualified domain name or IP address followed by `/?accountId=20-digit-account-id`). Instead, you are immediately redirected to your organization's SSO sign-in page, where you can [sign in with your SSO credentials](#).

2. Indicate whether you want to access the Grid Manager or the Tenant Manager:

- To access the Grid Manager, leave the **Account ID** field blank, enter **0** as the account ID, or select **Grid Manager** if it appears in the list of recent accounts.
- To access the Tenant Manager, enter the 20-digit tenant account ID or select a tenant by name if it appears in the list of recent accounts.

3. Select **Sign in**

StorageGRID redirects you to your organization's SSO sign-in page. For example:

4. Sign in with your SSO credentials.

If your SSO credentials are correct:

- a. The identity provider (IdP) provides an authentication response to StorageGRID.
- b. StorageGRID validates the authentication response.
- c. If the response is valid and you belong to a federated group with StorageGRID access permissions, you are signed in to the Grid Manager or the Tenant Manager, depending on which account you selected.



If the service account is inaccessible, you can still sign in, as long as you are an existing user that belongs to a federated group with StorageGRID access permissions.

- 5. Optionally, access other Admin Nodes, or access the Grid Manager or the Tenant Manager, if you have adequate permissions.

You don't need to reenter your SSO credentials.

Sign out when SSO is enabled

When SSO is enabled for StorageGRID, what happens when you sign out depends on what you are signed in to and where you are signing out from.

Steps

- 1. Locate the **Sign out** link in the top-right corner of the user interface.
- 2. Select **Sign out**.

The StorageGRID Sign in page appears. The **Recent Accounts** drop-down is updated to include **Grid Manager** or the name of the tenant, so you can access these user interfaces more quickly in the future.

If you are signed in to...	And you sign out from...	You are signed out of...
Grid Manager on one or more Admin Nodes	Grid Manager on any Admin Node	Grid Manager on all Admin Nodes Note: If you use Azure for SSO, it might take a few minutes to be signed out of all Admin Nodes.
Tenant Manager on one or more Admin Nodes	Tenant Manager on any Admin Node	Tenant Manager on all Admin Nodes
Both Grid Manager and Tenant Manager	Grid Manager	The Grid Manager only. You must also sign out of the Tenant Manager to sign out of SSO.
	Tenant Manager	The Tenant Manager only. You must also sign out of the Grid Manager to sign out of SSO.



The table summarizes what happens when you sign out if you are using a single browser session. If you are signed in to StorageGRID across multiple browser sessions, you must sign out of all browser sessions separately.

Requirements and considerations for single sign-on

Before enabling single sign-on (SSO) for a StorageGRID system, review the requirements and considerations.

Identity provider requirements

StorageGRID supports the following SSO identity providers (IdP):

- Active Directory Federation Service (AD FS)
- Azure Active Directory (Azure AD)
- PingFederate

You must configure identity federation for your StorageGRID system before you can configure an SSO identity provider. The type of LDAP service you use for identity federation controls which type of SSO you can implement.

Configured LDAP service type	Options for SSO identity provider
Active Directory	• Active Directory • Azure • PingFederate
Azure	Azure

AD FS requirements

You can use any of the following versions of AD FS:

- Windows Server 2022 AD FS
- Windows Server 2019 AD FS
- Windows Server 2016 AD FS



Windows Server 2016 should be using the [KB3201845 update](#), or higher.

Additional requirements

- Transport Layer Security (TLS) 1.2 or 1.3
- Microsoft .NET Framework, version 3.5.1 or higher

Considerations for Azure

If you use Azure as the SSO type and users have user principal names that don't use the sAMAccountName as the prefix, login issues can occur if StorageGRID loses its connection with the LDAP server. To allow users

to sign in, you must restore the connection to the LDAP server.

Server certificate requirements

By default, StorageGRID uses a management interface certificate on each Admin Node to secure access to the Grid Manager, the Tenant Manager, the Grid Management API, and the Tenant Management API. When you configure relying party trusts (AD FS), enterprise applications (Azure), or service provider connections (PingFederate) for StorageGRID, you use the server certificate as the signature certificate for StorageGRID requests.

If you have not already [configured a custom certificate for the management interface](#), you should do so now. When you install a custom server certificate, it is used for all Admin Nodes, and you can use it in all StorageGRID relying party trusts, enterprise applications, or SP connections.



Using an Admin Node's default server certificate in a relying party trust, enterprise application, or SP connection is not recommended. If the node fails and you recover it, a new default server certificate is generated. Before you can sign in to the recovered node, you must update the relying party trust, enterprise application, or SP connection with the new certificate.

You can access an Admin Node's server certificate by logging in to the command shell of the node and going to the `/var/local/mgmt-api` directory. A custom server certificate is named `custom-server.crt`. The node's default server certificate is named `server.crt`.

Port requirements

Single sign-on (SSO) is not available on the restricted Grid Manager or Tenant Manager ports. You must use the default HTTPS port (443) if you want users to authenticate with single sign-on. See [Control access at external firewall](#).

Confirm federated users can sign in

Before you enable single sign-on (SSO), you must confirm that at least one federated user can sign in to the Grid Manager and in to the Tenant Manager for any existing tenant accounts.

Before you begin

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

Steps

1. If there are existing tenant accounts, confirm that none of the tenants is using its own identity source.



When you enable SSO, an identity source configured in the Tenant Manager is overridden by the identity source configured in the Grid Manager. Users belonging to the tenant's identity source will no longer be able to sign in unless they have an account with the Grid Manager identity source.

- a. Sign in to the Tenant Manager for each tenant account.
- b. Select **ACCESS MANAGEMENT > Identity federation**.

- c. Confirm that the **Enable identity federation** checkbox is not selected.
 - d. If it is, confirm that any federated groups that might be in use for this tenant account are no longer required, clear the checkbox, and select **Save**.
2. Confirm that a federated user can access the Grid Manager:
 - a. From Grid Manager, select **CONFIGURATION > Access control > Admin groups**.
 - b. Ensure that at least one federated group has been imported from the Active Directory identity source and that it has been assigned the Root access permission.
 - c. Sign out.
 - d. Confirm you can sign back in to the Grid Manager as a user in the federated group.
3. If there are existing tenant accounts, confirm that a federated user who has Root access permission can sign in:
 - a. From the Grid Manager, select **TENANTS**.
 - b. Select the tenant account, and select **Actions > Edit**.
 - c. On the Enter details tab, select **Continue**.
 - d. If the **Use own identity source** checkbox is selected, uncheck the box and select **Save**.

The screenshot shows the 'Edit the tenant' interface. At the top, a blue header bar contains the title 'Edit the tenant' and a progress indicator with two steps: 'Enter details' (marked with a checkmark) and '2 Select permissions' (marked with a circle containing the number 2). Below the header, the main content area is titled 'Select permissions' and includes the instruction 'Select the permissions for this tenant account.' There are three checkboxes listed: 'Allow platform services' with a help icon, 'Use own identity source' with a help icon, and 'Allow S3 Select' with a help icon. The 'Use own identity source' checkbox is highlighted with a green rectangular border.

The Tenant page appears.

- e. Select the tenant account, select **Sign in**, and sign in to the tenant account as the local root user.
 - f. From the Tenant Manager, select **ACCESS MANAGEMENT > Groups**.
 - g. Ensure that at least one federated group from the Grid Manager has been assigned the Root access permission for this tenant.
 - h. Sign out.

- i. Confirm you can sign back in to the tenant as a user in the federated group.

Related information

- [Requirements and considerations for single sign-on](#)
- [Manage admin groups](#)
- [Use a tenant account](#)

Use sandbox mode

You can use sandbox mode to configure and test single sign-on (SSO) before enabling it for all StorageGRID users. After SSO has been enabled, you can return to sandbox mode whenever you need to change or retest the configuration.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).
- You have configured identity federation for your StorageGRID system.
- For the identity federation **LDAP service type**, you selected either Active Directory or Azure, based on the SSO identity provider you plan to use.

Configured LDAP service type	Options for SSO identity provider
Active Directory	• Active Directory • Azure • PingFederate
Azure	Azure

About this task

When SSO is enabled and a user attempts to sign in to an Admin Node, StorageGRID sends an authentication request to the SSO identity provider. In turn, the SSO identity provider sends an authentication response back to StorageGRID, indicating whether the authentication request was successful. For successful requests:

- The response from Active Directory or PingFederate includes a universally unique identifier (UUID) for the user.
- The response from Azure includes a User Principal Name (UPN).

To allow StorageGRID (the service provider) and the SSO identity provider to communicate securely about user authentication requests, you must configure certain settings in StorageGRID. Next, you must use the SSO identity provider's software to create a relying party trust (AD FS), Enterprise Application (Azure) or Service Provider (PingFederate) for each Admin Node. Finally, you must return to StorageGRID to enable SSO.


Sandbox mode makes it easy to perform this back-and-forth configuration and to test all of your settings before you enable SSO. When you are using sandbox mode, users can't sign in using SSO.

Access sandbox mode

Steps

1. Select **CONFIGURATION** > **Access control** > **Single sign-on**.

The Single Sign-on page appears, with the **Disabled** option selected.



Single Sign-on

You can enable single sign-on (SSO) if you want an external identity provider (IdP) to authorize all user access to StorageGRID. To start, enable [identity federation](#) and confirm that at least one federated user has Root Access permission to the Grid Manager and to the Tenant Manager for any existing tenant accounts. Next, select Sandbox Mode to configure, save, and then test your SSO settings. After verifying the connections, select Enabled and click Save to start using SSO.

SSO status ⓘ ☒ Disabled ☐ Sandbox Mode ☐ Enabled

[Save](#)



If the SSO Status options don't appear, confirm you have configured the identity provider as the federated identity source. See [Requirements and considerations for single sign-on](#).

2. Select **Sandbox Mode**.

The Identity Provider section appears.

Enter identity provider details

Steps

1. Select the **SSO type** from the drop-down list.
2. Complete the fields in the Identity Provider section based on the SSO type you selected.

Active Directory

- a. Enter the **Federation service name** for the identity provider, exactly as it appears in Active Directory Federation Service (AD FS).



To locate the federation service name, go to Windows Server Manager. Select **Tools > AD FS Management**. From the Action menu, select **Edit Federation Service Properties**. The Federation Service Name is shown in the second field.

- b. Specify which TLS certificate will be used to secure the connection when the identity provider sends SSO configuration information in response to StorageGRID requests.

- **Use operating system CA certificate:** Use the default CA certificate installed on the operating system to secure the connection.
- **Use custom CA certificate:** Use a custom CA certificate to secure the connection.

If you select this setting, copy the text of the custom certificate and paste it in the **CA Certificate** text box.

- **Do not use TLS:** Do not use a TLS certificate to secure the connection.



If you change the CA certificate, immediately [restart the mgmt-api service on the Admin Nodes](#) and test for a successful SSO into the Grid Manager.

- c. In the Relying Party section, specify the **Relying party identifier** for StorageGRID. This value controls the name you use for each relying party trust in AD FS.

- For example, if your grid has only one Admin Node and you don't anticipate adding more Admin Nodes in the future, enter `SG` or `StorageGRID`.
- If your grid includes more than one Admin Node, include the string `[HOSTNAME]` in the identifier. For example, `SG-[HOSTNAME]`. This generates a table that shows the relying party identifier for each Admin Node in your system, based on the node's hostname.



You must create a relying party trust for each Admin Node in your StorageGRID system. Having a relying party trust for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

- d. Select **Save**.

A green check mark appears on the **Save** button for a few seconds.



Azure

- a. Specify which TLS certificate will be used to secure the connection when the identity provider sends SSO configuration information in response to StorageGRID requests.

- **Use operating system CA certificate:** Use the default CA certificate installed on the operating system to secure the connection.
- **Use custom CA certificate:** Use a custom CA certificate to secure the connection.

If you select this setting, copy the text of the custom certificate and paste it in the **CA Certificate** text box.

- **Do not use TLS:** Do not use a TLS certificate to secure the connection.



If you change the CA certificate, immediately [restart the mgmt-api service on the Admin Nodes](#) and test for a successful SSO into the Grid Manager.

b. In the Enterprise Application section, specify the **Enterprise application name** for StorageGRID. This value controls the name you use for each enterprise application in Azure AD.

- For example, if your grid has only one Admin Node and you don't anticipate adding more Admin Nodes in the future, enter `SG` or `StorageGRID`.
- If your grid includes more than one Admin Node, include the string `[HOSTNAME]` in the identifier. For example, `SG-[HOSTNAME]`. This generates a table that shows an enterprise application name for each Admin Node in your system, based on the node's hostname.



You must create an enterprise application for each Admin Node in your StorageGRID system. Having an enterprise application for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

- c. Follow the steps in [Create enterprise applications in Azure AD](#) to create an enterprise application for each Admin Node listed in the table.
- d. From Azure AD, copy the federation metadata URL for each enterprise application. Then, paste this URL into the corresponding **Federation metadata URL** field in StorageGRID.
- e. After you have copied and pasted a federation metadata URL for all Admin Nodes, select **Save**.

A green check mark appears on the **Save** button for a few seconds.



PingFederate

a. Specify which TLS certificate will be used to secure the connection when the identity provider sends SSO configuration information in response to StorageGRID requests.

- **Use operating system CA certificate:** Use the default CA certificate installed on the operating system to secure the connection.
- **Use custom CA certificate:** Use a custom CA certificate to secure the connection.

If you select this setting, copy the text of the custom certificate and paste it in the **CA Certificate** text box.

- **Do not use TLS:** Do not use a TLS certificate to secure the connection.



If you change the CA certificate, immediately [restart the mgmt-api service on the Admin Nodes](#) and test for a successful SSO into the Grid Manager.

b. In the Service Provider (SP) section, specify the **SP connection ID** for StorageGRID. This value controls the name you use for each SP connection in PingFederate.

- For example, if your grid has only one Admin Node and you don't anticipate adding more Admin Nodes in the future, enter `SG` or `StorageGRID`.
- If your grid includes more than one Admin Node, include the string `[HOSTNAME]` in the identifier. For example, `SG-[HOSTNAME]`. This generates a table that shows the SP connection ID for each Admin Node in your system, based on the node's hostname.



You must create an SP connection for each Admin Node in your StorageGRID system. Having an SP connection for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

- c. Specify the federation metadata URL for each Admin Node in the **Federation metadata URL** field.

Use the following format:

```
https://<Federation Service  
Name>:<port>/pf/federation_metadata.ping?PartnerSpId=<SP  
Connection ID>
```

- d. Select **Save**.

A green check mark appears on the **Save** button for a few seconds.

Save ✓

Configure relying party trusts, enterprise applications, or SP connections

When the configuration is saved, the Sandbox mode confirmation notice appears. This notice confirms that sandbox mode is now enabled and provides overview instructions.

StorageGRID can remain in sandbox mode as long as required. However, when **Sandbox Mode** is selected on the Single Sign-on page, SSO is disabled for all StorageGRID users. Only local users can sign in.

Follow these steps to configure relying party trusts (Active Directory), complete enterprise applications (Azure), or configure SP connections (PingFederate).

Active Directory

Steps

1. Go to Active Directory Federation Services (AD FS).
2. Create one or more relying party trusts for StorageGRID, using each relying party identifier shown in the table on the StorageGRID Single Sign-on page.

You must create one trust for each Admin Node shown in the table.

For instructions, go to [Create relying party trusts in AD FS](#).

Azure

Steps

1. From the Single sign-on page for the Admin Node you are currently signed in to, select the button to download and save the SAML metadata.
2. Then, for any other Admin Nodes in your grid, repeat these steps:
 - a. Sign in to the node.
 - b. Select **CONFIGURATION > Access control > Single sign-on**.
 - c. Download and save the SAML metadata for that node.
3. Go to the Azure Portal.
4. Follow the steps in [Create enterprise applications in Azure AD](#) to upload the SAML metadata file for each Admin Node into its corresponding Azure enterprise application.

PingFederate

Steps

1. From the Single sign-on page for the Admin Node you are currently signed in to, select the button to download and save the SAML metadata.
2. Then, for any other Admin Nodes in your grid, repeat these steps:
 - a. Sign in to the node.
 - b. Select **CONFIGURATION > Access control > Single sign-on**.
 - c. Download and save the SAML metadata for that node.
3. Go to PingFederate.
4. [Create one or more service provider \(SP\) connections for StorageGRID](#). Use the SP connection ID for each Admin Node (shown in the table on the StorageGRID Single Sign-on page) and the SAML metadata you downloaded for that Admin Node.

You must create one SP connection for each Admin Node shown in the table.

Test SSO connections

Before you enforce the use of single sign-on for your entire StorageGRID system, you should confirm that single sign-on and single logout are correctly configured for each Admin Node.

Active Directory

Steps

1. From the StorageGRID Single Sign-on page, locate the link in the Sandbox mode message.

The URL is derived from the value you entered in the **Federation service name** field.

Sandbox mode

Sandbox mode is currently enabled. Use this mode to configure relying party trusts and to confirm that single sign-on (SSO) and single logout (SLO) are correctly configured for the StorageGRID system.

1. Use Active Directory Federation Services (AD FS) to create relying party trusts for StorageGRID. Create one trust for each Admin Node, using the relying party identifier(s) shown below.
2. Go to your identity provider's sign-on page: <https://ad2016.saml.sgws/adfs/ls/idpinitiatedsignon.htm>
3. From this page, sign in to each StorageGRID relying party trust. If the SSO operation is successful, StorageGRID displays a page with a success message. Otherwise, an error message is displayed.

When you have confirmed SSO for each of the relying party trusts and you are ready to enforce the use of SSO for StorageGRID, change the SSO Status to Enabled, and click Save.

2. Select the link, or copy and paste the URL into a browser, to access your identity provider's sign-on page.
3. To confirm you can use SSO to sign in to StorageGRID, select **Sign in to one of the following sites**, select the relying party identifier for your primary Admin Node, and select **Sign in**.

You are not signed in.

☐ Sign in to this site.

☒ Sign in to one of the following sites:

SG-DC1-ADM1

Sign in

4. Enter your federated username and password.
 - If the SSO sign-in and logout operations are successful, a success message appears.

✓ Single sign-on authentication and logout test completed successfully.

- If the SSO operation is unsuccessful, an error message appears. Fix the issue, clear the browser's cookies, and try again.
5. Repeat these steps to verify the SSO connection for each Admin Node in your grid.

Azure

Steps

1. Go to the Single sign-on page in the Azure portal.
2. Select **Test this application**.
3. Enter the credentials of a federated user.
 - If the SSO sign-in and logout operations are successful, a success message appears.

✓ Single sign-on authentication and logout test completed successfully.

- If the SSO operation is unsuccessful, an error message appears. Fix the issue, clear the browser's cookies, and try again.
4. Repeat these steps to verify the SSO connection for each Admin Node in your grid.

PingFederate

Steps

1. From the StorageGRID Single Sign-on page, select the first link in the Sandbox mode message.

Select and test one link at a time.

2. Enter the credentials of a federated user.
 - If the SSO sign-in and logout operations are successful, a success message appears.

✓ Single sign-on authentication and logout test completed successfully.

- If the SSO operation is unsuccessful, an error message appears. Fix the issue, clear the browser's cookies, and try again.
3. Select the next link to verify the SSO connection for each Admin Node in your grid.

If you see a Page Expired message, select the **Back** button in your browser and resubmit your credentials.

Enable single sign-on

When you have confirmed you can use SSO to sign in to each Admin Node, you can enable SSO for your entire StorageGRID system.



When SSO is enabled, all users must use SSO to access the Grid Manager, the Tenant Manager, the Grid Management API, and the Tenant Management API. Local users can no longer access StorageGRID.

Steps

1. Select **CONFIGURATION > Access control > Single sign-on**.
2. Change the SSO Status to **Enabled**.
3. Select **Save**.
4. Review the warning message, and select **OK**.

Single sign-on is now enabled.



If you are using the Azure Portal and you access StorageGRID from the same computer you use to access Azure, ensure that the Azure Portal user is also an authorized StorageGRID user (a user in a federated group that has been imported into StorageGRID) or log out of the Azure Portal before attempting to sign in to StorageGRID.

Create relying party trusts in AD FS

You must use Active Directory Federation Services (AD FS) to create a relying party trust for each Admin Node in your system. You can create relying party trusts using PowerShell commands, by importing SAML metadata from StorageGRID, or by entering the data manually.

Before you begin

- You have configured single sign-on for StorageGRID and you selected **AD FS** as the SSO type.
- **Sandbox mode** is selected on the Single sign-on page in Grid Manager. See [Use sandbox mode](#).
- You know the fully qualified domain name (or the IP address) and the relying party identifier for each Admin Node in your system. You can find these values in the Admin Nodes detail table on the StorageGRID Single Sign-on page.



You must create a relying party trust for each Admin Node in your StorageGRID system. Having a relying party trust for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

- You have experience creating relying party trusts in AD FS, or you have access to the Microsoft AD FS documentation.
- You are using the AD FS Management snap-in, and you belong to the Administrators group.
- If you are creating the relying party trust manually, you have the custom certificate that was uploaded for the StorageGRID management interface, or you know how to log in to an Admin Node from the command shell.

About this task

These instructions apply to Windows Server 2016 AD FS. If you are using a different version of AD FS, you will notice slight differences in the procedure. See the Microsoft AD FS documentation if you have questions.

Create a relying party trust using Windows PowerShell

You can use Windows PowerShell to quickly create one or more relying party trusts.

Steps

1. From the Windows start menu, right-select the PowerShell icon, and select **Run as Administrator**.
2. At the PowerShell command prompt, enter the following command:

```
Add-AdfsRelyingPartyTrust -Name "Admin_Node_Identifer" -MetadataURL  
"https://Admin_Node_FQDN/api/saml-metadata"
```

- For *Admin_Node_Identifier*, enter the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page. For example, SG-DC1-ADM1.

- For *Admin_Node_FQDN*, enter the fully qualified domain name for the same Admin Node. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be aware that you must update or recreate this relying party trust if that IP address ever changes.)

3. From Windows Server Manager, select **Tools > AD FS Management**.

The AD FS management tool appears.

4. Select **AD FS > Relying Party Trusts**.

The list of relying party trusts appears.

5. Add an Access Control Policy to the newly created relying party trust:

- a. Locate the relying party trust you just created.
- b. Right-click the trust, and select **Edit Access Control Policy**.
- c. Select an Access Control Policy.
- d. Select **Apply**, and select **OK**

6. Add a Claim Issuance Policy to the newly created Relying Party Trust:

- a. Locate the relying party trust you just created.
- b. Right-click the trust, and select **Edit claim issuance policy**.
- c. Select **Add rule**.
- d. On the Select Rule Template page, select **Send LDAP Attributes as Claims** from the list, and select **Next**.
- e. On the Configure Rule page, enter a display name for this rule.

For example, **ObjectGUID to Name ID** or **UPN to Name ID**.

- f. For the Attribute Store, select **Active Directory**.
 - g. In the LDAP Attribute column of the Mapping table, type **objectGUID** or select **User-Principal-Name**.
 - h. In the Outgoing Claim Type column of the Mapping table, select **Name ID** from the drop-down list.
 - i. Select **Finish**, and select **OK**.
7. Confirm that the metadata was imported successfully.
- a. Right-click the relying party trust to open its properties.
 - b. Confirm that the fields on the **Endpoints**, **Identifiers**, and **Signature** tabs are populated.

If the metadata is missing, confirm that the Federation metadata address is correct, or enter the values manually.

8. Repeat these steps to configure a relying party trust for all of the Admin Nodes in your StorageGRID system.

9. When you are done, return to StorageGRID and test all relying party trusts to confirm they are configured correctly. See [Use Sandbox mode](#) for instructions.

Create a relying party trust by importing federation metadata

You can import the values for each relying party trust by accessing the SAML metadata for each Admin Node.

Steps

1. In Windows Server Manager, select **Tools**, and then select **AD FS Management**.
2. Under Actions, select **Add Relying Party Trust**.
3. On the Welcome page, choose **Claims aware**, and select **Start**.
4. Select **Import data about the relying party published online or on a local network**.
5. In **Federation metadata address (host name or URL)**, type the location of the SAML metadata for this Admin Node:

`https://Admin_Node_FQDN/api/saml-metadata`

For *Admin_Node_FQDN*, enter the fully qualified domain name for the same Admin Node. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be aware that you must update or recreate this relying party trust if that IP address ever changes.)

6. Complete the Relying Party Trust wizard, save the relying party trust, and close the wizard.



When entering the display name, use the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page in the Grid Manager. For example, SG-DC1-ADM1.

7. Add a claim rule:
 - a. Right-click the trust, and select **Edit claim issuance policy**.
 - b. Select **Add rule**:
 - c. On the Select Rule Template page, select **Send LDAP Attributes as Claims** from the list, and select **Next**.
 - d. On the Configure Rule page, enter a display name for this rule.

For example, **ObjectGUID to Name ID** or **UPN to Name ID**.
 - e. For the Attribute Store, select **Active Directory**.
 - f. In the LDAP Attribute column of the Mapping table, type **objectGUID** or select **User-Principal-Name**.
 - g. In the Outgoing Claim Type column of the Mapping table, select **Name ID** from the drop-down list.
 - h. Select **Finish**, and select **OK**.

8. Confirm that the metadata was imported successfully.
 - a. Right-click the relying party trust to open its properties.
 - b. Confirm that the fields on the **Endpoints**, **Identifiers**, and **Signature** tabs are populated.

If the metadata is missing, confirm that the Federation metadata address is correct, or enter the values manually.

9. Repeat these steps to configure a relying party trust for all of the Admin Nodes in your StorageGRID system.
10. When you are done, return to StorageGRID and test all relying party trusts to confirm they are configured correctly. See [Use Sandbox mode](#) for instructions.

Create a relying party trust manually

If you choose not to import the data for the relying part trusts, you can enter the values manually.

Steps

1. In Windows Server Manager, select **Tools**, and then select **AD FS Management**.
2. Under Actions, select **Add Relying Party Trust**.
3. On the Welcome page, choose **Claims aware**, and select **Start**.
4. Select **Enter data about the relying party manually**, and select **Next**.
5. Complete the Relying Party Trust wizard:

- a. Enter a display name for this Admin Node.

For consistency, use the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page in the Grid Manager. For example, SG-DC1-ADM1.

- b. Skip the step to configure an optional token encryption certificate.
- c. On the Configure URL page, select the **Enable support for the SAML 2.0 WebSSO protocol** checkbox.
- d. Type the SAML service endpoint URL for the Admin Node:

`https://Admin_Node_FQDN/api/saml-response`

For *Admin_Node_FQDN*, enter the fully qualified domain name for the Admin Node. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be aware that you must update or recreate this relying party trust if that IP address ever changes.)

- e. On the Configure Identifiers page, specify the Relying Party Identifier for the same Admin Node:

Admin_Node_Identifier

For *Admin_Node_Identifier*, enter the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page. For example, SG-DC1-ADM1.

- f. Review the settings, save the relying party trust, and close the wizard.

The Edit Claim Issuance Policy dialog box appears.



If the dialog box does not appear, right-click the trust, and select **Edit claim issuance policy**.

6. To start the Claim Rule wizard, select **Add rule**:
 - a. On the Select Rule Template page, select **Send LDAP Attributes as Claims** from the list, and select **Next**.
 - b. On the Configure Rule page, enter a display name for this rule.

For example, **ObjectGUID to Name ID** or **UPN to Name ID**.
 - c. For the Attribute Store, select **Active Directory**.
 - d. In the LDAP Attribute column of the Mapping table, type **objectGUID** or select **User-Principal-Name**.
 - e. In the Outgoing Claim Type column of the Mapping table, select **Name ID** from the drop-down list.
 - f. Select **Finish**, and select **OK**.

7. Right-click the relying party trust to open its properties.
8. On the **Endpoints** tab, configure the endpoint for single logout (SLO):
 - a. Select **Add SAML**.
 - b. Select **Endpoint Type > SAML Logout**.
 - c. Select **Binding > Redirect**.
 - d. In the **Trusted URL** field, enter the URL used for single logout (SLO) from this Admin Node:

`https://Admin_Node_FQDN/api/saml-logout`

For *Admin_Node_FQDN*, enter the Admin Node's fully qualified domain name. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be aware that you must update or recreate this relying party trust if that IP address ever changes.)

- e. Select **OK**.
9. On the **Signature** tab, specify the signature certificate for this relying party trust:
 - a. Add the custom certificate:
 - If you have the custom management certificate you uploaded to StorageGRID, select that certificate.
 - If you don't have the custom certificate, log in to the Admin Node, go the `/var/local/mgmt-api` directory of the Admin Node, and add the `custom-server.crt` certificate file.



Using the Admin Node's default certificate (`server.crt`) is not recommended. If the Admin Node fails, the default certificate will be regenerated when you recover the node, and you will need to update the relying party trust.

- b. Select **Apply**, and select **OK**.

The Relying Party properties are saved and closed.

10. Repeat these steps to configure a relying party trust for all of the Admin Nodes in your StorageGRID system.
11. When you are done, return to StorageGRID and test all relying party trusts to confirm they are configured correctly. See [Use sandbox mode](#) for instructions.

Create enterprise applications in Azure AD

You use Azure AD to create an enterprise application for each Admin Node in your system.

Before you begin

- You have started configuring single sign-on for StorageGRID and you selected **Azure** as the SSO type.
- **Sandbox mode** is selected on the Single sign-on page in Grid Manager. See [Use sandbox mode](#).
- You have the **Enterprise application name** for each Admin Node in your system. You can copy these values from the Admin Node details table on the StorageGRID Single Sign-on page.



You must create an enterprise application for each Admin Node in your StorageGRID system. Having an enterprise application for each Admin Node ensures that users can securely sign in to and out of any Admin Node.

- You have experience creating enterprise applications in Azure Active Directory.
- You have an Azure account with an active subscription.
- You have one of the following roles in the Azure account: Global Administrator, Cloud Application Administrator, Application Administrator, or owner of the service principal.

Access Azure AD

Steps

1. Log in to the [Azure Portal](#).
2. Navigate to [Azure Active Directory](#).
3. Select [Enterprise applications](#).

Create enterprise applications and save StorageGRID SSO configuration

To save the SSO configuration for Azure in StorageGRID, you must use Azure to create an enterprise application for each Admin Node. You will copy the federation metadata URLs from Azure and paste them into the corresponding **Federation metadata URL** fields on the StorageGRID Single Sign-on page.

Steps

1. Repeat the following steps for each Admin Node.
 - a. In the Azure Enterprise applications pane, select **New application**.
 - b. Select **Create your own application**.
 - c. For the name, enter the **Enterprise application name** you copied from the Admin Node details table on the StorageGRID Single Sign-on page.
 - d. Leave the **Integrate any other application you don't find in the gallery (Non-gallery)** radio button selected.
 - e. Select **Create**.
 - f. Select the **Get started** link in the **2. Set up single sign on** box, or select the **Single sign-on** link in the left margin.
 - g. Select the **SAML** box.
 - h. Copy the **App Federation Metadata Url**, which you can find under **Step 3 SAML Signing Certificate**.
 - i. Go to the StorageGRID Single Sign-on page, and paste the URL in the **Federation metadata URL** field that corresponds to the **Enterprise application name** you used.
2. After you have pasted a federation metadata URL for each Admin Node and made all other needed changes to the SSO configuration, select **Save** on the StorageGRID Single Sign-on page.

Download SAML metadata for every Admin Node

After the SSO configuration is saved, you can download a SAML metadata file for each Admin Node in your StorageGRID system.

Steps

1. Repeat these steps for each Admin Node.
 - a. Sign in to StorageGRID from the Admin Node.
 - b. Select **CONFIGURATION > Access control > Single sign-on**.
 - c. Select the button to download the SAML metadata for that Admin Node.
 - d. Save the file, which you will upload into Azure AD.

Upload SAML metadata to each enterprise application

After downloading a SAML metadata file for each StorageGRID Admin Node, perform the following steps in Azure AD:

Steps

1. Return to the Azure Portal.
2. Repeat these steps for each enterprise application:



You might need to refresh the Enterprise applications page to see applications you previously added in the list.

- a. Go to the Properties page for the enterprise application.
 - b. Set **Assignment required** to **No** (unless you want to separately configure assignments).
 - c. Go to the Single sign-on page.
 - d. Complete the SAML configuration.
 - e. Select the **Upload metadata file** button and select the SAML metadata file you downloaded for the corresponding Admin Node.
 - f. After the file loads, select **Save** and then select **X** to close the pane. You are returned to the Set up Single Sign-On with SAML page.
3. Follow the steps in [Use sandbox mode](#) to test each application.

Create service provider (SP) connections in PingFederate

You use PingFederate to create a service provider (SP) connection for each Admin Node in your system. To speed up the process, you will import the SAML metadata from StorageGRID.

Before you begin

- You have configured single sign-on for StorageGRID and you selected **Ping Federate** as the SSO type.
- **Sandbox mode** is selected on the Single sign-on page in Grid Manager. See [Use sandbox mode](#).
- You have the **SP connection ID** for each Admin Node in your system. You can find these values in the Admin Nodes detail table on the StorageGRID Single Sign-on page.
- You have downloaded the **SAML metadata** for each Admin Node in your system.
- You have experience creating SP connections in PingFederate Server.
- You have the [Administrator's Reference Guide](#) for PingFederate Server. The PingFederate documentation provides detailed step-by-step instructions and explanations.
- You have the [Admin permission](#) for PingFederate Server.

About this task

These instructions summarize how to configure PingFederate Server version 10.3 as an SSO provider for StorageGRID. If you are using another version of PingFederate, you might need to adapt these instructions. Refer to the PingFederate Server documentation for detailed instructions for your release.

Complete prerequisites in PingFederate

Before you can create the SP connections you will use for StorageGRID, you must complete prerequisite tasks in PingFederate. You will use information from these prerequisites when you configure the SP connections.

Create data store

If you haven't already, create a data store to connect PingFederate to the AD FS LDAP server. Use the values you used when [configuring identity federation](#) in StorageGRID.

- **Type:** Directory (LDAP)
- **LDAP Type:** Active Directory
- **Binary Attribute Name:** Enter **objectGUID** on the LDAP Binary Attributes tab exactly as shown.

Create password credential validator

If you haven't already, create a password credential validator.

- **Type:** LDAP Username Password Credential Validator
- **Data store:** Select the data store you created.
- **Search base:** Enter information from LDAP (for example, DC=saml,DC=sgws).
- **Search filter:** sAMAccountName=\${username}
- **Scope:** Subtree

Create IdP adapter instance

If you haven't already, create an IdP adapter instance.

Steps

1. Go to **Authentication > Integration > IdP Adapters**.
2. Select **Create New Instance**.
3. On the Type tab, select **HTML Form IdP Adapter**.
4. On the IdP Adapter tab, select **Add a new row to 'Credential Validators'**.
5. Select the [password credential validator](#) you created.
6. On the Adapter Attributes tab, select the **username** attribute for **Pseudonym**.
7. Select **Save**.

Create or import signing certificate

If you haven't already, create or import the signing certificate.

Steps

1. Go to **Security > Signing & Decryption Keys & Certificates**.

2. Create or import the signing certificate.

Create an SP connection in PingFederate

When you create an SP connection in PingFederate, you import the SAML metadata you downloaded from StorageGRID for the Admin Node. The metadata file contains many of the specific values you need.



You must create an SP connection for each Admin Node in your StorageGRID system, so that users can securely sign in to and out of any node. Use these instructions to create the first SP connection. Then, go to [Create additional SP connections](#) to create any additional connections you need.

Choose SP connection type

Steps

1. Go to **Applications > Integration > SP Connections**.
2. Select **Create Connection**.
3. Select **Do not use a template for this connection**.
4. Select **Browser SSO Profiles** and **SAML 2.0** as the protocol.

Import SP metadata

Steps

1. On the Import Metadata tab, select **File**.
2. Choose the SAML metadata file you downloaded from the StorageGRID Single sign-on page for the Admin Node.
3. Review the Metadata Summary and the information provided on the General Info tab.

The Partner's Entity ID and the Connection Name are set to the StorageGRID SP connection ID. (for example, 10.96.105.200-DC1-ADM1-105-200). The Base URL is the IP of the StorageGRID Admin Node.

4. Select **Next**.

Configure IdP Browser SSO

Steps

1. From the Browser SSO tab, select **Configure Browser SSO**.
2. On the SAML profiles tab, select the **SP-initiated SSO**, **SP-initial SLO**, **IdP-initiated SSO**, and **IdP-initiated SLO** options.
3. Select **Next**.
4. On the Assertion Lifetime tab, make no changes.
5. On the Assertion Creation tab, select **Configure Assertion Creation**.
 - a. On the Identity Mapping tab, select **Standard**.
 - b. On the Attribute Contract tab, use the **SAML_SUBJECT** as the Attribute Contract and the unspecified name format that was imported.
6. For Extend the Contract, select **Delete** to remove the `urn:oid`, which is not used.

Map adapter instance

Steps

1. On the Authentication Source Mapping tab, select **Map New Adapter Instance**.
2. On the Adapter instance tab, select the [adapter instance](#) you created.
3. On the Mapping Method tab, select **Retrieve Additional Attributes From a Data Store**.
4. On the Attribute Source & User Lookup tab, select **Add Attribute Source**.
5. On the Data Store tab, provide a description and select the [data store](#) you added.
6. On the LDAP Directory Search tab:
 - Enter the **Base DN**, which should exactly match the value you entered in StorageGRID for the LDAP server.
 - For the Search Scope, select **Subtree**.
 - For the Root Object Class, search for and add either of these attributes: **objectGUID** or **userPrincipalName**.
7. On the LDAP Binary Attribute Encoding Types tab, select **Base64** for the **objectGUID** attribute.
8. On the LDAP Filter tab, enter **sAMAccountName=\${username}**.
9. On the Attribute Contract Fulfillment tab, select **LDAP (attribute)** from the Source drop-down and select either **objectGUID** or **userPrincipalName** from the Value drop-down.
10. Review and then save the attribute source.
11. On the Failsave Attribute Source tab, select **Abort the SSO Transaction**.
12. Review the summary and select **Done**.
13. Select **Done**.

Configure protocol settings

Steps

1. On the **SP Connection > Browser SSO > Protocol Settings** tab, select **Configure Protocol Settings**.
2. On the Assertion Consumer Service URL tab, accept the default values, which were imported from the StorageGRID SAML metadata (**POST** for Binding and `/api/saml-response` for Endpoint URL).
3. On the SLO Service URLs tab, accept the default values, which were imported from the StorageGRID SAML metadata (**REDIRECT** for Binding and `/api/saml-logout` for Endpoint URL).
4. On the Allowable SAML Bindings tab, clear **ARTIFACT** and **SOAP**. Only **POST** and **REDIRECT** are required.
5. On the Signature Policy tab, leave the **Require Authn Requests to be Signed** and **Always Sign Assertion** checkboxes selected.
6. On the Encryption Policy tab, select **None**.
7. Review the summary and select **Done** to save the protocol settings.
8. Review the summary and select **Done** to save the Browser SSO settings.

Configure credentials

Steps

1. From the SP Connection tab, select **Credentials**.

2. From the Credentials tab, select **Configure Credentials**.
3. Select the [signing certificate](#) you created or imported.
4. Select **Next** to go to **Manage Signature Verification Settings**.
 - a. On the Trust Model tab, select **Unanchored**.
 - b. On the Signature Verification Certificate tab, review the signing certificate information, which was imported from the StorageGRID SAML metadata.
5. Review the summary screens and select **Save** to save the SP connection.

Create additional SP connections

You can copy the first SP connection to create the SP connections you need for each Admin Node in your grid. You upload new metadata for each copy.



The SP connections for different Admin Nodes use identical settings, with the exception of the Partner's Entity ID, Base URL, Connection ID, Connection Name, Signature Verification, and SLO Response URL.

Steps

1. Select **Action > Copy** to create a copy of the initial SP connection for each additional Admin Node.
2. Enter the Connection ID and Connection Name for the copy, and select **Save**.
3. Choose the metadata file corresponding to the Admin Node:
 - a. Select **Action > Update with Metadata**.
 - b. Select **Choose File** and upload the metadata.
 - c. Select **Next**.
 - d. Select **Save**.
4. Resolve the error due to the unused attribute:
 - a. Select the new connection.
 - b. Select **Configure Browser SSO > Configure Assertion Creation > Attribute Contract**.
 - c. Delete the entry for **urn:oid**.
 - d. Select **Save**.

Disable single sign-on

You can disable single sign-on (SSO) if you no longer want to use this functionality. You must disable single sign-on before you can disable identity federation.

Before you begin

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

Steps

1. Select **CONFIGURATION > Access control > Single sign-on**.

The Single Sign-on page appears.

2. Select the **Disabled** option.
3. Select **Save**.

A warning message appears indicating that local users will now be able to sign in.

4. Select **OK**.

The next time you sign in to StorageGRID, the StorageGRID Sign in page appears and you must enter the username and password for a local or federated StorageGRID user.

Temporarily disable and reenable single sign-on for one Admin Node

You might not be able to sign in to the Grid Manager if the single sign-on (SSO) system goes down. In this case, you can temporarily disable and reenable SSO for one Admin Node. To disable and then reenable SSO, you must access the node's command shell.

Before you begin

- You have [specific access permissions](#).
- You have the `Passwords.txt` file.
- You know the password for the local root user.

About this task

After you disable SSO for one Admin Node, you can sign in to the Grid Manager as the local root user. To secure your StorageGRID system, you must use the node's command shell to reenable SSO on the Admin Node as soon as you sign out.



Disabling SSO for one Admin Node does not affect the SSO settings for any other Admin Nodes in the grid. The **Enable SSO** checkbox on the Single Sign-on page in the Grid Manager remains selected, and all existing SSO settings are maintained unless you update them.

Steps

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

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

2. Run the following command: `disable-saml`

A message indicates that the command applies to this Admin Node only.

3. Confirm that you want to disable SSO.

A message indicates that single sign-on is disabled on the node.

4. From a web browser, access the Grid Manager on the same Admin Node.

The Grid Manager sign-in page is now displayed because SSO has been disabled.

5. Sign in with the username root and the local root user's password.
6. If you disabled SSO temporarily because you needed to correct the SSO configuration:
 - a. Select **CONFIGURATION > Access control > Single sign-on**.
 - b. Change the incorrect or out-of-date SSO settings.
 - c. Select **Save**.

Selecting **Save** from the Single Sign-on page automatically reenables SSO for the entire grid.

7. If you disabled SSO temporarily because you needed to access the Grid Manager for some other reason:
 - a. Perform whatever task or tasks you need to perform.
 - b. Select **Sign out**, and close the Grid Manager.
 - c. Reenable SSO on the Admin Node. You can perform either of the following steps:
 - Run the following command: `enable-saml`

A message indicates that the command applies to this Admin Node only.

Confirm that you want to enable SSO.

A message indicates that single sign-on is enabled on the node.

- Reboot the grid node: `reboot`

8. From a web browser, access the Grid Manager from the same Admin Node.
9. Confirm that the StorageGRID Sign in page appears and that you must enter your SSO credentials to access the Grid Manager.

Use grid federation

What is grid federation?


You can use grid federation to clone tenants and replicate their objects between two StorageGRID systems for disaster recovery.

What is a grid federation connection?

A grid federation connection is a bidirectional, trusted, and secure connection between Admin and Gateway Nodes in two StorageGRID systems.

Workflow for grid federation

The workflow diagram summarizes the steps for configuring a grid federation connection between two grids.



Considerations and requirements for grid federation connections

- The grids used for grid federation must be running StorageGRID versions that are either identical or have no more than one major version difference between them.

For details about version requirements, refer to the [Release notes](#).

- A grid can have one or more grid federation connections to other grids. Each grid federation connection is independent of any other connections. For example, if Grid 1 has one connection with Grid 2 and a second connection with Grid 3, there is no implied connection between Grid 2 and Grid 3.
- Grid federation connections are bidirectional. After the connection is established, you can monitor and manage the connection from either grid.
- At least one grid federation connection must exist before you can use [account clone](#) or [cross-grid replication](#).

Networking and IP address requirements

- Grid federation connections can occur on the Grid Network, Admin Network, or Client Network.
- A grid federation connection connects one grid to another grid. The configuration for each grid specifies a grid federation endpoint on the other grid that consists of Admin Nodes, Gateway Nodes, or both.
- The best practice is to connect [high availability \(HA\) groups](#) of Gateway and Admin Nodes on each grid. Using HA groups helps ensure that grid federation connections will remain online if nodes become unavailable. If the active interface in either HA group fails, the connection can use a backup interface.

- Creating a grid federation connection that uses the IP address of a single Admin Node or Gateway Node is not recommended. If the node becomes unavailable, the grid federation connection will also become unavailable.
- [Cross-grid replication](#) of objects requires that the Storage Nodes on each grid be able to access the configured Admin and Gateway Nodes on the other grid. For each grid, confirm that all Storage Nodes have a high bandwidth route to as the Admin Nodes or Gateway Nodes used for the connection.

Use FQDNs to load balance the connection

For a production environment, use fully qualified domain names (FQDNs) to identify each grid in the connection. Then, create the appropriate DNS entries, as follows:

- The FQDN for Grid 1 mapped to one or more virtual IP (VIP) addresses for HA groups in Grid 1 or to the IP address of one or more Admin or Gateway Nodes in Grid 1.
- The FQDN for Grid 2 mapped to one or more VIP addresses for Grid 2 or to the IP address of one or more Admin or Gateway Nodes in Grid 2.

When you use multiple DNS entries, requests to use the connection are load balanced, as follows:

- DNS entries that map to the VIP addresses of multiple HA groups are load balanced between the active nodes in the HA groups.
- DNS entries that map to the IP addresses of multiple Admin Nodes or Gateway Nodes are load balanced between the mapped nodes.

Port requirements


When creating a grid federation connection, you can specify any unused port number from 23000 to 23999. Both grids in this connection will use the same port.

You must ensure that no node in either grid uses this port for other connections.

Certificate requirements

When you configure a grid federation connection, StorageGRID automatically generates four SSL certificates:

- Server and client certificates to authenticate and encrypt information sent from grid 1 to grid 2
- Server and client certificates to authenticate and encrypt information sent from grid 2 to grid 1



By default, the certificates are valid for 730 days (2 years). When these certificates near their expiration date,

the **Expiration of grid federation certificate** alert reminds you to rotate the certificates, which you can do using the Grid Manager.



If the certificates on either end of the connection expire, the connection will stop working. Data replication will be pending until the certificates are updated.

Learn more

- [Create grid federation connections](#)
- [Manage grid federation connections](#)
- [Troubleshoot grid federation errors](#)


What is account clone?

Account clone is the automatic replication of a tenant account, tenant groups, tenant users, and, optionally, S3 access keys between the StorageGRID systems in a [grid federation connection](#).

Account clone is required for [cross-grid replication](#). Cloning account information from a source StorageGRID system to a destination StorageGRID system ensures that tenant users and groups can access the corresponding buckets and objects on either grid.

Workflow for account clone

The workflow diagram shows the steps that grid administrators and permitted tenants will perform to set up account clone. These steps are performed after the [grid federation connection is configured](#).



Grid admin workflow

The steps that grid admins perform depend on whether the StorageGRID systems in the [grid federation connection](#) use single sign-on (SSO) or identity federation.

Configure SSO for account clone (optional)

If either StorageGRID system in the grid federation connection uses SSO, both grids must use SSO. Before creating the tenant accounts for grid federation, the grid admins for the tenant's source and destination grids must perform these steps.

Steps

1. Configure the same identity source for both grids. See [Use identity federation](#).
2. Configure the same SSO identity provider (IdP) for both grids. See [Configure single sign-on](#).
3. [Create the same admin group](#) on both grids by importing the same federated group.

When you create the tenant, you will select this group to have the initial Root access permission for both the source and destination tenant accounts.



If this admin group doesn't exist on both grids before you create the tenant, the tenant isn't replicated to the destination.

Configure grid-level identity federation for account clone (optional)

If either StorageGRID system uses identity federation without SSO, both grids must use identity federation. Before creating the tenant accounts for grid federation, the grid admins for the tenant's source and destination grids must perform these steps.

Steps

1. Configure the same identity source for both grids. See [Use identity federation](#).
2. Optionally, if a federated group will have initial Root access permission for both the source and destination tenant accounts, [create the same admin group](#) on both grids by importing the same federated group.



If you assign Root access permission to a federated group that doesn't exist on both grids, the tenant isn't replicated to the destination grid.

3. If you don't want a federated group to have initial Root access permission for both accounts, specify a password for the local root user.

Create permitted S3 tenant account

After optionally configuring SSO or identity federation, a grid admin performs these steps to determine which tenants can replicate bucket objects to other StorageGRID systems.

Steps

1. Determine which grid you want to be the tenant's source grid for account clone operations.

The grid where the tenant is originally created is known as the tenant's *source grid*. The grid where the tenant is replicated is known as the tenant's *destination grid*.

2. On that grid, create a new S3 tenant account or edit an existing account.
3. Assign the **Use grid federation connection** permission.
4. If the tenant account will manage its own federated users, assign the **Use own identity source** permission.

If this permission is assigned, both the source and destination tenant accounts must configure the same identity source before creating federated groups. Federated groups added to the source tenant can't be cloned to the destination tenant unless both grids use the same identity source.

5. Select a specific grid federation connection.
6. Save the new or modified tenant.

When a new tenant with the **Use grid federation connection** permission is saved, StorageGRID automatically creates a replica of that tenant on the other grid, as follows:

- Both tenant accounts have the same account ID, name, storage quota, and assigned permissions.
- If you selected a federated group to have Root access permission for the tenant, that group is cloned to the destination tenant.
- If you selected a local user to have Root access permission for the tenant, that user is cloned to the destination tenant. However, the password for that user is not cloned.

For details, see [Manage permitted tenants for grid federation](#).

Permitted tenant account workflow

After a tenant with the **Use grid federation connection** permission is replicated to the destination grid, permitted tenant accounts can perform these steps to clone tenant groups, users, and S3 access keys.

Steps

1. Sign in to the tenant account on the tenant's source grid.
2. If permitted, configure identify federation on both the source and destination tenant accounts.
3. Create groups and users on the source tenant.

When new groups or users are created on the source tenant, StorageGRID automatically clones them to the destination tenant, but no cloning occurs from the destination back to the source.

4. Create S3 access keys.
5. Optionally, clone S3 access keys from the source tenant to the destination tenant.


For details about the permitted tenant account workflow and to learn how groups, users, and S3 access keys are cloned, see [Clone tenant groups and users](#) and [Clone S3 access keys using the API](#).

What is cross-grid replication?

Cross-grid replication is the automatic replication of objects between selected S3 buckets in two StorageGRID systems that are connected in a [grid federation connection](#). [Account clone](#) is required for cross-grid replication.

Workflow for cross-grid replication

The workflow diagram summarize the steps for configuring cross-grid replication between buckets on two grids.



Requirements for cross-grid replication

If a tenant account has the **Use grid federation connection** permission to use one or more [grid federation connections](#), a tenant user with Root access permission can create identical buckets in the corresponding tenant accounts on each grid. These buckets:

- Must have the same name but can have different regions
- Must have versioning enabled
- Must have S3 Object Lock disabled
- Must be empty

After both buckets have been created, cross-grid replication can be configured for either or both buckets.

Learn more


[Manage cross-grid replication](#)

How cross-grid replication works

Cross-grid replication can be configured to occur in one direction or in both directions.


Replication in one direction

If you enable cross-grid replication for a bucket on only one grid, objects added to that bucket (the source bucket) are replicated to the corresponding bucket on the other grid (the destination bucket). However, objects added to the destination bucket aren't replicated back to the source. In the figure, cross-grid replication is enabled for `my-bucket` from Grid 1 to Grid 2, but it is not enabled in the other direction.



Replication in both directions

If you enable cross-grid replication for the same bucket on both grids, objects added to either bucket are replicated to the other grid. In the figure, cross-grid replication is enabled for `my-bucket` in both directions.



What happens when objects are ingested?

When an S3 client adds an object to a bucket that has cross-grid replication enabled, the following happens:

1. StorageGRID automatically replicates the object from the source bucket to the destination bucket. The time to perform this background replication operation depends on several factors, including the number of other replication operations that are pending.

The S3 client can verify an object's replication status by issuing a `GetObject` or `HeadObject` request. The response includes a StorageGRID-specific `x-ntap-sg-cgr-replication-status` response header, which will have one of the following values:

The S3 client can verify an object's replication status by issuing a `GetObject` or `HeadObject` request. The response includes a StorageGRID-specific `x-ntap-sg-cgr-replication-status` response header, which will have one of the following values:

Grid	Replication status
Source	• COMPLETED: The replication was successful for all grid connections. • PENDING: The object hasn't been replicated to at least one grid connection. • FAILURE: Replication is not pending for any grid connection and at least one failed with a permanent failure. A user must resolve the error.
Destination	REPLICA: The object was replicated from the source grid.



StorageGRID does not support the `x-amz-replication-status` header.

2. StorageGRID uses each grid's active ILM policies to manage the objects, just as it would any other object. For example, Object A on Grid 1 might be stored as two replicated copies and retained forever, while the copy of Object A that was replicated to Grid 2 might be stored using 2+1 erasure coding and deleted after three years.

What happens when objects are deleted?

As described in [Delete data flow](#), StorageGRID can delete an object for any of these reasons:

- The S3 client issues a delete request.
- A Tenant Manager user selects the [Delete objects in bucket](#) option to remove all objects from a bucket.
- The bucket has a lifecycle configuration, which expires.
- The last time period in the ILM rule for the object ends, and there are no further placements specified.

When StorageGRID deletes an object because of a Delete objects in bucket operation, bucket lifecycle expiration, or ILM placement expiration, the replicated object is never deleted from the other grid in a grid federation connection. However, delete markers added to the source bucket by S3 client deletes can optionally be replicated to the destination bucket.


To understand what happens when an S3 client deletes objects from a bucket that has cross-grid replication enabled, review how S3 clients delete objects from buckets that have versioning enabled, as follows:

- If an S3 client issues a delete request that includes a version ID, that version of the object is permanently removed. No delete marker is added to the bucket.
- If an S3 client issues a delete request that does not include a version ID, StorageGRID does not delete any object versions. Instead, it adds a delete marker to the bucket. The delete marker causes StorageGRID to act as if the object was deleted:
 - A GetObject request without a version ID will fail with 404 No Object Found
 - A GetObject request with a valid version ID will succeed and return the requested object version.

When an S3 client deletes an object from a bucket that has cross-grid replication enabled, StorageGRID determines whether to replicate the delete request to the destination, as follows:

- If the delete request includes a version ID, that object version is permanently removed from the source grid. However, StorageGRID does not replicate delete requests that include a version ID, so the same object version is not deleted from the destination.
- If the delete request does not include a version ID, StorageGRID can optionally replicate the delete marker, based on how cross-grid replication is configured for the bucket:
 - If you choose to replicate delete markers (default), a delete marker is added to the source bucket and replicated to the destination bucket. In effect, the object appears to be deleted on both grids.
 - If you choose not to replicate delete markers, a delete marker is added to the source bucket but is not replicated to the destination bucket. In effect, objects that are deleted on the source grid aren't deleted on the destination grid.

In the figure, **Replicate delete markers** was set to **Yes** when **cross-grid replication was enabled**. Delete requests for the source bucket that include a version ID will not delete objects from the destination bucket. Delete requests for the source bucket that don't include a version ID will appear to delete objects in the destination bucket.



If you want to keep object deletions synchronized between grids, create corresponding [S3 lifecycle configurations](#) for the buckets on both grids.

How encrypted objects are replicated

When you use cross-grid replication to replicate objects between grids, you can encrypt individual objects, use default bucket encryption, or configure grid-wide encryption. You can add, modify, or remove default bucket or

grid-wide encryption settings before or after you enable cross-grid replication for a bucket.

To encrypt individual objects, you can use SSE (server-side encryption with StorageGRID-managed keys) when adding the objects to the source bucket. Use the `x-amz-server-side-encryption` request header and specify `AES256`. See [Use server-side encryption](#).



Using SSE-C (server-side encryption with customer-provided keys) is not supported for cross-grid replication. The ingest operation will fail.

To use default encryption for a bucket, use a `PutBucketEncryption` request and set the `SSEAlgorithm` parameter to `AES256`. Bucket-level encryption applies to any objects ingested without the `x-amz-server-side-encryption` request header. See [Operations on buckets](#).

To use grid-level encryption, set the **Stored object encryption** option to **AES-256**. Grid-level encryption applies to any objects that aren't encrypted at the bucket level or that are ingested without the `x-amz-server-side-encryption` request header. See [Configure network and object options](#).



SSE does not support AES-128. If the **Stored object encryption** option is enabled for the source grid using the **AES-128** option, the use of the AES-128 algorithm will not be propagated to the replicated object. Instead, the replicated object will use the destination's default bucket or grid-level encryption setting, if available.

When determining how to encrypt source objects, StorageGRID applies these rules:

1. Use the `x-amz-server-side-encryption` ingest header, if present.
2. If an ingest header is not present, use the bucket default encryption setting, if configured.
3. If a bucket setting is not configured, use the grid-wide encryption setting, if configured.
4. If a grid-wide setting is not present, don't encrypt the source object.

When determining how to encrypt replicated objects, StorageGRID applies these rules in this order:

1. Use the same encryption as the source object, unless that object uses AES-128 encryption.
2. If the source object is not encrypted or it uses AES-128, use the destination bucket's default encryption setting, if configured.
3. If the destination bucket does not have an encryption setting, use the destination's grid-wide encryption setting, if configured.
4. If a grid-wide setting is not present, don't encrypt the destination object.

PutObjectTagging and DeleteObjectTagging aren't supported

`PutObjectTagging` and `DeleteObjectTagging` requests aren't supported for objects in buckets that have cross-grid replication enabled.

If an S3 client issues a `PutObjectTagging` or `DeleteObjectTagging` request, `501 Not Implemented` is returned. The message is `Put (Delete) ObjectTagging is not available for buckets that have cross-grid replication configured`.

How segmented objects are replicated

The source grid's maximum segment size applies to objects replicated to the destination grid. When objects

are replicated to another grid, the **Maximum Segment Size** setting (**CONFIGURATION > System > Storage options**) of the source grid will be used on both grids. For example, suppose the maximum segment size for the source grid is 1 GB, while the maximum segment size of the destination grid is 50 MB. If you ingest a 2-GB object on the source grid, that object is saved as two 1-GB segments. It will also be replicated to the destination grid as two 1-GB segments, even though that grid's maximum segment size is 50 MB.

Compare cross-grid replication and CloudMirror replication

As you begin using grid federation, review the similarities and differences between [cross-grid replication](#) and the [StorageGRID CloudMirror replication service](#).

	Cross-grid replication	CloudMirror replication service
What is the primary purpose?	One StorageGRID system acts as a disaster recovery system. Objects in a bucket can be replicated between the grids in one or both directions.	Enables a tenant to automatically replicate objects from a bucket in StorageGRID (source) to an external S3 bucket (destination). CloudMirror replication creates an independent copy of an object in an independent S3 infrastructure. This independent copy is not used as a backup, but often further processed in the cloud.
How is it set up?	1. Configure a grid federation connection between two grids. 2. Add new tenant accounts, which are automatically cloned to the other grid. 3. Add new tenant groups and users, which are also cloned. 4. Create corresponding buckets on each grid and enable cross-grid replication to occur in one or both directions.	1. A tenant user configures CloudMirror replication by defining a CloudMirror endpoint (IP address, credentials, and so on) using the Tenant Manager or the S3 API. 2. Any bucket owned by that tenant account can be configured to point to the CloudMirror endpoint.
Who is responsible for setting it up?	• A grid admin configures the connection and the tenants. • Tenant users configure the groups, users, keys, and buckets.	Typically, a tenant user.
What is the destination?	A corresponding and identical S3 bucket on the other StorageGRID system in the grid federation connection.	• Any compatible S3 infrastructure (including Amazon S3). • Google Cloud Platform (GCP)
Is object versioning required?	Yes, both the source and destination buckets must have object versioning enabled.	No, CloudMirror replication supports any combination of unversioned and versioned buckets on both the source and destination.


	Cross-grid replication	CloudMirror replication service
What causes objects to be moved to the destination?	Objects are automatically replicated when they are added to a bucket that has cross-grid replication enabled.	Objects are automatically replicated when they are added to a bucket that has been configured with a CloudMirror endpoint. Objects that existed in the source bucket before the bucket was configured with the CloudMirror endpoint aren't replicated, unless they are modified.
How are objects replicated?	Cross-grid replication creates versioned objects, and it replicates the version ID from the source bucket to the destination bucket. This allows the version order to be maintained across both grids.	CloudMirror replication doesn't require versioning-enabled buckets, so CloudMirror can only maintain ordering for a key within a site. There are no guarantees that ordering will be maintained for requests to an object at different site.
What if an object can't be replicated?	The object is queued for replication, subject to metadata storage limits.	The object is queued for replication, subject to platform services limits (see Recommendations for using platform services).
Is the object's system metadata replicated?	Yes, when an object is replicated to the other grid, its system metadata is also replicated. The metadata will be identical on both grids.	No, when an object is replicated to the external bucket, its system metadata is updated. The metadata will differ between locations, depending on time of ingest and the behavior of the independent S3 infrastructure.
How are objects retrieved?	Applications can retrieve or read objects by making a request to the bucket on either grid.	Applications can retrieve or read objects by making a request either to StorageGRID or to the S3 destination. For example, suppose you use CloudMirror replication to mirror objects to a partner organization. The partner can use its own applications to read or update objects directly from the S3 destination. Using StorageGRID is not required.

	Cross-grid replication	CloudMirror replication service
What happens if an object is deleted?	• Delete requests that include a version ID are never replicated to the destination grid. • Delete requests that don't include a version ID add a delete marker to the source bucket, which can optionally be replicated to the destination grid. • If cross-grid replication is configured for only one direction, objects in the destination bucket can be deleted without affecting the source.	The results will vary based on the versioning state of the source and destination buckets (which don't need to be the same): • If both buckets are versioned, a delete request will add a delete marker in both locations. • If only the source bucket is versioned, a delete request will add a delete marker to the source but not to the destination. • If neither bucket is versioned, a delete request will delete the object from the source but not from the destination. Similarly, objects in the destination bucket can be deleted without affecting the source.

Create grid federation connections

You can create a grid federation connection between two StorageGRID systems if you want to clone tenant details and replicate object data.

As shown in the figure, creating a grid federation connection includes steps on both grids. You add the connection on one grid and complete it on the other grid. You can start from either grid.



Before you begin

- You have reviewed the [considerations and requirements](#) for configuring grid federation connections.
- If you plan to use fully qualified domain names (FQDNs) for each grid instead of IP or VIP addresses, you know which names to use and you have confirmed that the DNS server for each grid has the appropriate entries.
- You are using a [supported web browser](#).
- You have Root access permission and the provisioning passphrase for both grids.

Add connection

Perform these steps on either of the two StorageGRID systems.

Steps

1. Sign in to the Grid Manager from the primary Admin Node on either grid.
2. Select **CONFIGURATION > System > Grid federation**.
3. Select **Add connection**.
4. Enter details for the connection.

Field	Description
Connection name	A unique name to help you recognize this connection, for example, "Grid 1-Grid 2."
FQDN or IP for this grid	One of the following: • The FQDN of the grid you are currently signed into • A VIP address of an HA group on this grid • An IP address of an Admin Node or Gateway Node on this grid. The IP can be on any network that the destination grid can reach.
Port	The port you want to use for this connection. You can enter any unused port number from 23000 to 23999. Both grids in this connection will use the same port. You must ensure that no node in either grid uses this port for other connections.
Certificate valid days for this grid	The number of days you want the security certificates for this grid in the connection to be valid. The default value is 730 days (2 years), but you can enter any value from 1 to 762 days. StorageGRID automatically generates client and server certificates for each grid when you save the connection.
Provisioning passphrase for this grid	The provisioning passphrase for the grid you are signed in to.
FQDN or IP for the other grid	One of the following: • The FQDN of the grid you want to connect to • A VIP address of an HA group on the other grid • An IP address of an Admin Node or Gateway Node on the other grid. The IP can be on any network that the source grid can reach.

5. Select **Save and continue**.
6. For the Download verification file step, select **Download verification file**.

After the connection is completed on the other grid, you can no longer download the verification file from either grid.

7. Locate the downloaded file (*connection-name.grid-federation*), and save it to a safe location.



This file contains secrets (masked as *****) and other sensitive details and must be securely stored and transmitted.

8. Select **Close** to return to the Grid federation page.
9. Confirm that the new connection is shown and that its **Connection status** is **Waiting to connect**.
10. Provide the *connection-name.grid-federation* file to the grid admin for the other grid.

Complete connection

Perform these steps on the StorageGRID system you are connecting to (the other grid).

Steps

1. Sign in to the Grid Manager from the primary Admin Node.
2. Select **CONFIGURATION > System > Grid federation**.
3. Select **Upload verification file** to access the Upload page.
4. Select **Upload verification file**. Then, browse to and select the file that was downloaded from the first grid (*connection-name.grid-federation*).

The details for the connection are shown.

5. Optionally, enter a different number of valid days for the security certificates for this grid. The **Certificate valid days** entry defaults to the value you entered on the first grid, but each grid can use different expiration dates.

In general, use the same number of days for the certificates on both sides of the connection.



If the certificates on either end of the connection expire, the connection will stop working and replications will be pending until the certificates are updated.

6. Enter the provisioning passphrase for the grid you are currently signed in to.
7. Select **Save and test**.

The certificates are generated and the connection is tested. If the connection is valid, a success message appears and the new connection is listed on the Grid federation page. The **Connection status** will be **Connected**.

If an error message appears, address any issues. See [Troubleshoot grid federation errors](#).

8. Go to the Grid federation page on the first grid and refresh the browser. Confirm that the **Connection status** is now **Connected**.
9. After the connection has been established, securely delete all copies of the verification file.

If you edit this connection, a new verification file will be created. The original file can't be reused.

After you finish

- Review the considerations for [managing permitted tenants](#).
- [Create one or more new tenant accounts](#), assign the **Use grid federation connection** permission, and select the new connection.
- [Manage the connection](#) as required. You can edit connection values, test a connection, rotate connection certificates, or remove a connection.
- [Monitor the connection](#) as part of your normal StorageGRID monitoring activities.
- [Troubleshoot the connection](#), including resolving any alerts and errors related to account clone and cross-grid replication.

Manage grid federation connections

Managing grid federation connections between StorageGRID systems includes editing connection details, rotating the certificates, removing tenant permissions, and removing unused connections.

Before you begin

- You are signed in to the Grid Manager on either grid using a [supported web browser](#).
- You have the [Root access permission](#) for the grid you are signed in to.

Edit a grid federation connection

You can edit a grid federation connection by signing in to the primary Admin Node on either grid in the connection. After you make changes to the first grid, you must download a new verification file and upload it to the other grid.



While the connection is being edited, account clone or cross-grid replication requests will continue to use the existing connection settings. Any edits you make to the first grid are saved locally but aren't used until they have been uploaded to the second grid, saved, and tested.

Start editing the connection

Steps

1. Sign in to the Grid Manager from the primary Admin Node on either grid.
2. Select **NODES** and confirm that all other Admin Nodes in your system are online.



When you edit a grid federation connection, StorageGRID attempts to save a "candidate configuration" file on all Admin Nodes on the first grid. If this file can't be saved to all Admin Nodes, a warning message appears when you select **Save and test**.

3. Select **CONFIGURATION > System > Grid federation**.
4. Edit the connection details using the **Actions** menu on the Grid federation page or the details page for a specific connection. See [Create grid federation connections](#) for what to enter.

Actions menu

- a. Select the radio button for the connection.
- b. Select **Actions > Edit**.
- c. Enter the new information.

Details page

- a. Select a connection name to display its details.
- b. Select **Edit**.
- c. Enter the new information.

5. Enter the provisioning passphrase for the grid you are signed in to.
6. Select **Save and continue**.

The new values are saved, but they will not be applied to the connection until you have uploaded the new verification file on the other grid.

7. Select **Download verification file**.

To download this file at a later time, go to the details page for the connection.

8. Locate the downloaded file (*connection-name.grid-federation*), and save it to a safe location.



The verification file contains secrets and must be securely stored and transmitted.

9. Select **Close** to return to the Grid federation page.
10. Confirm that the **Connection status** is **Pending edit**.



If the connection status was something other than **Connected** when you started editing the connection, it will not change to **Pending edit**.

11. Provide the *connection-name.grid-federation* file to the grid admin for the other grid.

Finish editing the connection

Finish editing the connection by uploading the verification file on the other grid.

Steps

1. Sign in to the Grid Manager from the primary Admin Node.
2. Select **CONFIGURATION > System > Grid federation**.
3. Select **Upload verification file** to access the upload page.
4. Select **Upload verification file**. Then, browse to and select the file that was downloaded from the first grid.
5. Enter the provisioning passphrase for the grid you are currently signed in to.
6. Select **Save and test**.

If the connection can be established using the edited values, a success message appears. Otherwise, an error message appears. Review the message and address any issues.

7. Close the wizard to return to the Grid federation page.
8. Confirm that the **Connection status** is **Connected**.
9. Go to the Grid federation page on the first grid and refresh the browser. Confirm that the **Connection status** is now **Connected**.
10. After the connection has been established, securely delete all copies of the verification file.

Test a grid federation connection

Steps

1. Sign in to the Grid Manager from the primary Admin Node.
2. Select **CONFIGURATION > System > Grid federation**.
3. Test the connection using the **Actions** menu on the Grid federation page or the details page for a specific connection.

Actions menu

- a. Select the radio button for the connection.
- b. Select **Actions > Test**.

Details page

- a. Select a connection name to display its details.
- b. Select **Test connection**.

4. Review the connection status:

Connection status	Description
Connected	Both grids are connected and communicating normally.
Error	The connection is in an error state. For example, a certificate has expired or a configuration value is no longer valid.
Pending edit	You have edited the connection on this grid, but the connection is still using the existing configuration. To complete the edit, upload the new verification file to the other grid.
Waiting to connect	You have configured the connection on this grid, but the connection hasn't been completed on the other grid. Download the verification file from this grid and upload it to the other grid.
Unknown	The connection is in an unknown state, possibly because of a networking issue or an offline node.

5. If the Connection status is **Error**, resolve any issues. Then, select **Test connection** again to confirm the issue has been fixed.

Rotate connection certificates

Each grid federation connection uses four automatically-generated SSL certificates to secure the connection. When the two certificates for each grid near their expiration date, the **Expiration of grid federation certificate** alert reminds you to rotate the certificates.



If the certificates on either end of the connection expire, the connection will stop working and replications will be pending until the certificates are updated.


Steps

1. Sign in to the Grid Manager from the primary Admin Node on either grid.
2. Select **CONFIGURATION > System > Grid federation**.
3. From either tab on the Grid federation page, select the connection name to display its details.
4. Select the **Certificates** tab.
5. Select **Rotate certificates**.
6. Specify how many days the new certificates should be valid.
7. Enter the provisioning passphrase for the grid you are signed in to.
8. Select **Rotate certificates**.
9. As required, repeat these steps on the other grid in the connection.

In general, use the same number of days for the certificates on both sides of the connection.

Remove a grid federation connection

You can remove a grid federation connection from either grid in the connection. As shown in the figure, you must perform prerequisite steps on both grids to confirm that the connection is not being used by any tenant on either grid.



Before removing a connection, note the following:

- Removing a connection does not delete any items that have already been copied between grids. For example, tenant users, groups, and objects that exist on both grids aren't deleted from either grid when the tenant's permission is removed. If you want to delete these items, you must manually delete them from both grids.
- When you remove a connection, any objects that are pending replication (ingested but not yet replicated to the other grid) will have their replication permanently failed.

Disable replication for all tenant buckets

Steps

1. Starting from either grid, sign in to the Grid Manager from the primary Admin Node.
2. Select **CONFIGURATION > System > Grid federation**.
3. Select the connection name to display its details.

4. On the **Permitted tenants** tab, determine if the connection is being used by any tenants.
5. If any tenants are listed, instruct all tenants to [disable cross-grid replication](#) for all of their buckets on both grids in the connection.



You can't remove the **Use grid federation connection** permission if any tenant buckets have cross-grid replication enabled. Each tenant account must disable cross-grid replication for their buckets on both grids.

Remove permission for each tenant

After cross-grid replication has been disabled for all tenant buckets, remove the **Use grid federation permission** from all tenants on both grids.

Steps

1. Select **CONFIGURATION > System > Grid federation**.
2. Select the connection name to display its details.
3. For each tenant on the **Permitted tenants** tab, remove the **Use grid federation connection** permission from each tenant. See [Manage permitted tenants](#).
4. Repeat these steps for the permitted tenants on the other grid.

Remove connection

Steps

1. When no tenants on either grid are using the connection, select **Remove**.
2. Review the confirmation message, and select **Remove**.
 - If the connection can be removed, a success message is shown. The grid federation connection is now removed from both grids.
 - If the connection can't be removed (for example, it is still in use or there is a connection error), an error message is displayed. You can do either of the following:
 - Resolve the error (recommended). See [Troubleshoot grid federation errors](#).
 - Remove the connection by force. See the next section.

Remove a grid federation connection by force

If necessary, you can force the removal of a connection that does not have **Connected** status.

Force removal only deletes the connection from the local grid. To completely remove the connection, perform the same steps on both grids.

Steps

1. From the confirmation dialog box, select **Force remove**.

A success message appears. This grid federation connection can no longer be used. However, tenant buckets might still have cross-grid replication enabled and some object copies might have already been replicated between the grids in the connection.

2. From the other grid in the connection, sign in to the Grid Manager from the primary Admin Node.
3. Select **CONFIGURATION > System > Grid federation**.

4. Select the connection name to display its details.
5. Select **Remove** and **Yes**.
6. Select **Force remove** to remove the connection from this grid.

Manage the permitted tenants for grid federation

You can allow S3 tenant accounts to use a grid federation connection between two StorageGRID systems. When tenants are allowed to use a connection, special steps are required to edit tenant details or to permanently remove a tenant's permission to use the connection.

Before you begin

- You are signed in to the Grid Manager on either grid using a [supported web browser](#).
- You have the [Root access permission](#) for the grid you are signed in to.
- You have [created a grid federation connection](#) between two grids.
- You have reviewed the workflows for [account clone](#) and [cross-grid replication](#).
- As required, you have already configured single sign-on (SSO) or identify federation for both grids in the connection. See [What is account clone](#).

Create a permitted tenant

If you want to allow a new or existing tenant account to use a grid federation connection for account clone and cross-grid replication, follow the general instructions to [create a new S3 tenant](#) or [edit a tenant account](#) and note the following:

- You can create the tenant from either grid in the connection. The grid where a tenant is created is the *tenant's source grid*.
- The status of the connection must be **Connected**.
- When the tenant is created or edited to enable the **Use grid federation connection** permission and then saved on the first grid, an identical tenant is automatically replicated to the other grid. The grid where the tenant is replicated is the *tenant's destination grid*.
- The tenants on both grids will have the same 20-digit account ID, name, description, quota, and permissions. Optionally, you can use the **Description** field to help identify which is the source tenant and which is the destination tenant. For example, this description for a tenant created on Grid 1 will also appear for the tenant replicated to Grid 2: "This tenant was created on Grid 1."
- For security reasons, the password for a local root user is not copied to the destination grid.



Before a local root user can sign in to the replicated tenant on the destination grid, a grid administrator for that grid must [change the password for the local root user](#).

- After the new or edited tenant is available on both grids, tenant users can perform these operations:
 - From the tenant's source grid, create groups and local users, which are automatically cloned to the tenant's destination grid. See [Clone tenant groups and users](#).
 - Create new S3 access keys, which can be optionally cloned to the tenant's destination grid. See [Clone S3 access keys using the API](#).
 - Create identical buckets on both grids in the connection and enable cross-grid replication in one direction or in both directions. See [Manage cross-grid replication](#).

View a permitted tenant

You can see details for a tenant that is permitted to use a grid federation connection.

Steps

1. Select **TENANTS**.
2. From the Tenants page, select the tenant name to view the tenant details page.

If this is the source grid for the tenant (that is, if the tenant was created on this grid), a banner appears to remind you that the tenant was cloned to another grid. If you edit or delete this tenant, your changes will not be synced to the other grid.

Tenants > tenant A for grid federation

tenant A for grid federation

Tenant ID: 0899 6970 1700 0930 0009

Protocol: S3

Object count: 0

Quota utilization: —

Logical space used: 0 bytes

Quota: —

Description: this tenant was created on Grid 1

[Sign in](#) [Edit](#) [Actions](#) ▾

This tenant has been cloned to another grid. If you edit or delete this tenant, your changes will not be synced to the other grid.

[Space breakdown](#) [Allowed features](#) [Grid federation](#)

[Remove permission](#) [Clear error](#) Displaying one result

Connection name	Connection status	Remote grid hostname	Last error
Grid 1 to Grid 2	Connected	10.96.106.230	Check for errors

3. Optionally select the **Grid federation** tab to [monitor the grid federation connection](#).

Edit a permitted tenant

If you need to edit a tenant that has the **Use grid federation connection** permission, follow the general instructions for [editing a tenant account](#) and note the following:

- If a tenant has the **Use grid federation connection** permission, you can edit tenant details from either grid in the connection. However, any changes you make will not be copied to the other grid. If you want to keep the tenant details synchronized between grids, you must make the same edits on both grids.

- You can't clear the **Use grid federation connection** permission when you are editing a tenant.
- You can't select a different grid federation connection when you are editing a tenant.

Delete a permitted tenant


If you need to remove a tenant that has the **Use grid federation connection** permission, follow the general instructions for [deleting a tenant account](#) and note the following:

- Before you can remove the original tenant on the source grid, you must remove all buckets for the account on the source grid.
- Before you can remove the cloned tenant on the destination grid, you must remove all buckets for the account on the destination grid.
- If you remove either the original or the cloned tenant, the account can no longer be used for cross-grid replication.
- If you are removing the original tenant on the source grid, any tenant groups, users, or keys that were cloned to the destination grid will be unaffected. You can either delete the cloned tenant or allow it to manage its own groups, users, access keys, and buckets.
- If you are removing the cloned tenant on the destination grid, clone errors will occur if new groups or users are added to the original tenant.

To avoid these errors, remove the tenant's permission to use the grid federation connection before deleting the tenant from this grid.

Remove Use grid federation connection permission

To prevent a tenant from using a grid federation connection, you must remove the **Use grid federation connection** permission.



Before removing a tenant's permission to use a grid federation connection, note the following:

- You can't remove the **Use grid federation connection** permission if any of the tenant's buckets have cross-grid replication enabled. The tenant account must disable cross-grid replication for all of their buckets first.
- Removing the **Use grid federation connection** permission does not delete any items that have already been replicated between grids. For example, any tenant users, groups, and objects that exist on both grids aren't deleted from either grid when the tenant's permission is removed. If you want to delete these items, you must manually delete them from both grids.
- If you want to re-enable this permission with the same grid federation connection, delete this tenant on the destination grid first; otherwise, re-enabling this permission will result in an error.



Re-enabling the **Use grid federation connection** permission makes the local grid the source grid and triggers cloning to the remote grid specified by the selected grid federation connection. If the tenant account already exists on the remote grid, cloning will result in a conflict error.

Before you begin

- You are using a [supported web browser](#).
- You have the [Root access permission](#) for both grids.

Disable replication for tenant buckets

As a first step, disable cross-grid replication for all tenant buckets.

Steps

1. Starting from either grid, sign in to the Grid Manager from the primary Admin Node.
2. Select **CONFIGURATION > System > Grid federation**.
3. Select the connection name to display its details.
4. On the **Permitted tenants** tab, determine if the tenant is using the connection.
5. If the tenant is listed, instruct them to [disable cross-grid replication](#) for all of their buckets on both grids in the connection.



You can't remove the **Use grid federation connection** permission if any tenant buckets have cross-grid replication enabled. The tenant must disable cross-grid replication for their buckets on both grids.

Remove permission for tenant

After cross-grid replication is disabled for tenant buckets, you can remove the tenant's permission to use the grid federation connection.

Steps

1. Sign in to the Grid Manager from the primary Admin Node.
2. Remove the permission from the Grid federation page or the Tenants page.

Grid federation page


- a. Select **CONFIGURATION > System > Grid federation**.
- b. Select the connection name to display its details page.
- c. On the **Permitted tenants** tab, select radio button for the tenant.
- d. Select **Remove permission**.

Tenants page

- a. Select **TENANTS**.
- b. Select the tenant's name to display the details page.
- c. On the **Grid federation** tab, select radio button for the connection.
- d. Select **Remove permission**.

3. Review the warnings in the confirmation dialog box, and select **Remove**.
 - If the permission can be removed, you are returned to the details page and a success message is shown. This tenant can no longer use the grid federation connection.

- If one or more tenant buckets still have cross-grid replication enabled, an error is displayed.



You can do either of the following:

- (Recommended.) Sign in to the Tenant Manager and disable replication for each of the tenant's buckets. See [Manage cross-grid replication](#). Then, repeat the steps to remove the **Use grid connection** permission.
- Remove the permission by force. See the next section.

4. Go to the other grid and repeat these steps to remove the permission for the same tenant on the other grid.

Remove the permission by force

If necessary, you can force the removal of a tenant's permission to use a grid federation connection even if tenant buckets have cross-grid replication enabled.

Before removing a tenant's permission by force, note the general considerations for [removing the permission](#) as well as these additional considerations:

- If you remove the **Use grid federation connection** permission by force, any objects that are pending replication to the other grid (ingested but not yet replicated) will continue to be replicated. To prevent these

in-process objects from reaching the destination bucket, you must remove the tenant's permission on the other grid as well.

- Any objects ingested into the source bucket after you remove the **Use grid federation connection** permission will never be replicated to the destination bucket.

Steps

1. Sign in to the Grid Manager from the primary Admin Node.
2. Select **CONFIGURATION > System > Grid federation**.
3. Select the connection name to display its details page.
4. On the **Permitted tenants** tab, select radio button for the tenant.
5. Select **Remove permission**.
6. Review the warnings in the confirmation dialog box, and select **Force remove**.

A success message appears. This tenant can no longer use the grid federation connection.

7. As required, go to the other grid and repeat these steps to force-remove the permission for the same tenant account on the other grid. For example, you should repeat these steps on the other grid to prevent in-process objects from reaching the destination bucket.

Troubleshoot grid federation errors

You might need to troubleshoot alerts and errors related to grid federation connections, account clone, and cross-grid replication.

Grid federation connection alerts and errors

You might receive alerts or experience errors with your grid federation connections.

After making any changes to resolve a connection issue, test the connection to ensure that the connection status returns to **Connected**. For instructions, see [Manage grid federation connections](#).

Grid federation connection failure alert

Issue

The **Grid federation connection failure** alert was triggered.

Details

This alert indicates that the grid federation connection between the grids is not working.

Recommended actions

1. Review the settings on the Grid Federation page for both grids. Confirm that all values are correct. See [Manage grid federation connections](#).
2. Review the certificates used for the connection. Make sure there are no alerts for expired grid federation certificates and that the details for each certificate are valid. See the instructions for rotating connection certificates in [Manage grid federation connections](#).
3. Confirm that all Admin and Gateway Nodes in both grids are online and available. Resolve any alerts that might be affecting these nodes and try again.
4. If you provided a fully qualified domain name (FQDN) for the local or remote grid, confirm the DNS server is online and available. See [What is grid federation?](#) for networking, IP address, and DNS requirements.

Expiration of grid federation certificate alert

Issue

The **Expiration of grid federation certificate** alert was triggered.

Details

This alert indicates that one or more grid federation certificates are about to expire.

Recommended actions

See the instructions for rotating connection certificates in [Manage grid federation connections](#).

Error editing a grid federation connection

Issue

When editing a grid federation connection, you see the following warning message when you select **Save and test**: "Failed to create a candidate configuration file on one or more nodes."

Details

When you edit a grid federation connection, StorageGRID attempts to save a "candidate configuration" file on all Admin Nodes on the first grid. A warning message appears if this file can't be saved to all Admin Nodes, for example, because an Admin Node is offline.

Recommended actions

1. From the grid you are using to edit the connection, select **NODES**.
2. Confirm that all Admin Nodes for that grid are online.
3. If any nodes are offline, bring them back online and try editing the connection again.

Account clone errors

Can't sign in to a cloned tenant account

Issue

You can't sign in to a cloned tenant account. The error message on the Tenant Manager sign-in page is "Your credentials for this account were invalid. Please try again."

Details

For security reasons, when a tenant account is cloned from the tenant's source grid to the tenant's destination grid, the password you set for the tenant's local root user is not cloned. Similarly, when a tenant creates local users on its source grid, the local user passwords aren't cloned to the destination grid.

Recommended actions

Before the root user can sign in to the tenant's destination grid, a grid administrator must first [change the password for the local root user](#) on the destination grid.

Before a cloned local user can sign in to the tenant's destination grid, the root user for the cloned tenant must add a password for the user on the destination grid. For instructions, see [Manage local users](#) in the instructions for using the Tenant Manager.

Tenant created without a clone

Issue

You see the message "Tenant created without a clone" after creating a new tenant with the **Use grid**


federation connection permission.

Details

This issue can occur if updates to the Connection status are delayed, which might cause an unhealthy connection to be listed as **Connected**.

Recommended actions

- 1. Review the reason listed in the error message and resolve any networking or other issues that might be preventing the connection from working. See [Grid federation connection alerts and errors](#).
- 2. Follow the instructions to test a grid federation connection in [Manage grid federation connections](#) to confirm the issue has been fixed.
- 3. From the tenant’s source grid, select **TENANTS**.
- 4. Locate the tenant account that failed to be cloned.
- 5. Select the tenant name to display the details page.
- 6. Select **Retry account clone**.



If the error has been resolved, the tenant account will now be cloned to the other grid.

Cross-grid replication alerts and errors

Last error shown for connection or tenant

Issue

When [viewing a grid federation connection](#) (or when [managing the permitted tenants](#) for a connection), you notice an error in the **Last error** column on the connection details page. For example:

Grid 1 - Grid 2

Local hostname (this grid): 10.96.130.64
Port: 23000
Remote hostname (other grid): 10.96.130.76
Connection status:  **Connected**

[Edit](#) [Download file](#) [Test connection](#) [Remove](#)

Permitted tenants

[Certificates](#)

[Remove permission](#)

[Clear error](#)



Displaying one result

Tenant
name



Last error



Tenant A

2022-12-22 16:19:20 MST

Cross-grid replication has encountered an error. Failed to send cross-grid replication request from source bucket 'my-bucket' to destination bucket 'my-bucket'. Error code: DestinationRequestError. Detail: InvalidBucketState. Confirm that the source and destination buckets have object versioning enabled and S3 Object Lock disabled. (logID 13916508109026943924)

[Check for errors](#)

Details

For each grid federation connection, the **Last error** column shows the most recent error to occur, if any, when a tenant's data was being replicated to the other grid. This column only shows the last cross-grid replication error to occur; previous errors that might have occurred will not be shown.

An error in this column might occur for one of these reasons:

- The source object version was not found.
- The source bucket was not found.
- The destination bucket was deleted.
- The destination bucket was re-created by a different account.
- The destination bucket has versioning suspended.
- The destination bucket was re-created by the same account but is now unversioned.

Recommended actions

If an error message appears in the **Last error** column, follow these steps:

1. Review the message text.
2. Perform any recommended actions. For example, if versioning was suspended on the destination bucket for cross-grid replication, reenable versioning for that bucket.
3. Select the connection or tenant account from the table.
4. Select **Clear error**.
5. Select **Yes** to clear the message and update the system's status.

- Wait 5-6 minutes and then ingest a new object into the bucket. Confirm that the error message does not reappear.



To ensure the error message is cleared, wait at least 5 minutes after the timestamp in the message before ingesting a new object.



After you clear the error, a new **Last error** might appear if objects are ingested in a different bucket that also has an error.

- To determine if any objects failed to be replicated because of the bucket error, see [Identify and retry failed replication operations](#).

Cross-grid replication permanent failure alert

Issue

The **Cross-grid replication permanent failure** alert was triggered.

Details

This alert indicates that tenant objects can't be replicated between the buckets on two grids for a reason that requires user intervention to resolve. This alert is typically caused by a change to either the source or the destination bucket.

Recommended actions

- Sign in to the grid where the alert was triggered.
- Go to **CONFIGURATION > System > Grid federation**, and locate the connection name listed in the alert.
- On the Permitted tenants tab, look at the **Last error** column to determine which tenant accounts have errors.
- To learn more about the failure, see the instructions in [Monitor grid federation connections](#) to review the cross-grid replication metrics.
- For each affected tenant account:
 - See the instructions in [Monitor tenant activity](#) to confirm that the tenant has not exceeded its quota on the destination grid for cross-grid replication.
 - As required, increase the tenant's quota on the destination grid to allow new objects to be saved.
- For each affected tenant, sign in to Tenant Manager on both grids, so you can compare the list of buckets.
- For each bucket that has cross-grid replication enabled, confirm the following:
 - There is a corresponding bucket for the same tenant on the other grid (must use the exact name).
 - Both buckets have object versioning enabled (versioning can't be suspended on either grid).
 - Both buckets have S3 Object Lock disabled.
 - Neither bucket is in the **Deleting objects: read-only** state.
- To confirm that the issue was resolved, see the instructions in [Monitor grid federation connections](#) to review the cross-grid replication metrics, or perform these steps:
 - Go back to the Grid federation page.
 - Select the affected tenant, and select **Clear Error** in the **Last error** column.
 - Select **Yes** to clear the message and update the system's status.

- d. Wait 5-6 minutes and then ingest a new object into the bucket. Confirm that the error message does not reappear.



To ensure the error message is cleared, wait at least 5 minutes after the timestamp in the message before ingesting a new object.



It might take up to a day for the alert to clear after it is resolved.

- e. Go to [Identify and retry failed replication operations](#) to identify any objects or delete markers that failed to be replicated to the other grid and to retry replication as needed.

Cross-grid replication resource unavailable alert

Issue

The **Cross-grid replication resource unavailable** alert was triggered.

Details

This alert indicates that cross-grid replication requests are pending because a resource is unavailable. For example, there might be a network error.

Recommended actions

1. Monitor the alert to see if the issue resolves on its own.
2. If the issue persists, determine if either grid has a **Grid federation connection failure** alert for the same connection or an **Unable to communicate with node** alert for a node. This alert might be resolved when you resolve those alerts.
3. To learn more about the failure, see the instructions in [Monitor grid federation connections](#) to review the cross-grid replication metrics.
4. If you can't resolve the alert, contact technical support.

Cross-grid replication will proceed as normal after the issue is resolved.

Identify and retry failed replication operations

After resolving the **Cross-grid replication permanent failure** alert, you should determine if any objects or delete markers failed to be replicated to the other grid. You can then reingest these objects or use the Grid Management API to retry replication.

The **Cross-grid replication permanent failure** alert indicates that tenant objects can't be replicated between the buckets on two grids for a reason that requires user intervention to resolve. This alert is typically caused by a change to either the source or the destination bucket. For details, see [Troubleshoot grid federation errors](#).

Determine if any objects failed to be replicated

To determine if any objects or delete markers have not been replicated to the other grid, you can search the audit log for [CGRR \(Cross-Grid Replication Request\)](#) messages. This message is added to the log when StorageGRID fails to replicate an object, multipart object, or delete marker to the destination bucket.

You can use the [audit-explain tool](#) to translate the results into an easier-to-read format.

Before you begin

- You have Root access permission.
- You have the `Passwords.txt` file.
- You know the IP address of the primary Admin Node.

Steps

1. Log in to the primary Admin Node:

- a. Enter the following command: `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 `#`.

2. Search the audit.log for CGRR messages, and use the audit-explain tool to format the results.

For example, this command greps for all CGRR messages in the past 30 minutes and uses the audit-explain tool.

```
# awk -vdate=$(date -d "30 minutes ago" '+%Y-%m-%dT%H:%M:%S') '$1$2 >= date {  
print }' audit.log | grep CGRR | audit-explain
```

The results of the command will look like this example, which has entries for six CGRR messages. In the example, all cross-grid replication requests returned a general error because the object could not be replicated. The first three errors are for "replicate object" operations, and the last three errors are for "replicate delete marker" operations.

```

CGRR Cross-Grid Replication Request tenant:50736445269627437748
connection:447896B6-6F9C-4FB2-95EA-AEBF93A774E9 operation:"replicate
object" bucket:bucket123 object:"audit-0"
version:QjRBNDIzODAtNjQ3My0xMUVELTg2QjEtODJBMjAwQkI3NEM4 error:general
error
CGRR Cross-Grid Replication Request tenant:50736445269627437748
connection:447896B6-6F9C-4FB2-95EA-AEBF93A774E9 operation:"replicate
object" bucket:bucket123 object:"audit-3"
version:QjRDOTRCOUMtNjQ3My0xMUVELTkzM0YtOTg1MTAwQkI3NEM4 error:general
error
CGRR Cross-Grid Replication Request tenant:50736445269627437748
connection:447896B6-6F9C-4FB2-95EA-AEBF93A774E9 operation:"replicate
delete marker" bucket:bucket123 object:"audit-1"
version:NUQ0OEYxMDAtNjQ3NC0xMUVELTg2NjMtOTY5NzAwQkI3NEM4 error:general
error
CGRR Cross-Grid Replication Request tenant:50736445269627437748
connection:447896B6-6F9C-4FB2-95EA-AEBF93A774E9 operation:"replicate
delete marker" bucket:bucket123 object:"audit-5"
version:NUQ1ODUwQkUtNjQ3NC0xMUVELTg1NTItRDkwNzAwQkI3NEM4 error:general
error

```

Each entry contains the following information:

Field	Description
CGRR Cross-Grid Replication Request	The name of the request
tenant	The tenant's account ID
connection	The ID of the grid federation connection
operation	The type of replication operation that was being attempted: • replicate object • replicate delete marker • replicate multipart object
bucket	The bucket name
object	The object name
version	The version ID for the object

Field	Description
error	The type of error. If cross-grid replication failed, the error is "General error".

Retry failed replications

After generating a list of objects and delete markers that were not replicated to the destination bucket and resolving the underlying issues, you can retry replication in either of two ways:

- Reingest each object into the source bucket.
- Use the Grid Management private API, as described.

Steps

1. From the top of the Grid Manager, select the help icon and select **API documentation**.
2. Select **Go to private API documentation**.



The StorageGRID API endpoints that are marked "Private" are subject to change without notice. StorageGRID private endpoints also ignore the API version of the request.

3. In the **cross-grid-replication-advanced** section, select the following endpoint:

```
POST /private/cross-grid-replication-retry-failed
```

4. Select **Try it out**.
5. In the **body** text box, replace the example entry for **versionID** with a version ID from the audit.log that corresponds to a failed cross-grid-replication request.

Be sure to retain the double quotes around the string.

6. Select **Execute**.
7. Confirm that the server response code is **204**, indicating that the object or delete marker has been marked as pending for cross-grid replication to the other grid.



Pending means the cross-grid replication request has been added to the internal queue for processing.

Monitor replication retries

You should monitor the replication retry operations to make sure they complete.



It might take several hours or longer for an object or delete marker to be replicated to the other grid.

You can monitor retry operations in either of two ways:

- Use an S3 [HeadObject](#) or [GetObject](#) request. The response includes the StorageGRID-specific `x-ntap-sg-cgr-replication-status` response header, which will have one of the following values:

Grid	Replication status
Source	• COMPLETED: The replication was successful. • PENDING: The object hasn't been replicated yet. • FAILURE: The replication failed with a permanent failure. A user must resolve the error.
Destination	REPLICA: The object was replicated from the source grid.

- Use the Grid Management private API, as described.

Steps

1. In the **cross-grid-replication-advanced** section of the private API documentation, select the following endpoint:

```
GET /private/cross-grid-replication-object-status/{id}
```

2. Select **Try it out**.
3. In the Parameter section, enter the version ID you used in the `cross-grid-replication-retry-failed` request.
4. Select **Execute**.
5. Confirm that the server response code is **200**.
6. Review the replication status, which will be one of the following:
 - **PENDING:** The object hasn't been replicated yet.
 - **COMPLETED:** The replication was successful.
 - **FAILED:** The replication failed with a permanent failure. A user must resolve the error.

Manage security

Manage security

You can configure various security settings from the Grid Manager to help secure your StorageGRID system.

Manage encryption

StorageGRID provides several options for encrypting data. You should [review the available encryption methods](#) to determine which ones meet your data-protection requirements.

Manage certificates

You can [configure and manage the server certificates](#) used for HTTP connections or the client certificates used to authenticate a client or user identity to the server.

Configure key management servers

Using a [key management server](#) lets you protect StorageGRID data even if an appliance is removed from the data center. After the appliance volumes are encrypted, you can't access any data on the appliance unless the

node can communicate with the KMS.



To use encryption key management, you must enable the **Node Encryption** setting for each appliance during installation, before the appliance is added to the grid.

Manage proxy settings

If you are using S3 platform services or Cloud Storage Pools, you can configure a [storage proxy server](#) between Storage Nodes and the external S3 endpoints. If you send AutoSupport packages using HTTPS or HTTP, you can configure an [admin proxy server](#) between Admin Nodes and technical support.

Control firewalls

To enhance the security of your system, you can control access to StorageGRID Admin Nodes by opening or closing specific ports at the [external firewall](#). You can also control network access to each node by configuring its [internal firewall](#). You can prevent access on all ports except those needed for your deployment.

Review StorageGRID encryption methods

StorageGRID provides several options for encrypting data. You should review the available methods to determine which methods meet your data-protection requirements.

The table provides a high-level summary of the encryption methods available in StorageGRID.

Encryption option	How it works	Applies to
Key management server (KMS) in Grid Manager	You configure a key management server for the StorageGRID site and enable node encryption for the appliance . Then, an appliance node connects to the KMS to request a key encryption key (KEK). This key encrypts and decrypts the data encryption key (DEK) on each volume.	Appliance nodes that have Node Encryption enabled during installation. All data on the appliance is protected against physical loss or removal from the data center. Note: Managing encryption keys with a KMS is only supported for Storage Nodes and services appliances.
Drive Encryption page in StorageGRID Appliance Installer	If the appliance contains drives that support hardware encryption, you can set a drive passphrase during installation. When you set a drive passphrase, it's impossible for anyone to recover valid data from drives that have been removed from the system, unless they know the passphrase. Before starting installation, go to Configure Hardware > Drive Encryption to set a drive passphrase that applies to all StorageGRID-managed, self-encrypting drives in a node.	Appliances that contain self-encrypting drives. All data on the secured drives is protected against physical loss or removal from the data center. Drive encryption doesn't apply to SANtricity-managed drives. If you have a storage appliance with self-encrypting drives and SANtricity controllers, you can enable drive security in SANtricity.

Encryption option	How it works	Applies to
Drive security in SANtricity System Manager	If the Drive Security feature is enabled for your StorageGRID appliance, you can use SANtricity System Manager to create and manage the security key. The key is required to access the data on the secured drives.	Storage appliances that have Full Disk Encryption (FDE) drives or self-encrypting drives. All data on the secured drives is protected against physical loss or removal from the data center. Can't be used with some storage appliances or with any services appliances.
Stored object encryption	You enable the Stored object encryption option in the Grid Manager. When enabled, any new objects that aren't encrypted at the bucket level or at the object level are encrypted during ingest.	Newly ingested S3 object data. Existing stored objects aren't encrypted. Object metadata and other sensitive data aren't encrypted.
S3 bucket encryption	You issue a PutBucketEncryption request to enable encryption for the bucket. Any new objects that aren't encrypted at the object level are encrypted during ingest.	Newly ingested S3 object data only. Encryption must be specified for the bucket. Existing bucket objects aren't encrypted. Object metadata and other sensitive data aren't encrypted. Operations on buckets
S3 object server-side encryption (SSE)	You issue an S3 request to store an object and include the x-amz-server-side-encryption request header.	Newly ingested S3 object data only. Encryption must be specified for the object. Object metadata and other sensitive data aren't encrypted. StorageGRID manages the keys. Use server-side encryption
S3 object server-side encryption with customer-provided keys (SSE-C)	You issue an S3 request to store an object and include three request headers. - x-amz-server-side-encryption-customer-algorithm - x-amz-server-side-encryption-customer-key - x-amz-server-side-encryption-customer-key-MD5	Newly ingested S3 object data only. Encryption must be specified for the object. Object metadata and other sensitive data aren't encrypted. Keys are managed outside of StorageGRID. Use server-side encryption

Encryption option	How it works	Applies to
External volume or datastore encryption	You use an encryption method outside of StorageGRID to encrypt an entire volume or datastore, if your deployment platform supports it.	All object data, metadata, and system configuration data, assuming every volume or datastore is encrypted. An external encryption method provides tighter control over encryption algorithms and keys. Can be combined with the other methods listed.
Object encryption outside of StorageGRID	You use an encryption method outside of StorageGRID to encrypt object data and metadata before they are ingested into StorageGRID.	Object data and metadata only (system configuration data is not encrypted). An external encryption method provides tighter control over encryption algorithms and keys. Can be combined with the other methods listed. Amazon Simple Storage Service - User Guide: Protecting data using client-side encryption

Use multiple encryption methods

Depending on your requirements, you can use more than one encryption method at a time. For example:

- You can use a KMS to protect appliance nodes and also use the drive security feature in SANtricity System Manager to "double encrypt" data on the self-encrypting drives in the same appliances.
- You can use a KMS to secure data on appliance nodes and also use the Stored object encryption option to encrypt all objects when they are ingested.

If only a small portion of your objects require encryption, consider controlling encryption at the bucket or individual object level instead. Enabling multiple levels of encryption has an additional performance cost.

Manage certificates

Manage security certificates

Security certificates are small data files used to create secure, trusted connections between StorageGRID components and between StorageGRID components and external systems.

StorageGRID uses two types of security certificates:

- **Server certificates** are required when you use HTTPS connections. Server certificates are used to establish secure connections between clients and servers, authenticating the identity of a server to its clients and providing a secure communication path for data. The server and the client each have a copy of

the certificate.

- **Client certificates** authenticate a client or user identity to the server, providing more secure authentication than passwords alone. Client certificates don't encrypt data.

When a client connects to the server using HTTPS, the server responds with the server certificate, which contains a public key. The client verifies this certificate by comparing the server signature to the signature on its copy of the certificate. If the signatures match, the client starts a session with the server using the same public key.

StorageGRID functions as the server for some connections (such as the load balancer endpoint) or as the client for other connections (such as the CloudMirror replication service).

Default Grid CA certificate

StorageGRID includes a built-in certificate authority (CA) that generates an internal Grid CA certificate during system installation. The Grid CA certificate is used, by default, to secure internal StorageGRID traffic. An external certificate authority (CA) can issue custom certificates that are fully compliant with your organization's information security policies. Although you can use the Grid CA certificate for a non-production environment, the best practice for a production environment is to use custom certificates signed by an external certificate authority. Unsecured connections with no certificate are also supported but aren't recommended.

- Custom CA certificates don't remove the internal certificates; however, the custom certificates should be the ones specified for verifying server connections.
- All custom certificates must meet the [system hardening guidelines for server certificates](#).
- StorageGRID supports bundling of certificates from a CA into a single file (known as a CA certificate bundle).



StorageGRID also includes operating system CA certificates that are the same on all grids. In production environments, make sure that you specify a custom certificate signed by an external certificate authority in place of the operating system CA certificate.

Variants of the server and client certificate types are implemented in several ways. You should have all the certificates needed for your specific StorageGRID configuration ready before you configure the system.

Access security certificates

You can access information about all StorageGRID certificates in a single location, along with links to the configuration workflow for each certificate.

Steps

1. From Grid Manager, select **CONFIGURATION > Security > Certificates**.

Certificates

View and manage the certificates that secure HTTPS connections between StorageGRID and external clients, such as S3 or Swift, and external servers, such as a key management server (KMS).

Global

Grid CA

Client

Load balancer endpoints

Tenants

Other

The StorageGRID certificate authority ("grid CA") generates and signs two global certificates during installation. The management interface certificate on Admin Nodes secures the management interface. The S3 and Swift API certificate on Storage and Gateway Nodes secures client access. You should replace each default certificate with your own custom certificate signed by an external certificate authority.

Name	Description	Type 	Expiration date  
Management interface certificate	Secures the connection between client web browsers and the Grid Manager, Tenant Manager, Grid Management API, and Tenant Management API.	Custom	Jun 4th, 2022
S3 and Swift API certificate	Secures the connections between S3 and Swift clients and Storage Nodes or between clients and the deprecated CLB service on Gateway Nodes. You can optionally use this certificate for a load balancer endpoint as well.	Custom	Jun 4th, 2022

2. Select a tab on the Certificates page for information about each certificate category and to access the certificate settings. You can access a tab if you have the [appropriate permission](#).
 - **Global:** Secures StorageGRID access from web browsers and external API clients.
 - **Grid CA:** Secures internal StorageGRID traffic.
 - **Client:** Secures connections between external clients and the StorageGRID Prometheus database.
 - **Load balancer endpoints:** Secures connections between S3 clients and the StorageGRID Load Balancer.
 - **Tenants:** Secures connections to identity federation servers or from platform service endpoints to S3 storage resources.
 - **Other:** Secures StorageGRID connections requiring specific certificates.

Each tab is described below with links to additional certificate details.

Global

The global certificates secure StorageGRID access from web browsers and external S3 API clients. Two global certificates are initially generated by the StorageGRID certificate authority during installation. The best practice for a production environment is to use custom certificates signed by an external certificate authority.

- [Management interface certificate](#): Secures client web-browser connections to StorageGRID management interfaces.
- [S3 API certificate](#): Secures client API connections to Storage Nodes, Admin Nodes, and Gateway Nodes, which S3 client applications use to upload and download object data.

Information about the global certificates that are installed includes:

- **Name**: Certificate name with link to managing the certificate.
- **Description**
- **Type**: Custom or default.
You should always use a custom certificate for improved grid security.
- **Expiration date**: If using the default certificate, no expiration date is shown.

You can:

- Replace the default certificates with custom certificates signed by an external certificate authority for improved grid security:
 - [Replace the default StorageGRID-generated management interface certificate](#) used for Grid Manager and Tenant Manager connections.
 - [Replace the S3 API certificate](#) used for Storage Node and load balancer endpoint (optional) connections.
- [Restore the default management interface certificate](#).
- [Restore the default S3 API certificate](#).
- [Use a script to generate a new self-signed management interface certificate](#).
- Copy or download the [management interface certificate](#) or [S3 API certificate](#).

Grid CA

The [Grid CA certificate](#), generated by the StorageGRID certificate authority during StorageGRID installation, secures all internal StorageGRID traffic.

Certificate information includes the certificate expiration date and the certificate contents.

You can [copy or download the Grid CA certificate](#), but you can't change it.

Client

[Client certificates](#), generated by an external certificate authority, secure the connections between external monitoring tools and the StorageGRID Prometheus database.

The certificate table has a row for each configured client certificate and indicates whether the certificate can be used for Prometheus database access, along with the certificate expiration date.

You can:

- [Upload or generate a new client certificate.](#)
- Select a certificate name to display the certificate details where you can:
 - [Change the client certificate name.](#)
 - [Set the Prometheus access permission.](#)
 - [Upload and replace the client certificate.](#)
 - [Copy or download the client certificate.](#)
 - [Remove the client certificate.](#)
- Select **Actions** to quickly [edit](#), [attach](#), or [remove](#) a client certificate. You can select up to 10 client certificates and remove them at one time using **Actions > Remove**.

Load balancer endpoints

[Load balancer endpoint certificates](#) secure the connections between S3 clients and the StorageGRID Load Balancer service on Gateway Nodes and Admin Nodes.

The load balancer endpoint table has a row for each configured load balancer endpoint and indicates whether the global S3 API certificate or a custom load balancer endpoint certificate is being used for the endpoint. The expiration date for each certificate is also displayed.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

You can:

- [View a load balancer endpoint](#), including its certificate details.
- [Specify a load balancer endpoint certificate for FabricPool.](#)
- [Use the global S3 API certificate](#) instead of generating a new load balancer endpoint certificate.

Tenants

Tenants can use [identity federation server certificates](#) or [platform service endpoint certificates](#) to secure their connections with StorageGRID.

The tenant table has a row for each tenant and indicates if each tenant has permission to use its own identity source or platform services.

You can:

- [Select a tenant name to sign in to the Tenant Manager](#)
- [Select a tenant name to view the tenant identity federation details](#)
- [Select a tenant name to view tenant platform services details](#)
- [Specify a platform service endpoint certificate during endpoint creation](#)

Other

StorageGRID uses other security certificates for specific purposes. These certificates are listed by their functional name. Other security certificates include:

- [Cloud Storage Pool certificates](#)

- [Email alert notification certificates](#)
- [External syslog server certificates](#)
- [Grid federation connection certificates](#)
- [Identity federation certificates](#)
- [Key management server \(KMS\) certificates](#)
- [Single sign-on certificates](#)

Information indicates the type of certificate a function uses and its server and client certificate expiration dates, as applicable. Selecting a function name opens a browser tab where you can view and edit the certificate details.



You can only view and access information for other certificates if you have the [appropriate permission](#).

You can:

- [Specify a Cloud Storage Pool certificate for S3, C2S S3, or Azure](#)
- [Specify a certificate for alert email notifications](#)
- [Use a certificate for an external syslog server](#)
- [Rotate grid federation connection certificates](#)
- [View and edit an identity federation certificate](#)
- [Upload key management server \(KMS\) server and client certificates](#)
- [Manually specify an SSO certificate for a relying party trust](#)

Security certificate details

Each type of security certificate is described below, with links to the implementation instructions.

Management interface certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection between client web browsers and the StorageGRID management interface, allowing users to access the Grid Manager and Tenant Manager without security warnings. This certificate also authenticates Grid Management API and Tenant Management API connections. You can use the default certificate created during installation or upload a custom certificate.	CONFIGURATION > Security > Certificates , select the Global tab, and then select Management interface certificate	Configure management interface certificates

S3 API certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates secure S3 client connections to a Storage Node and to load balancer endpoints (optional).	CONFIGURATION > Security > Certificates , select the Global tab, and then select S3 API certificate	Configure S3 API certificates

Grid CA certificate

See the [Default Grid CA certificate description](#).

Administrator client certificate

Certificate type	Description	Navigation location	Details
Client	Installed on each client, allowing StorageGRID to authenticate external client access. - Allows authorized external clients to access the StorageGRID Prometheus database. - Allows secure monitoring of StorageGRID using external tools.	CONFIGURATION > Security > Certificates and then select the Client tab	Configure client certificates

Load balancer endpoint certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection between S3 clients and the StorageGRID Load Balancer service on Gateway Nodes and Admin Nodes. You can upload or generate a load balancer certificate when you configure a load balancer endpoint. Client applications use the load balancer certificate when connecting to StorageGRID to save and retrieve object data. You can also use a custom version of the global S3 API certificate to authenticate connections to the Load Balancer service. If the global certificate is used to authenticate load balancer connections, you don't need to upload or generate a separate certificate for each load balancer endpoint. Note: The certificate used for load balancer authentication is the most used certificate during normal StorageGRID operation.	CONFIGURATION > Network > Load balancer endpoints	• Configure load balancer endpoints • Create a load balancer endpoint for FabricPool

Cloud Storage Pool endpoint certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection from a StorageGRID Cloud Storage Pool to an external storage location, such as S3 Glacier or Microsoft Azure Blob storage. A different certificate is required for each cloud provider type.	ILM > Storage pools	Create a Cloud Storage Pool

Email alert notification certificate

Certificate type	Description	Navigation location	Details
Server and client	Authenticates the connection between an SMTP email server and StorageGRID that is used for alert notifications. • If communications with the SMTP server requires Transport Layer Security (TLS), you must specify the email server CA certificate. • Specify a client certificate only if the SMTP email server requires client certificates for authentication.	ALERTS > Email setup	Set up email notifications for alerts

External syslog server certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the TLS or RELP/TLS connection between an external syslog server that logs events in StorageGRID. Note: An external syslog server certificate is not required for TCP, RELP/TCP, and UDP connections to an external syslog server.	CONFIGURATION > Monitoring > Audit and syslog server	Use an external syslog server

Grid federation connection certificate

Certificate type	Description	Navigation location	Details
Server and client	Authenticate and encrypt information sent between the current StorageGRID system and another grid in a grid federation connection.	CONFIGURATION > System > Grid federation	• Create grid federation connections • Rotate connection certificates

Identity federation certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection between StorageGRID and an external identity provider, such as Active Directory, OpenLDAP, or Oracle Directory Server. Used for identity federation, which allows admin groups and users to be managed by an external system.	CONFIGURATION > Access Control > Identity federation	Use identity federation

Key management server (KMS) certificate

Certificate type	Description	Navigation location	Details
Server and client	Authenticates the connection between StorageGRID and an external key management server (KMS), which provides encryption keys to StorageGRID appliance nodes.	CONFIGURATION > Security > Key management server	Add key management server (KMS)

Platform services endpoint certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection from the StorageGRID platform service to an S3 storage resource.	Tenant Manager > STORAGE (S3) > Platform services endpoints	Create platform services endpoint Edit platform services endpoint

Single sign-on (SSO) certificate

Certificate type	Description	Navigation location	Details
Server	Authenticates the connection between identity federation services, such as Active Directory Federation Services (AD FS), and StorageGRID that are used for single sign-on (SSO) requests.	CONFIGURATION > Access control > Single sign-on	Configure single sign-on

Certificate examples

Example 1: Load Balancer service

In this example, StorageGRID acts as the server.

1. You configure a load balancer endpoint and upload or generate a server certificate in StorageGRID.
2. You configure an S3 client connection to the load balancer endpoint and upload the same certificate to the client.
3. When the client wants to save or retrieve data, it connects to the load balancer endpoint using HTTPS.
4. StorageGRID responds with the server certificate, which contains a public key, and with a signature based on the private key.
5. The client verifies this certificate by comparing the server signature to the signature on its copy of the certificate. If the signatures match, the client starts a session using the same public key.

6. The client sends object data to StorageGRID.

Example 2: External key management server (KMS)

In this example, StorageGRID acts as the client.

1. Using external Key Management Server software, you configure StorageGRID as a KMS client and obtain a CA-signed server certificate, a public client certificate, and the private key for the client certificate.
2. Using the Grid Manager, you configure a KMS server and upload the server and client certificates and the client private key.
3. When a StorageGRID node needs an encryption key, it makes a request to the KMS server that includes data from the certificate and a signature based on the private key.
4. The KMS server validates the certificate signature and decides that it can trust StorageGRID.
5. The KMS server responds using the validated connection.

Supported server certificate types

The StorageGRID system supports custom certificates encrypted with RSA or ECDSA (Elliptic Curve Digital Signature Algorithm).



The cipher type for the security policy must match the server certificate type. For example, RSA ciphers require RSA certificates, and ECDSA ciphers require ECDSA certificates. See [Manage security certificates](#). If you configure a custom security policy that is not compatible with the server certificate, you can [temporarily revert to the default security policy](#).

For more information about how StorageGRID secures client connections, see [Security for S3 clients](#).

Configure management interface certificates

You can replace the default management interface certificate with a single custom certificate that allows users to access the Grid Manager and the Tenant Manager without encountering security warnings. You can also revert to the default management interface certificate or generate a new one.

About this task

By default, every Admin Node is issued a certificate signed by the grid CA. These CA signed certificates can be replaced by a single common custom management interface certificate and corresponding private key.

Because a single custom management interface certificate is used for all Admin Nodes, you must specify the certificate as a wildcard or multi-domain certificate if clients need to verify the hostname when connecting to the Grid Manager and Tenant Manager. Define the custom certificate such that it matches all Admin Nodes in the grid.

You need to complete configuration on the server, and depending on the root certificate authority (CA) you are using, users might also need to install the Grid CA certificate in the web browser they will use to access the Grid Manager and the Tenant Manager.



To ensure that operations aren't disrupted by a failed server certificate, the **Expiration of server certificate for Management Interface** alert is triggered when this server certificate is about to expire. As required, you can view when the current certificate expires by selecting **CONFIGURATION > Security > Certificates** and looking at the Expiration date for the management interface certificate on the Global tab.



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

- Your custom management interface certificate expires.
- You [revert from a custom management interface certificate to the default server certificate](#).

Add a custom management interface certificate

To add a custom management interface certificate, you can provide your own certificate or generate one using the Grid Manager.

Steps

1. Select **CONFIGURATION > Security > Certificates**.
2. On the **Global** tab, select **Management interface certificate**.
3. Select **Use custom certificate**.
4. Upload or generate the certificate.

Upload certificate

Upload the required server certificate files.

a. Select **Upload certificate**.

b. Upload the required server certificate files:

- **Server certificate**: The custom server certificate file (PEM encoded).
- **Certificate private key**: The custom server certificate private key file (.key).



EC private keys must be 224 bits or larger. RSA private keys must be 2048 bits or larger.

- **CA bundle**: A single optional file containing the certificates from each intermediate issuing certificate authority (CA). The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.

c. Expand **Certificate details** to see the metadata for each certificate you uploaded. If you uploaded an optional CA bundle, each certificate displays on its own tab.

- Select **Download certificate** to save the certificate file or select **Download CA bundle** to save the certificate bundle.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select **Copy certificate PEM** or **Copy CA bundle PEM** to copy the certificate contents for pasting elsewhere.

d. Select **Save**.

The custom management interface certificate is used for all subsequent new connections to the Grid Manager, Tenant Manager, Grid Manager API or Tenant Manager API.

Generate certificate

Generate the server certificate files.



The best practice for a production environment is to use a custom management interface certificate signed by an external certificate authority.

a. Select **Generate certificate**.

b. Specify the certificate information:

Field	Description
Domain name	One or more fully qualified domain names to include in the certificate. Use an * as a wildcard to represent multiple domain names.
IP	One or more IP addresses to include in the certificate.

Field	Description
Subject (optional)	X.509 subject or distinguished name (DN) of the certificate owner. If no value is entered in this field, the generated certificate uses the first domain name or IP address as the subject common name (CN).
Days valid	Number of days after creation that the certificate expires.
Add key usage extensions	If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate. These extensions define the purpose of the key contained in the certificate. Note: Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.

c. Select **Generate**.

d. Select **Certificate details** to see the metadata for the generated certificate.

- Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension `.pem`.

For example: `storagegrid_certificate.pem`

- Select **Copy certificate PEM** to copy the certificate contents for pasting elsewhere.

e. Select **Save**.

The custom management interface certificate is used for all subsequent new connections to the Grid Manager, Tenant Manager, Grid Manager API or Tenant Manager API.

5. Refresh the page to ensure the web browser is updated.



After uploading or generating a new certificate, allow up to one day for any related certificate expiration alerts to clear.

6. After you add a custom management interface certificate, the Management interface certificate page displays detailed certificate information for the certificates that are in use.

You can download or copy the certificate PEM as required.

Restore the default management interface certificate

You can revert to using the default management interface certificate for Grid Manager and Tenant Manager connections.

Steps

1. Select **CONFIGURATION > Security > Certificates**.
2. On the **Global** tab, select **Management interface certificate**.
3. Select **Use default certificate**.

When you restore the default management interface certificate, the custom server certificate files you configured are deleted and can't be recovered from the system. The default management interface certificate is used for all subsequent new client connections.

4. Refresh the page to ensure the web browser is updated.

Use a script to generate a new self-signed management interface certificate

If strict hostname validation is required, you can use a script to generate the management interface certificate.

Before you begin

- You have [specific access permissions](#).
- You have the `Passwords.txt` file.

About this task

The best practice for a production environment is to use a certificate signed by an external certificate authority.

Steps

1. Obtain the fully qualified domain name (FQDN) of each Admin Node.
2. Log in to the primary Admin Node:
 - a. Enter the following command: `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 `#`.

3. Configure StorageGRID with a new self-signed certificate.

```
$ sudo make-certificate --domains wildcard-admin-node-fqdn --type management
```

- For `--domains`, use wildcards to represent the fully qualified domain names of all Admin Nodes. For example, `*.ui.storagegrid.example.com` uses the `*` wildcard to represent `admin1.ui.storagegrid.example.com` and `admin2.ui.storagegrid.example.com`.
- Set `--type` to `management` to configure the management interface certificate, which is used by Grid Manager and Tenant Manager.
- By default, generated certificates are valid for one year (365 days) and must be recreated before they expire. You can use the `--days` argument to override the default validity period.



A certificate's validity period begins when `make-certificate` is run. You must ensure the management client is synchronized to the same time source as StorageGRID; otherwise, the client might reject the certificate.

```
$ sudo make-certificate --domains *.ui.storagegrid.example.com --type management --days 720
```

The resulting output contains the public certificate needed by your management API client.

4. Select and copy the certificate.

Include the BEGIN and the END tags in your selection.

5. Log out of the command shell. `$ exit`
6. Confirm the certificate was configured:
 - a. Access the Grid Manager.
 - b. Select **CONFIGURATION > Security > Certificates**
 - c. On the **Global** tab, select **Management interface certificate**.
7. Configure your management client to use the public certificate you copied. Include the BEGIN and END tags.

Download or copy the management interface certificate

You can save or copy the management interface certificate contents for use elsewhere.

Steps

1. Select **CONFIGURATION > Security > Certificates**.
2. On the **Global** tab, select **Management interface certificate**.
3. Select the **Server** or **CA bundle** tab and then download or copy the certificate.

Download certificate file or CA bundle

Download the certificate or CA bundle .pem file. If you are using an optional CA bundle, each certificate in the bundle displays on its own sub-tab.

- a. Select **Download certificate** or **Download CA bundle**.

If you are downloading a CA bundle, all the certificates in the CA bundle secondary tabs download as a single file.

- b. Specify the certificate file name and download location. Save the file with the extension .pem.

For example: `storagegrid_certificate.pem`

Copy certificate or CA bundle PEM

Copy the certificate text to paste elsewhere. If you are using an optional CA bundle, each certificate in the bundle displays on its own sub-tab.

- a. Select **Copy certificate PEM** or **Copy CA bundle PEM**.

If you are copying a CA bundle, all the certificates in the CA bundle secondary tabs copy together.

- b. Paste the copied certificate into a text editor.
- c. Save the text file with the extension .pem.

For example: `storagegrid_certificate.pem`

Configure S3 API certificates

You can replace or restore the server certificate that is used for S3 client connections to Storage Nodes or to load balancer endpoints. The replacement custom server certificate is specific to your organization.



Swift details have been removed from this version of the doc site. See [StorageGRID 11.8: Configure S3 and Swift API certificates](#).

About this task

By default, every Storage Node is issued a X.509 server certificate signed by the grid CA. These CA signed certificates can be replaced by a single common custom server certificate and corresponding private key.

A single custom server certificate is used for all Storage Nodes, so you must specify the certificate as a wildcard or multi-domain certificate if clients need to verify the hostname when connecting to the storage endpoint. Define the custom certificate such that it matches all Storage Nodes in the grid.

After completing configuration on the server, you might also need to install the Grid CA certificate in the S3 API client you will use to access the system, depending on the root certificate authority (CA) you are using.



To ensure that operations aren't disrupted by a failed server certificate, the **Expiration of global server certificate for S3 API** alert is triggered when the root server certificate is about to expire. As required, you can view when the current certificate expires by selecting **CONFIGURATION > Security > Certificates** and looking at the Expiration date for the S3 API certificate on the Global tab.

You can upload or generate a custom S3 API certificate.

Add a custom S3 API certificate

Steps

1. Select **CONFIGURATION > Security > Certificates**.
2. On the **Global** tab, select **S3 API certificate**.
3. Select **Use custom certificate**.
4. Upload or generate the certificate.

Upload certificate

Upload the required server certificate files.

a. Select **Upload certificate**.

b. Upload the required server certificate files:

- **Server certificate**: The custom server certificate file (PEM encoded).
- **Certificate private key**: The custom server certificate private key file (.key).



EC private keys must be 224 bits or larger. RSA private keys must be 2048 bits or larger.

- **CA bundle**: A single optional file containing the certificates from each intermediate issuing certificate authority. The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.

c. Select the certificate details to display the metadata and PEM for each custom S3 API certificate that was uploaded. If you uploaded an optional CA bundle, each certificate displays on its own tab.

- Select **Download certificate** to save the certificate file or select **Download CA bundle** to save the certificate bundle.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select **Copy certificate PEM** or **Copy CA bundle PEM** to copy the certificate contents for pasting elsewhere.

d. Select **Save**.

The custom server certificate is used for subsequent new S3 client connections.

Generate certificate

Generate the server certificate files.

a. Select **Generate certificate**.

b. Specify the certificate information:

Field	Description
Domain name	One or more fully qualified domain names to include in the certificate. Use an * as a wildcard to represent multiple domain names.
IP	One or more IP addresses to include in the certificate.
Subject (optional)	X.509 subject or distinguished name (DN) of the certificate owner. If no value is entered in this field, the generated certificate uses the first domain name or IP address as the subject common name (CN).

Field	Description
Days valid	Number of days after creation that the certificate expires.
Add key usage extensions	If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate. These extensions define the purpose of the key contained in the certificate. Note: Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.

c. Select **Generate**.

d. Select **Certificate Details** to display the metadata and PEM for the custom S3 API certificate that was generated.

- Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension `.pem`.

For example: `storagegrid_certificate.pem`

- Select **Copy certificate PEM** to copy the certificate contents for pasting elsewhere.

e. Select **Save**.

The custom server certificate is used for subsequent new S3 client connections.

5. Select a tab to display metadata for the default StorageGRID server certificate, a CA signed certificate that was uploaded, or a custom certificate that was generated.



After uploading or generating a new certificate, allow up to one day for any related certificate expiration alerts to clear.

6. Refresh the page to ensure the web browser is updated.

7. After you add a custom S3 API certificate the S3 API certificate page displays detailed certificate information for the custom S3 API certificate that is in use.
You can download or copy the certificate PEM as required.

Restore the default S3 API certificate

You can revert to using the default S3 API certificate for S3 client connections to Storage Nodes. However, you can't use the default S3 API certificate for a load balancer endpoint.

Steps

1. Select **CONFIGURATION > Security > Certificates**.
2. On the **Global** tab, select **S3 API certificate**.
3. Select **Use default certificate**.

When you restore the default version of the global S3 API certificate, the custom server certificate files you configured are deleted and can't be recovered from the system. The default S3 API certificate will be used for subsequent new S3 client connections to Storage Nodes.

4. Select **OK** to confirm the warning and restore the default S3 API certificate.

If you have Root access permission and the custom S3 API certificate was used for load balancer endpoint connections, a list is displayed of load balancer endpoints that will no longer be accessible using the default S3 API certificate. Go to [Configure load balancer endpoints](#) to edit or remove the affected endpoints.

5. Refresh the page to ensure the web browser is updated.

Download or copy the S3 API certificate

You can save or copy the S3 API certificate contents for use elsewhere.

Steps

1. Select **CONFIGURATION > Security > Certificates**.
2. On the **Global** tab, select **S3 API certificate**.
3. Select the **Server** or **CA bundle** tab and then download or copy the certificate.

Download certificate file or CA bundle

Download the certificate or CA bundle .pem file. If you are using an optional CA bundle, each certificate in the bundle displays on its own sub-tab.

- a. Select **Download certificate** or **Download CA bundle**.

If you are downloading a CA bundle, all the certificates in the CA bundle secondary tabs download as a single file.

- b. Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

Copy certificate or CA bundle PEM

Copy the certificate text to paste elsewhere. If you are using an optional CA bundle, each certificate in the bundle displays on its own sub-tab.

- a. Select **Copy certificate PEM** or **Copy CA bundle PEM**.

If you are copying a CA bundle, all the certificates in the CA bundle secondary tabs copy together.

- b. Paste the copied certificate into a text editor.
- c. Save the text file with the extension .pem.

For example: storagegrid_certificate.pem

Related information

- [Use S3 REST API](#)
- [Configure S3 endpoint domain names](#)

Copy the Grid CA certificate

StorageGRID uses an internal certificate authority (CA) to secure internal traffic. This certificate does not change if you upload your own certificates.

Before you begin

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

About this task

If a custom server certificate has been configured, client applications should verify the server using the custom server certificate. They should not copy the CA certificate from the StorageGRID system.

Steps

1. Select **CONFIGURATION > Security > Certificates** and then select the **Grid CA** tab.
2. In the **Certificate PEM** section, download or copy the certificate.

Download certificate file

Download the certificate .pem file.

- a. Select **Download certificate**.
- b. Specify the certificate file name and download location. Save the file with the extension .pem.

For example: `storagegrid_certificate.pem`

Copy certificate PEM

Copy the certificate text to paste elsewhere.

- a. Select **Copy certificate PEM**.
- b. Paste the copied certificate into a text editor.
- c. Save the text file with the extension .pem.

For example: `storagegrid_certificate.pem`

Configure StorageGRID certificates for FabricPool

For S3 clients that perform strict hostname validation and don't support disabling strict hostname validation, such as ONTAP clients using FabricPool, you can generate or upload a server certificate when you configure the load balancer endpoint.

Before you begin

- You have [specific access permissions](#).

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

About this task

When you create a load balancer endpoint, you can generate a self-signed server certificate or upload a certificate that is signed by a known certificate authority (CA). In production environments, you should use a certificate that is signed by a known CA. Certificates signed by a CA can be rotated non-disruptively. They are also more secure because they provide better protection against man-in-the-middle attacks.

The following steps provide general guidelines for S3 clients that use FabricPool. For more detailed information and procedures, see [Configure StorageGRID for FabricPool](#).

Steps

1. Optionally, configure a high availability (HA) group for FabricPool to use.
2. Create an S3 load balancer endpoint for FabricPool to use.

When you create an HTTPS load balancer endpoint, you are prompted to upload your server certificate, certificate private key, and optional CA bundle.

3. Attach StorageGRID as a cloud tier in ONTAP.

Specify the load balancer endpoint port and the fully qualified domain name used in the CA certificate you uploaded. Then, provide the CA certificate.



If an intermediate CA issued the StorageGRID certificate, you must provide the intermediate CA certificate. If the StorageGRID certificate was issued directly by the Root CA, you must provide the Root CA certificate.

Configure client certificates

Client certificates allow authorized external clients to access the StorageGRID Prometheus database, providing a secure way for external tools to monitor StorageGRID.

If you need to access StorageGRID using an external monitoring tool, you must upload or generate a client certificate using the Grid Manager and copy the certificate information to the external tool.

See [Manage security certificates](#) and [Configure custom server certificates](#).



To ensure that operations aren't disrupted by a failed server certificate, the **Expiration of client certificates configured on the Certificates page** alert is triggered when this server certificate is about to expire. As required, you can view when the current certificate expires by selecting **CONFIGURATION > Security > Certificates** and looking at the Expiration date for the client certificate on the Client tab.



If you are using a key management server (KMS) to protect the data on specially configured appliance nodes, see the specific information about [uploading a KMS client certificate](#).

Before you begin

- You have Root access permission.
- You are signed in to the Grid Manager using a [supported web browser](#).
- To configure a client certificate:

- You have the IP address or domain name of the Admin Node.
- If you have configured the StorageGRID management interface certificate, you have the CA, client certificate, and private key used to configure the management interface certificate.
- To upload your own certificate, the private key for the certificate is available on your local computer.
- The private key must have been saved or recorded at the time it was created. If you don't have the original private key, you must create a new one.
- To edit a client certificate:
 - You have the IP address or domain name of the Admin Node.
 - To upload your own certificate or a new certificate, the private key, client certificate, and CA (if used) are available on your local computer.

Add client certificates

To add the client certificate, use one of these procedures:

- [Management interface certificate already configured](#)
- [CA issued client certificate](#)
- [Generated certificate from Grid Manager](#)

Management interface certificate already configured

Use this procedure to add a client certificate if a management interface certificate is already configured using a customer-supplied CA, client certificate, and private key.

Steps

1. In the Grid Manager, select **CONFIGURATION** > **Security** > **Certificates** and then select the **Client** tab.
2. Select **Add**.
3. Enter a certificate name.
4. To access Prometheus metrics using your external monitoring tool, select **Allow prometheus**.
5. Select **Continue**.
6. For the **Attach certificates** step, upload the management interface certificate.
 - a. Select **Upload certificate**.
 - b. Select **Browse** and select the management interface certificate file (.pem).
 - Select **Client certificate details** to display the certificate metadata and certificate PEM.
 - Select **Copy certificate PEM** to copy the certificate contents for pasting elsewhere.
 - c. Select **Create** to save the certificate in the Grid Manager.

The new certificate appears on the Client tab.

7. [Configure an external monitoring tool](#), such as Grafana.

CA issued client certificate

Use this procedure to add an administrator client certificate if a management interface certificate was not configured and you plan to add a client certificate for Prometheus that uses a CA issued client certificate and private key.

Steps

1. Perform the steps to [configure a management interface certificate](#).
2. In the Grid Manager, select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.
3. Select **Add**.
4. Enter a certificate name.
5. To access Prometheus metrics using your external monitoring tool, select **Allow prometheus**.
6. Select **Continue**.
7. For the **Attach certificates** step, upload the client certificate, private key, and CA bundle files:
 - a. Select **Upload certificate**.
 - b. Select **Browse** and select the client certificate, private key, and CA bundle files (.pem).
 - Select **Client certificate details** to display the certificate metadata and certificate PEM.
 - Select **Copy certificate PEM** to copy the certificate contents for pasting elsewhere.
 - c. Select **Create** to save the certificate in the Grid Manager.

The new certificates appear on the Client tab.

8. [Configure an external monitoring tool](#), such as Grafana.

Generated certificate from Grid Manager

Use this procedure to add an administrator client certificate if a management interface certificate was not configured and you plan to add a client certificate for Prometheus that uses the generate certificate function in Grid Manager.

Steps

1. In the Grid Manager, select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.
2. Select **Add**.
3. Enter a certificate name.
4. To access Prometheus metrics using your external monitoring tool, select **Allow prometheus**.
5. Select **Continue**.
6. For the **Attach certificates** step, select **Generate certificate**.
7. Specify the certificate information:
 - **Subject** (optional): X.509 subject or distinguished name (DN) of the certificate owner.
 - **Days valid**: The number of days the generated certificate is valid, starting at the time it is generated.
 - **Add key usage extensions**: If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate.

These extensions define the purpose of the key contained in the certificate.



Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.

8. Select **Generate**.

9. Select **Client certificate details** to display the certificate metadata and certificate PEM.



You will not be able to view the certificate private key after you close the dialog. Copy or download the key to a safe location.

- Select **Copy certificate PEM** to copy the certificate contents for pasting elsewhere.
- Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension `.pem`.

For example: `storagegrid_certificate.pem`

- Select **Copy private key** to copy the certificate private key for pasting elsewhere.
- Select **Download private key** to save the private key as a file.

Specify the private key file name and download location.

10. Select **Create** to save the certificate in the Grid Manager.

The new certificate appears on the Client tab.

11. In the Grid Manager, select **CONFIGURATION > Security > Certificates** and then select the **Global** tab.

12. Select **Management Interface certificate**.

13. Select **Use custom certificate**.

14. Upload the `certificate.pem` and `private_key.pem` files from the [client certificate details](#) step. There is no need to upload CA bundle.
 - a. Select **Upload certificate** and then select **Continue**.
 - b. Upload each certificate file (`.pem`).
 - c. Select **Save** to save the certificate in the Grid Manager.

The new certificate appears on the Management Interface certificate page.

15. [Configure an external monitoring tool](#), such as Grafana.

Configure an external monitoring tool

Steps

1. Configure the following settings on your external monitoring tool, such as Grafana.

- a. **Name:** Enter a name for the connection.

StorageGRID does not require this information, but you must provide a name to test the connection.

- b. **URL:** Enter the domain name or IP address for the Admin Node. Specify HTTPS and port 9091.

For example: `https://admin-node.example.com:9091`

- c. Enable **TLS Client Auth** and **With CA Cert**.
- d. Under TLS/SSL Auth Details, copy and paste: +
 - The management interface CA certificate to **CA Cert**

- The client certificate to **Client Cert**
- The private key to **Client Key**

e. **ServerName**: Enter the domain name of the Admin Node.

ServerName must match the domain name as it appears in the management interface certificate.

2. Save and test the certificate and private key that you copied from StorageGRID or a local file.

You can now access the Prometheus metrics from StorageGRID with your external monitoring tool.

For information about the metrics, see the [instructions for monitoring StorageGRID](#).

Edit client certificates

You can edit an administrator client certificate to change its name, enable or disable Prometheus access, or upload a new certificate when the current one has expired.

Steps

1. Select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.

Certificate expiration dates and Prometheus access permissions are listed in the table. If a certificate will expire soon or is already expired, a message appears in the table and an alert is triggered.

2. Select the certificate you want to edit.

3. Select **Edit** and then select **Edit name and permission**

4. Enter a certificate name.

5. To access Prometheus metrics using your external monitoring tool, select **Allow prometheus**.

6. Select **Continue** to save the certificate in the Grid Manager.

The updated certificate displays on the Client tab.

Attach new client certificate

You can upload a new certificate when the current one has expired.

Steps

1. Select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.

Certificate expiration dates and Prometheus access permissions are listed in the table. If a certificate will expire soon or is already expired, a message appears in the table and an alert is triggered.

2. Select the certificate you want to edit.

3. Select **Edit** and then select an edit option.

Upload certificate

Copy the certificate text to paste elsewhere.

- a. Select **Upload certificate** and then select **Continue**.
- b. Upload the client certificate name (.pem).

Select **Client certificate details** to display the certificate metadata and certificate PEM.

- Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select **Copy certificate PEM** to copy the certificate contents for pasting elsewhere.
- c. Select **Create** to save the certificate in the Grid Manager.

The updated certificate displays on the Client tab.

Generate certificate

Generate the certificate text to paste elsewhere.

- a. Select **Generate certificate**.
- b. Specify the certificate information:
 - **Subject** (optional): X.509 subject or distinguished name (DN) of the certificate owner.
 - **Days valid**: The number of days the generated certificate is valid, starting at the time it is generated.
 - **Add key usage extensions**: If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate.

These extensions define the purpose of the key contained in the certificate.



Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.

- c. Select **Generate**.
- d. Select **Client certificate details** to display the certificate metadata and certificate PEM.



You will not be able to view the certificate private key after you close the dialog. Copy or download the key to a safe location.

- Select **Copy certificate PEM** to copy the certificate contents for pasting elsewhere.
- Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select **Copy private key** to copy the certificate private key for pasting elsewhere.
- Select **Download private key** to save the private key as a file.

Specify the private key file name and download location.

e. Select **Create** to save the certificate in the Grid Manager.

The new certificate appears on the Client tab.

Download or copy client certificates

You can download or copy a client certificate for use elsewhere.

Steps

1. Select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.
2. Select the certificate you want to copy or download.
3. Download or copy the certificate.

Download certificate file

Download the certificate .pem file.

- a. Select **Download certificate**.
- b. Specify the certificate file name and download location. Save the file with the extension .pem.

For example: `storagegrid_certificate.pem`

Copy certificate

Copy the certificate text to paste elsewhere.

- a. Select **Copy certificate PEM**.
- b. Paste the copied certificate into a text editor.
- c. Save the text file with the extension .pem.

For example: `storagegrid_certificate.pem`

Remove client certificates

If you no longer need an administrator client certificate, you can remove it.

Steps

1. Select **CONFIGURATION > Security > Certificates** and then select the **Client** tab.
2. Select the certificate you want to remove.
3. Select **Delete** and then confirm.



To remove up to 10 certificates, select each certificate to remove on the Client tab and then select **Actions** > **Delete**.

After a certificate is removed, clients that used the certificate must specify a new client certificate to access the StorageGRID Prometheus database.

Configure security settings

Manage the TLS and SSH policy

The TLS and SSH policy determines which protocols and ciphers are used to establish secure TLS connections with client applications and secure SSH connections to internal StorageGRID services.

The security policy controls how TLS and SSH encrypt data in motion. In general, use the Modern compatibility (default) policy, unless your system needs to be Common Criteria-compliant or you need to use other ciphers.



Some StorageGRID services have not been updated to use the ciphers in these policies.

Before you begin

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

Select a security policy

Steps

1. Select **CONFIGURATION** > **Security** > **Security settings**.

The **TLS and SSH policies** tab shows the available policies. The currently active policy is noted by a green check mark on the policy tile.



2. Review the tiles to learn about the available policies.

Policy	Description
Modern compatibility (default)	Use the default policy if you need strong encryption and unless you have special requirements. This policy is compatible with most TLS and SSH clients.

Policy	Description
Legacy compatibility	Use this policy if you need additional compatibility options for older clients. The additional options in this policy might make it less secure than the Modern compatibility policy.
Common Criteria	Use this policy if you require Common Criteria certification.
FIPS strict	Use this policy if you require Common Criteria certification and need to use the NetApp Cryptographic Security Module 3.0.8 for external client connections to load balancer endpoints, Tenant Manager, and Grid Manager. Using this policy might reduce performance. Note: After you select this policy, all nodes must be rebooted in a rolling fashion to activate the NetApp Cryptographic Security Module. Use Maintenance > Rolling reboot to initiate and monitor reboots.
Custom	Create a custom policy if you need to apply your own ciphers.

3. To see details about each policy's ciphers, protocols, and algorithms, select **View details**.
4. To change the current policy, select **Use policy**.


A green check mark appears next to **Current policy** on the policy tile.

Create a custom security policy

You can create a custom policy if you need to apply your own ciphers.

Steps

1. From the tile of the policy that is the most similar to the custom policy you want to create, select **View details**.
2. Select **Copy to clipboard**, and then select **Cancel**.



3. From the **Custom policy** tile, select **Configure and use**.

4. Paste the JSON you copied and make any changes required.
5. Select **Use policy**.

A green check mark appears next to **Current policy** on the Custom policy tile.

6. Optionally, select **Edit configuration** to make more changes to the new custom policy.

Temporarily revert to the default security policy

If you configured a custom security policy, you might not be able to sign in to the Grid Manager if the configured TLS policy is incompatible with the [configured server certificate](#).

You can temporarily revert to the default security policy.

Steps

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

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

2. Run the following command:

```
restore-default-cipher-configurations
```

3. From a web browser, access the Grid Manager on the same Admin Node.
4. Follow the steps in [Select a security policy](#) to configure the policy again.

Configure network and object security

You can configure network and object security to encrypt stored objects, to prevent certain S3 requests, or to allow client connections to Storage Nodes to use HTTP instead of HTTPS.

Stored object encryption

Stored object encryption enables the encryption of all object data as it is ingested through S3. By default, stored objects aren't encrypted but you can choose to encrypt objects using the AES-128 or AES-256 encryption algorithm. When you enable the setting, all newly ingested objects are encrypted but no change is made to existing stored objects. If you disable encryption, currently encrypted objects remain encrypted but newly ingested objects aren't encrypted.

The Stored object encryption setting applies only to S3 objects that have not been encrypted by bucket-level or object-level encryption.

For more details on StorageGRID encryption methods, see [Review StorageGRID encryption methods](#).

Prevent client modification

Prevent client modification is a system wide setting. When the **Prevent client modification** option is selected, the following requests are denied.

S3 REST API

- DeleteBucket requests
- Any requests to modify an existing object's data, user-defined metadata, or S3 object tagging

Enable HTTP for Storage Node connections

By default, client applications use the HTTPS network protocol for any direct connections to Storage Nodes. You can optionally enable HTTP for these connections, for example, when testing a non-production grid.

Use HTTP for Storage Node connections only if S3 clients need to make HTTP connections directly to Storage Nodes. You don't need to use this option for clients that only use HTTPS connections or for clients that connect to the Load Balancer service (because you can [configure each load balancer endpoint](#) to use either HTTP or HTTPS).

See [Summary: IP addresses and ports for client connections](#) to learn which ports S3 clients use when connecting to Storage Nodes using HTTP or HTTPS.

Select options

Before you begin

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

Steps

1. Select **CONFIGURATION > Security > Security settings**.
2. Select the **Network and objects** tab.
3. For Stored object encryption, use the **None** (default) setting if you don't want stored objects to be encrypted, or select **AES-128** or **AES-256** to encrypt stored objects.
4. Optionally select **Prevent client modification** if you want to prevent S3 clients from making specific requests.



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.

5. Optionally select **Enable HTTP for Storage Node connections** if clients connect directly to Storage Nodes and you want to use HTTP connections.



Be careful when enabling HTTP for a production grid because requests will be sent unencrypted.

6. Select **Save**.

Change interface security settings

The interface security settings let you control whether users are signed out if they are

inactive for more than the specified amount of time and whether a stack trace is included in API error responses.

Before you begin

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

About this task

The **Security settings** page includes the **Browser inactivity timeout** and **Management API stack trace** settings.

Browser inactivity timeout

Indicates how long a user's browser can be inactive before the user is signed out. The default is 15 minutes.

Browser inactivity timeout is also controlled by the following:

- A separate, non-configurable StorageGRID timer, which is included for system security. Each user's authentication token expires 16 hours after the user signs in. When a user's authentication expires, that user is automatically signed out, even if browser inactivity timeout is disabled or the value for the browser timeout has not been reached. To renew the token, the user must sign back in.
- Timeout settings for the identity provider, assuming single sign-on (SSO) is enabled for StorageGRID.

If SSO is enabled and a user's browser times out, the user must reenter their SSO credentials to access StorageGRID again. See [Configure single sign-on](#).

Management API stack trace

Controls whether a stack trace is returned in Grid Manager and Tenant Manager API error responses.

This option is disabled by default, but you might want to enable this functionality for a test environment. In general, you should leave stack trace disabled in production environments to avoid revealing internal software details when API errors occur.

Steps

1. Select **CONFIGURATION > Security > Security settings**.
2. Select the **Interface** tab.
3. To change the setting for browser inactivity timeout:
 - a. Expand the accordion.
 - b. To change the timeout period, specify a value between 60 seconds and 7 days. The default timeout is 15 minutes.
 - c. To disable this feature, unselect the checkbox.
 - d. Select **Save**.

The new setting doesn't affect users who are currently signed in. Users must sign in again or refresh their browsers for the new timeout setting to take effect.

4. To change the setting for Management API stack trace:
 - a. Expand the accordion.

- b. Select the checkbox to return a stack trace in Grid Manager and Tenant Manager API error responses.



Leave stack trace disabled in production environments to avoid revealing internal software details when API errors occur.

- c. Select **Save**.

Configure key management servers

What is a key management server (KMS)?

A key management server (KMS) is an external, third-party system that provides encryption keys to StorageGRID appliance nodes at the associated StorageGRID site using the Key Management Interoperability Protocol (KMIP).

StorageGRID supports only certain key management servers. For a list of supported products and versions, use the [NetApp Interoperability Matrix Tool \(IMT\)](#).

You can use one or more key management servers to manage the node encryption keys for any StorageGRID appliance nodes that have the **Node Encryption** setting enabled during installation. Using key management servers with these appliance nodes lets you protect your data even if an appliance is removed from the data center. After the appliance volumes are encrypted, you can't access any data on the appliance unless the node can communicate with the KMS.




StorageGRID does not create or manage the external keys used to encrypt and decrypt appliance nodes. If you plan to use an external key management server to protect StorageGRID data, you must understand how to set up that server, and you must understand how to manage the encryption keys. Performing key management tasks is beyond the scope of these instructions. If you need help, see the documentation for your key management server or contact technical support.

KMS and appliance configuration

Before you can use a key management server (KMS) to secure StorageGRID data on appliance nodes, you must complete two configuration tasks: setting up one or more KMS servers and enabling node encryption for the appliance nodes. When these two configuration tasks are complete, the key management process occurs automatically.

The flowchart shows the high-level steps for using a KMS to secure StorageGRID data on appliance nodes.



The flowchart shows KMS setup and appliance setup occurring in parallel; however, you can set up the key

management servers before or after you enable node encryption for new appliance nodes, based on your requirements.

Set up the key management server (KMS)

Setting up a key management server includes the following high-level steps.

Step	Refer to
Access the KMS software and add a client for StorageGRID to each KMS or KMS cluster.	Configure StorageGRID as a client in the KMS
Obtain the required information for the StorageGRID client on the KMS.	Configure StorageGRID as a client in the KMS
Add the KMS to the Grid Manager, assign it to a single site or to a default group of sites, upload the required certificates, and save the KMS configuration.	Add a key management server (KMS)

Set up the appliance

Setting up an appliance node for KMS use includes the following high-level steps.

1. During the hardware configuration stage of appliance installation, use the StorageGRID Appliance Installer to enable the **Node Encryption** setting for the appliance.



You can't enable the **Node Encryption** setting after an appliance is added to the grid, and you can't use external key management for appliances that don't have node encryption enabled.

2. Run the StorageGRID Appliance Installer. During installation, a random data encryption key (DEK) is assigned to each appliance volume, as follows:
 - The DEKs are used to encrypt the data on each volume. These keys are generated using Linux Unified Key Setup (LUKS) disk encryption in the appliance OS and can't be changed.
 - Each individual DEK is encrypted by a master key encryption key (KEK). The initial KEK is a temporary key that encrypts the DEKs until the appliance can connect to the KMS.
3. Add the appliance node to StorageGRID.

See [Enable node encryption](#) for details.

Key management encryption process (occurs automatically)

Key management encryption includes the following high-level steps that are performed automatically.

1. When you install an appliance that has node encryption enabled into the grid, StorageGRID determines if a KMS configuration exists for the site that contains the new node.
 - If a KMS has already been configured for the site, the appliance receives the KMS configuration.
 - If a KMS has not yet been configured for the site, data on the appliance continues to be encrypted by the temporary KEK until you configure a KMS for the site and the appliance receives the KMS configuration.

2. The appliance uses the KMS configuration to connect to the KMS and request an encryption key.
3. The KMS sends an encryption key to the appliance. The new key from the KMS replaces the temporary KEK and is now used to encrypt and decrypt the DEKs for the appliance volumes.



Any data that exists before the encrypted appliance node connects to the configured KMS is encrypted with a temporary key. However, the appliance volumes should not be considered protected from removal from the data center until the temporary key is replaced by the KMS encryption key.

4. If the appliance is powered on or rebooted, it reconnects to the KMS to request the key. The key, which is saved in volatile memory, can't survive a loss of power or a reboot.

Considerations and requirements for using a key management server

Before configuring an external key management server (KMS), you must understand the considerations and requirements.

Which version of KMIP is supported?

StorageGRID supports KMIP version 1.4.

[Key Management Interoperability Protocol Specification Version 1.4](#)

What are the network considerations?

The network firewall settings must allow each appliance node to communicate through the port used for Key Management Interoperability Protocol (KMIP) communications. The default KMIP port is 5696.

You must ensure that each appliance node that uses node encryption has network access to the KMS or KMS cluster you configured for the site.

Which versions of TLS are supported?

Communications between the appliance nodes and the configured KMS use secure TLS connections. StorageGRID can support either the TLS 1.2 or TLS 1.3 protocol when it makes KMIP connections to a KMS or KMS cluster, based on what the KMS supports and which [TLS and SSH policy](#) you are using.

StorageGRID negotiates the protocol and cipher (TLS 1.2) or cipher suite (TLS 1.3) with the KMS when it makes the connection. To see which protocol versions and ciphers/cipher suites are available, review the `tlsOutbound` section of the grid's active TLS and SSH policy (**CONFIGURATION > Security Security settings**).

Which appliances are supported?

You can use a key management server (KMS) to manage encryption keys for any StorageGRID appliance in your grid that has the **Node Encryption** setting enabled. This setting can only be enabled during the hardware configuration stage of appliance installation using the StorageGRID Appliance Installer.



You can't enable node encryption after an appliance is added to the grid, and you can't use external key management for appliances that don't have node encryption enabled.

You can use the configured KMS for StorageGRID appliances and appliance nodes.

You can't use the configured KMS for software-based (non-appliance) nodes, including the following:

- Nodes deployed as virtual machines (VMs)
- Nodes deployed within container engines on Linux hosts

Nodes deployed on these other platforms can use encryption outside of StorageGRID at the datastore or disk level.

When should I configure key management servers?

For a new installation, you should typically set up one or more key management servers in the Grid Manager before creating tenants. This order ensures that the nodes are protected before any object data is stored on them.

You can configure the key management servers in the Grid Manager before or after you install the appliance nodes.


How many key management servers do I need?

You can configure one or more external key management servers to provide encryption keys to the appliance nodes in your StorageGRID system. Each KMS provides a single encryption key to the StorageGRID appliance nodes at a single site or at a group of sites.

StorageGRID supports the use of KMS clusters. Each KMS cluster contains multiple, replicated key management servers that share configuration settings and encryption keys. Using KMS clusters for key management is recommended because it improves the failover capabilities of a high availability configuration.

For example, suppose your StorageGRID system has three data center sites. You might configure one KMS cluster to provide a key to all appliance nodes at Data Center 1 and a second KMS cluster to provide a key to all appliance nodes at all other sites. When you add the second KMS cluster, you can configure a default KMS for Data Center 2 and Data Center 3.

Note that you can't use a KMS for non-appliance nodes or for any appliance nodes that did not have the **Node Encryption** setting enabled during installation.



What happens when a key is rotated?

As a security best practice, you should periodically [rotate the encryption key](#) used by each configured KMS.

When the new key version is available:

- It is automatically distributed to the encrypted appliance nodes at the site or sites associated with the KMS. The distribution should occur within an hour of when the key is rotated.
- If the encrypted appliance node is offline when the new key version is distributed, the node will receive the new key as soon as it reboots.
- If the new key version can't be used to encrypt appliance volumes for any reason, the **KMS encryption key rotation failed** alert is triggered for the appliance node. You might need to contact technical support for help in resolving this alert.

Can I reuse an appliance node after it has been encrypted?

If you need to install an encrypted appliance into another StorageGRID system, you must first decommission the grid node to move object data to another node. Then, you can use the StorageGRID Appliance Installer to [clear the KMS configuration](#). Clearing the KMS configuration disables the **Node Encryption** setting and removes the association between the appliance node and the KMS configuration for the StorageGRID site.



With no access to the KMS encryption key, any data that remains on the appliance can no longer be accessed and is permanently locked.


Considerations for changing the KMS for a site

Each key management server (KMS) or KMS cluster provides an encryption key to all appliance nodes at a single site or at a group of sites. If you need to change which KMS is used for a site, you might need to copy the encryption key from one KMS to another.


If you change the KMS used for a site, you must ensure that the previously encrypted appliance nodes at that site can be decrypted using the key stored on the new KMS. In some cases, you might need to copy the current version of the encryption key from the original KMS to the new KMS. You must ensure that the KMS has the correct key to decrypt the encrypted appliance nodes at the site.

For example:


1. You initially configure a default KMS that applies to all sites that don't have a dedicated KMS.
2. When the KMS is saved, all appliance nodes that have the **Node Encryption** setting enabled connect to the KMS and request the encryption key. This key is used to encrypt the appliance nodes at all sites. This same key must also be used to decrypt those appliances.



3. You decide to add a site-specific KMS for one site (Data Center 3 in the figure). However, because the appliance nodes are already encrypted, a validation error occurs when you attempt to save the configuration for the site-specific KMS. The error occurs because the site-specific KMS does not have the correct key to decrypt the nodes at that site.



- To address the issue, you copy the current version of the encryption key from the default KMS to the new KMS. (Technically, you copy the original key to a new key with the same alias. The original key becomes a prior version of the new key.) The site-specific KMS now has the correct key to decrypt the appliance nodes at Data Center 3, so it can be saved in StorageGRID.



Use cases for changing which KMS is used for a site

The table summarizes the required steps for the most common cases for changing the KMS for a site.

Use case for changing a site's KMS	Required steps
You have one or more site-specific KMS entries, and you want to use one of them as the default KMS.	Edit the site-specific KMS. In the Manages keys for field, select Sites not managed by another KMS (default KMS). The site-specific KMS will now be used as the default KMS. It will apply to any sites that don't have a dedicated KMS. Edit a key management server (KMS)

Use case for changing a site's KMS	Required steps
You have a default KMS and you add a new site in an expansion. You don't want to use the default KMS for the new site.	1. If the appliance nodes at the new site have already been encrypted by the default KMS, use the KMS software to copy the current version of the encryption key from the default KMS to a new KMS. 2. Using the Grid Manager, add the new KMS and select the site. Add a key management server (KMS)
You want the KMS for a site to use a different server.	1. If the appliance nodes at the site have already been encrypted by the existing KMS, use the KMS software to copy the current version of the encryption key from the existing KMS to the new KMS. 2. Using the Grid Manager, edit the existing KMS configuration and enter the new host name or IP address. Add a key management server (KMS)

Configure StorageGRID as a client in the KMS

You must configure StorageGRID as a client for each external key management server or KMS cluster before you can add the KMS to StorageGRID.



These instructions apply to Thales CipherTrust Manager and Hashicorp Vault. For a list of supported products and versions, use the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Steps

1. From the KMS software, create a StorageGRID client for each KMS or KMS cluster you plan to use.

Each KMS manages a single encryption key for the StorageGRID appliances nodes at a single site or at a group of sites.

2. Create a key using one of the following two methods:

- Use the key management page of your KMS product. Create an AES encryption key for each KMS or KMS cluster.

The encryption key must be 2,048 bits or more, and it must be exportable.

- Have StorageGRID create the key. You will be prompted when you test and save after [uploading client certificates](#).

3. Record the following information for each KMS or KMS cluster.

You need this information when you add the KMS to StorageGRID:

- Host name or IP address for each server.
- KMIP port used by the KMS.
- Key alias for the encryption key in the KMS.

4. For each KMS or KMS cluster, obtain a server certificate signed by a certificate authority (CA) or a certificate bundle that contains each of the PEM-encoded CA certificate files, concatenated in certificate

chain order.

The server certificate allows the external KMS to authenticate itself to StorageGRID.

- The certificate must use the Privacy Enhanced Mail (PEM) Base-64 encoded X.509 format.
- The Subject Alternative Name (SAN) field in each server certificate must include the fully qualified domain name (FQDN) or IP address that StorageGRID will connect to.



When you configure the KMS in StorageGRID, you must enter the same FQDNs or IP addresses in the **Hostname** field.

- The server certificate must match the certificate used by the KMIP interface of the KMS, which typically uses port 5696.
5. Obtain the public client certificate issued to StorageGRID by the external KMS and the private key for the client certificate.

The client certificate allows StorageGRID to authenticate itself to the KMS.

Add a key management server (KMS)

You use the StorageGRID Key Management Server wizard to add each KMS or KMS cluster.

Before you begin

- You have reviewed the [considerations and requirements for using a key management server](#).
- You have [configured StorageGRID as a client in the KMS](#), and you have the required information for each KMS or KMS cluster.
- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).

About this task

If possible, configure any site-specific key management servers before configuring a default KMS that applies to all sites not managed by another KMS. If you create the default KMS first, all node-encrypted appliances in the grid will be encrypted by the default KMS. If you want to create a site-specific KMS later, you must first copy the current version of the encryption key from the default KMS to the new KMS. See [Considerations for changing the KMS for a site](#) for details.

Step 1: KMS details

In Step 1 (KMS details) of the Add a Key Management Server wizard, you provide details about the KMS or KMS cluster.

Steps

1. Select **CONFIGURATION > Security > Key management server**.

The Key management server page appears with the Configuration details tab selected.

2. Select **Create**.

Step 1 (KMS details) of the Add a Key Management Server wizard appears.

3. Enter the following information for the KMS and the StorageGRID client you configured in that KMS.

Field	Description
KMS name	A descriptive name to help you identify this KMS. Must be between 1 and 64 characters.
Key name	The exact key alias for the StorageGRID client in the KMS. Must be between 1 and 255 characters. Note: If you haven't created a key using your KMS product, you'll be prompted to have StorageGRID create the key.
Manages keys for	The StorageGRID site that will be associated with this KMS. If possible, you should configure any site-specific key management servers before configuring a default KMS that applies to all sites not managed by another KMS. • Select a site if this KMS will manage encryption keys for the appliance nodes at a specific site. • Select Sites not managed by another KMS (default KMS) to configure a default KMS that will apply to any sites that don't have a dedicated KMS and to any sites you add in subsequent expansions. Note: A validation error will occur when you save the KMS configuration if you select a site that was previously encrypted by the default KMS but you did not provide the current version of original encryption key to the new KMS.
Port	The port the KMS server uses for Key Management Interoperability Protocol (KMIP) communications. Defaults to 5696, which is the KMIP standard port.
Hostname	The fully qualified domain name or IP address for the KMS. Note: The Subject Alternative Name (SAN) field of the server certificate must include the FQDN or IP address you enter here. Otherwise, StorageGRID will not be able to connect to the KMS or to all servers in a KMS cluster.

4. If you are configuring a KMS cluster, select **Add another hostname** to add a hostname for each server in the cluster.

5. Select **Continue**.

Step 2: Upload server certificate

In Step 2 (Upload server certificate) of the Add a Key Management Server wizard, you upload the server certificate (or certificate bundle) for the KMS. The server certificate allows the external KMS to authenticate itself to StorageGRID.

Steps

1. From **Step 2 (Upload server certificate)**, browse to the location of the saved server certificate or certificate bundle.
2. Upload the certificate file.

The server certificate metadata appears.



If you uploaded a certificate bundle, the metadata for each certificate appears on its own tab.

3. Select **Continue**.

Step 3: Upload client certificates

In Step 3 (Upload client certificates) of the Add a Key Management Server wizard, you upload the client certificate and the client certificate private key. The client certificate allows StorageGRID to authenticate itself to the KMS.

Steps

1. From **Step 3 (Upload client certificates)**, browse to the location of the client certificate.
2. Upload the client certificate file.

The client certificate metadata appears.

3. Browse to the location of the private key for the client certificate.
4. Upload the private key file.
5. Select **Test and save**.

If a key doesn't exist, you are prompted to have StorageGRID create one.

The connections between the key management server and the appliance nodes are tested. If all connections are valid and the correct key is found on the KMS, the new key management server is added to the table on the Key Management Server page.



Immediately after you add a KMS, the certificate status on the Key Management Server page appears as Unknown. It might take StorageGRID as long as 30 minutes to get the actual status of each certificate. You must refresh your web browser to see the current status.

6. If an error message appears when you select **Test and save**, review the message details and then select **OK**.

For example, you might receive a 422: Unprocessable Entity error if a connection test failed.

7. If you need to save the current configuration without testing the external connection, select **Force save**.



Selecting **Force save** saves the KMS configuration, but it does not test the external connection from each appliance to that KMS. If there is an issue with the configuration, you might not be able to reboot appliance nodes that have node encryption enabled at the affected site. You might lose access to your data until the issues are resolved.

8. Review the confirmation warning, and select **OK** if you are sure you want to force save the configuration.

The KMS configuration is saved but the connection to the KMS is not tested.

Manage a KMS

Managing a key management server (KMS) involves viewing or editing details, managing certificates, viewing encrypted nodes, and removing a KMS when it is no longer needed.

Before you begin

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

View KMS details

You can view information about each key management server (KMS) in your StorageGRID system, including key details and the current status of the server and client certificates.

Steps

1. Select **CONFIGURATION > Security > Key management server**.

The Key management server page appears and shows the following information:

- The Configuration details tab lists any key management servers that are configured.
 - The Encrypted nodes tab lists any nodes that have node encryption enabled.
2. To view the details for a specific KMS and perform operations on that KMS, select the name of the KMS. The details page for the KMS lists the following information:

Field	Description
Manages keys for	The StorageGRID site associated with the KMS. This field displays the name of a specific StorageGRID site or Sites not managed by another KMS (default KMS).
Hostname	The fully qualified domain name or IP address of the KMS. If there is a cluster of two key management servers, the fully qualified domain name or IP address of both servers are listed. If there are more than two key management servers in a cluster, the fully qualified domain name or IP address of the first KMS is listed along with the number of additional key management servers in the cluster. For example: 10.10.10.10 and 10.10.10.11 or 10.10.10.10 and 2 others. To view all hostnames in a cluster, select a KMS and select Edit or Actions > Edit.

3. Select a tab on the KMS details page to view the following information:

Tab	Field	Description
Key details	Key name	The key alias for the StorageGRID client in the KMS.
	Key UID	The unique identifier of the latest version of the key.
	Last modified	The date and time of the latest version of the key.
Server certificate	Metadata	The metadata for the certificate, such as serial number, expiration date and time, and the certificate PEM.
	Certificate PEM	The contents of the PEM (privacy enhanced mail) file for the certificate.
Client certificate	Metadata	The metadata for the certificate, such as serial number, expiration date and time, and the certificate PEM.
	Certificate PEM	The contents of the PEM (privacy enhanced mail) file for the certificate.

- As often as required by your organization's security practices, select **Rotate key**, or use the KMS software, to create a new version of the key.

When key rotation is successful, the Key UID and Last modified fields are updated.



If you rotate the encryption key using the KMS software, rotate it from the last used version of the key to a new version of the same key. Don't rotate to an entirely different key.

Never attempt to rotate a key by changing the key name (alias) for the KMS. StorageGRID requires all previously used key versions (as well as any future ones) to be accessible from the KMS with the same key alias. If you change the key alias for a configured KMS, StorageGRID might not be able to decrypt your data.

Manage certificates

Promptly address any server or client certificate issues. If possible, replace certificates before they expire.



You must address any certificate issues as soon as possible to maintain data access.

Steps

- Select **CONFIGURATION > Security > Key management server**.
- In the table, look at the value for Certificate expiration for each KMS.
- If Certificate expiration for any KMS is Unknown, wait up to 30 minutes and then refresh your web browser.
- If the Certificate expiration column indicates that a certificate has expired or is nearing expiration, select the KMS to go to the KMS details page.
 - Select **Server certificate** and verify the value for the "Expires on" field.
 - To replace the certificate, select **Edit certificate** to upload a new certificate.

c. Repeat these sub-steps and select **Client certificate** instead of Server certificate.

5. When the **KMS CA certificate expiration**, **KMS client certificate expiration**, and **KMS server certificate expiration** alerts are triggered, note the description of each alert and perform the recommended actions.

It might take StorageGRID as long as 30 minutes to get updates to the certificate expiration. Refresh your web browser to see the current values.



If you get a status of **Server certificate status is unknown**, ensure your KMS allows obtaining a server certificate without requiring a client certificate.

View encrypted nodes

You can view information about the appliance nodes in your StorageGRID system that have the **Node Encryption** setting enabled.

Steps

1. Select **CONFIGURATION > Security > Key management server**.

The Key Management Server page appears. The Configuration Details tab shows any key management servers that have been configured.

2. From the top of the page, select the **Encrypted nodes** tab.

The Encrypted nodes tab lists the appliance nodes in your StorageGRID system that have the **Node Encryption** setting enabled.

3. Review the information in the table for each appliance node.

Column	Description
Node name	The name of the appliance node.
Node type	The type of node: Storage, Admin, or Gateway.
Site	The name of the StorageGRID site where the node is installed.
KMS name	The descriptive name of the KMS used for the node. If no KMS is listed, select the Configuration details tab to add a KMS. Add a key management server (KMS)
Key UID	The unique ID of the encryption key used to encrypt and decrypt data on the appliance node. To view an entire key UID, select the text. A dash (--) indicates the key UID is unknown, possibly because of a connection issue between the appliance node and the KMS.

Column	Description
Status	The status of the connection between the KMS and the appliance node. If the node is connected, the timestamp updates every 30 minutes. It can take several minutes for the connection status to update after the KMS configuration changes. Note: Refresh your web browser to see the new values.

- If the Status column indicates a KMS issue, address the issue immediately.

During normal KMS operations, the status will be **Connected to KMS**. If a node is disconnected from the grid, the node connection state is shown (Administratively Down or Unknown).

Other status messages correspond to StorageGRID alerts with the same names:

- KMS configuration failed to load
- KMS connectivity error
- KMS encryption key name not found
- KMS encryption key rotation failed
- KMS key failed to decrypt an appliance volume
- KMS is not configured

Perform the recommended actions for these alerts.



You must address any issues immediately to ensure that your data is fully protected.

Edit a KMS

You might need to edit the configuration of a key management server, for example, if a certificate is about to expire.

Before you begin

- If you plan to update the site selected for a KMS, you have reviewed the [considerations for changing the KMS for a site](#).
- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).

Steps

- Select **CONFIGURATION > Security > Key management server**.

The Key management server page appears and shows all key management servers that have been configured.

- Select the KMS you want to edit, and select **Actions > Edit**.

You can also edit a KMS by selecting the KMS name in the table and selecting **Edit** on the KMS details page.

- Optionally, update the details in **Step 1 (KMS details)** of the Edit a Key Management Server wizard.

Field	Description
KMS name	A descriptive name to help you identify this KMS. Must be between 1 and 64 characters.
Key name	The exact key alias for the StorageGRID client in the KMS. Must be between 1 and 255 characters. You only need to edit the key name in rare cases. For example, you must edit the key name if the alias is renamed in the KMS or if all versions of the previous key have been copied to the version history of the new alias.
Manages keys for	If you are editing a site-specific KMS and you don't already have a default KMS, optionally select Sites not managed by another KMS (default KMS). This selection converts a site-specific KMS to the default KMS, which will apply to all sites that don't have a dedicated KMS and to any sites added in an expansion. Note: If you are editing a site-specific KMS, you can't select another site. If you are editing the default KMS, you can't select a specific site.
Port	The port the KMS server uses for Key Management Interoperability Protocol (KMIP) communications. Defaults to 5696, which is the KMIP standard port.
Hostname	The fully qualified domain name or IP address for the KMS. Note: The Subject Alternative Name (SAN) field of the server certificate must include the FQDN or IP address you enter here. Otherwise, StorageGRID will not be able to connect to the KMS or to all servers in a KMS cluster.

- If you are configuring a KMS cluster, select **Add another hostname** to add a hostname for each server in the cluster.
- Select **Continue**.

Step 2 (Upload server certificate) of the Edit a Key Management Server wizard appears.

- If you need to replace the server certificate, select **Browse** and upload the new file.
- Select **Continue**.

Step 3 (Upload client certificates) of the Edit a Key Management Server wizard appears.

- If you need to replace the client certificate and the client certificate private key, select **Browse** and upload the new files.
- Select **Test and save**.

The connections between the key management server and all node-encrypted appliance nodes at the affected sites are tested. If all node connections are valid and the correct key is found on the KMS, the key management server is added to the table on the Key Management Server page.

- If an error message appears, review the message details, and select **OK**.

For example, you might receive a 422: Unprocessable Entity error if the site you selected for this KMS is already managed by another KMS, or if a connection test failed.

11. If you need to save the current configuration before resolving the connection errors, select **Force save**.



Selecting **Force save** saves the KMS configuration, but it does not test the external connection from each appliance to that KMS. If there is an issue with the configuration, you might not be able to reboot appliance nodes that have node encryption enabled at the affected site. You might lose access to your data until the issues are resolved.

The KMS configuration is saved.

12. Review the confirmation warning, and select **OK** if you are sure you want to force save the configuration.

The KMS configuration is saved, but the connection to the KMS is not tested.

Remove a key management server (KMS)

You might want to remove a key management server in some cases. For example, you might want to remove a site-specific KMS if you have decommissioned the site.

Before you begin

- You have reviewed the [considerations and requirements for using a key management server](#).
- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).

About this task

You can remove a KMS in these cases:

- You can remove a site-specific KMS if the site has been decommissioned or if the site includes no appliance nodes with node encryption enabled.
- You can remove the default KMS if a site-specific KMS already exists for each site that has appliance nodes with node encryption enabled.

Steps

1. Select **CONFIGURATION > Security > Key management server**.

The Key management server page appears and shows all key management servers that have been configured.

2. Select the KMS you want to remove, and select **Actions > Remove**.

You can also remove a KMS by selecting the KMS name in the table and selecting **Remove** from the KMS details page.

3. Confirm the following is true:

- You are removing a site-specific KMS for a site that has no appliance node with node encryption enabled.
- You are removing the default KMS, but a site-specific KMS already exists for each site with node encryption.

4. Select **Yes**.

The KMS configuration is removed.

Manage proxy settings

Configure storage proxy

If you are using platform services or Cloud Storage Pools, you can configure a non-transparent proxy between Storage Nodes and the external S3 endpoints. For example, you might need a non-transparent proxy to allow platform services messages to be sent to external endpoints, such as an endpoint on the internet.



Configured storage proxy settings do not apply to Kafka platform services endpoints.

Before you begin

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

About this task

You can configure the settings for a single storage proxy.

Steps

1. Select **CONFIGURATION** > **Security** > **Proxy settings**.
2. On the **Storage** tab, select the **Enable storage proxy** checkbox.
3. Select the protocol for the storage proxy.
4. Enter the hostname or IP address of the proxy server.
5. Optionally, enter the port used to connect to the proxy server.

Leave this field blank to use the default port for the protocol: 80 for HTTP or 1080 for SOCKS5.

6. Select **Save**.

After the storage proxy is saved, new endpoints for platform services or Cloud Storage Pools can be configured and tested.



Proxy changes can take up to 10 minutes to take effect.

7. Check the settings of your proxy server to ensure that platform service-related messages from StorageGRID will not be blocked.
8. If you need to disable a storage proxy, clear the checkbox, and select **Save**.

Configure admin proxy settings

If you send AutoSupport packages using HTTP or HTTPS, you can configure a non-transparent proxy server between Admin Nodes and technical support (AutoSupport).

For more information about AutoSupport, see [Configure AutoSupport](#).

Before you begin

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

About this task

You can configure the settings for a single admin proxy.

Steps

1. Select **CONFIGURATION > Security > Proxy settings**.

The Proxy Settings page appears. By default, Storage is selected in the tab menu.

2. Select the **Admin** tab.
3. Select the **Enable Admin Proxy** checkbox.
4. Enter the hostname or IP address of the proxy server.
5. Enter the port used to connect to the proxy server.
6. Optionally, enter a username and password for the proxy server.

Leave these fields blank if your proxy server does not require a username or a password.

7. Select one of the following:
 - If you want to secure the connection to the admin proxy, select **Verify proxy certificate**. Upload a CA bundle to verify the authenticity of SSL certificates presented by the admin proxy server.



AutoSupport on Demand, E-Series AutoSupport through StorageGRID, and Update Path determination on the StorageGRID Upgrade page will not work if a proxy certificate is verified.

After you upload the CA bundle, its metadata appears.

- If you don't want to validate certificates when communicating with the admin proxy server, select **Do not verify proxy certificate**.
8. Select **Save**.

After the admin proxy is saved, the proxy server between Admin Nodes and technical support is configured.



Proxy changes can take up to 10 minutes to take effect.

9. If you need to disable the admin proxy, clear the **Enable Admin Proxy** checkbox, and then select **Save**.

Control firewalls

Control access at external firewall

You can open or close specific ports at the external firewall.

You can control access to the user interfaces and APIs on StorageGRID Admin Nodes by opening or closing specific ports at the external firewall. For example, you might want to prevent tenants from being able to connect to the Grid Manager at the firewall, in addition to using other methods to control system access.

If you want to configure the StorageGRID internal firewall, see [Configure internal firewall](#).

Port	Description	If port is open...
443	Default HTTPS port for Admin Nodes	Web browsers and management API clients can access the Grid Manager, the Grid Management API, the Tenant Manager, and the Tenant Management API. Note: Port 443 is also used for some internal traffic.
8443	Restricted Grid Manager port on Admin Nodes	• Web browsers and management API clients can access the Grid Manager and the Grid Management API using HTTPS. • Web browsers and management API clients can't access the Tenant Manager or the Tenant Management API. • Requests for internal content will be rejected.
9443	Restricted Tenant Manager port on Admin Nodes	• Web browsers and management API clients can access the Tenant Manager and the Tenant Management API using HTTPS. • Web browsers and management API clients can't access the Grid Manager or the Grid Management API. • Requests for internal content will be rejected.



Single sign-on (SSO) is not available on the restricted Grid Manager or Tenant Manager ports. You must use the default HTTPS port (443) if you want users to authenticate with single sign-on.

Related information

- [Sign in to the Grid Manager](#)
- [Create tenant account](#)
- [External communications](#)

Manage internal firewall controls

StorageGRID includes an internal firewall on each node that enhances the security of your grid by enabling you to control network access to the node. Use the firewall to prevent network access on all ports except those necessary for your specific grid deployment. The configuration changes you make on the Firewall control page are deployed to each node.

Use the three tabs on the Firewall control page to customize the access you need for your grid.

- **Privileged address list:** Use this tab to allow selected access to closed ports. You can add IP addresses or subnets in CIDR notation that can access ports closed using the Manage external access tab.
- **Manage external access:** Use this tab to close ports that are open by default, or reopen ports previously

closed.

- **Untrusted Client Network:** Use this tab to specify whether a node trusts inbound traffic from the Client Network.

The settings on this tab override the settings in the Manage external access tab.

- A node with an untrusted Client Network will accept only connections on load balancer endpoint ports configured on that node (global, node interface and node type bound endpoints).
- Load balancer endpoint ports *are the only open ports* on untrusted Client Networks, regardless of the settings on the Manage external networks tab.
- When trusted, all ports opened under the Manage external access tab are accessible, as well as any load balancer endpoints opened on the Client Network.



The settings you make on one tab can affect the access changes you make on another tab. Be sure to check the settings on all tabs to ensure your network behaves in the way you expect.

To configure internal firewall controls, see [Configure firewall controls](#).

For more information about external firewalls and network security, see [Control access at external firewall](#).

Privileged address list and Manage external access tabs

The Privileged address list tab enables you to register one or more IP addresses that are granted access to grid ports that are closed. The Manage external access tab enables you to close external access to selected external ports or all open external ports (external ports are ports that are accessible by non-grid nodes by default). These two tabs often can be used together to customize the exact network access you need to allow for your grid.



Privileged IP addresses don't have internal grid port access by default.

Example 1: Use a jump host for maintenance tasks

Suppose you want to use a jump host (a security hardened host) for network administration. You could use these general steps:

1. Use the Privileged address list tab to add the IP address of the jump host.
2. Use the Manage external access tab to block all ports.



Add the privileged IP address before you block ports 443 and 8443. Any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list.

After you save your configuration, all external ports on the Admin Node in your grid will be blocked for all hosts except the jump host. You can then use the jump host to perform maintenance tasks on your grid more securely.

Example 2: Lock down sensitive ports

Suppose you want to lock down sensitive ports and the service on that port (for example, SSH on port 22). You could use the following general steps:

1. Use the Privileged address list tab to grant access only to the hosts that need access to the service.
2. Use the Manage external access tab to block all ports.



Add the privileged IP address before you block access to any ports assigned to access Grid Manager and Tenant manager (preset ports are 443 and 8443). Any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list.

After you save your configuration, port 22 and SSH service will be available to hosts on the privileged address list. All other hosts will be denied access to the service no matter what interface the request comes from.

Example 3: Disable access to unused services

At a network level, you could disable some services that you don't intend to use. For example, to block HTTP S3 client traffic, you would use the toggle on the Manage external access tab to block port 18084.

Untrusted Client Networks tab

If you are using a Client Network, you can help secure StorageGRID from hostile attacks by accepting inbound client traffic only on explicitly configured endpoints.

By default, the Client Network on each grid node is *trusted*. That is, by default, StorageGRID trusts inbound connections to each grid node on all [available external ports](#).

You can reduce the threat of hostile attacks on your StorageGRID system by specifying that the Client Network on each node be *untrusted*. If a node's Client Network is untrusted, the node only accepts inbound connections on ports explicitly configured as load balancer endpoints. See [Configure load balancer endpoints](#) and [Configure firewall controls](#).

Example 1: Gateway Node only accepts HTTPS S3 requests

Suppose you want a Gateway Node to refuse all inbound traffic on the Client Network except for HTTPS S3 requests. You would perform these general steps:

1. From the [Load balancer endpoints](#) page, configure a load balancer endpoint for S3 over HTTPS on port 443.
2. From the Firewall control page, select Untrusted to specify that the Client Network on the Gateway Node is untrusted.

After you save your configuration, all inbound traffic on the Gateway Node's Client Network is dropped except for HTTPS S3 requests on port 443 and ICMP echo (ping) requests.

Example 2: Storage Node sends S3 platform services requests

Suppose you want to enable outbound S3 platform services traffic from a Storage Node, but you want to prevent any inbound connections to that Storage Node on the Client Network. You would perform this general step:

- From the Untrusted Client Networks tab of the Firewall control page, indicate that the Client Network on the Storage Node is untrusted.

After you save your configuration, the Storage Node no longer accepts any incoming traffic on the Client Network, but it continues to allow outbound requests to configured platform services destinations.

Example 3: Limiting access to Grid Manager to a subnet

Suppose you want to allow Grid Manager access only on a specific subnet. You would perform the following steps:

1. Attach the Client Network of your Admin Nodes to the subnet.
2. Use the Untrusted Client Network tab to configure the Client Network as untrusted.
3. When you create a management interface load balancer endpoint, enter port and select the management interface that the port will access.
4. Select **Yes** for Untrusted Client Network.
5. Use the Manage external access tab to block all external ports (with or without privileged IP addresses set for hosts outside that subnet).

After you save your configuration, only hosts on the subnet you specified can access the Grid Manager. All other hosts are blocked.

Configure internal firewall

You can configure the StorageGRID firewall to control network access to specific ports on your StorageGRID nodes.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- You have reviewed the information in [Manage firewall controls](#) and [Networking guidelines](#).
- If you want an Admin Node or Gateway Node to accept inbound traffic only on explicitly configured endpoints, you have defined the load balancer endpoints.



When changing the configuration of the Client Network, existing client connections might fail if load balancer endpoints have not been configured.

About this task

StorageGRID includes an internal firewall on each node that enables you to open or close some of the ports on the nodes of your grid. You can use the Firewall control tabs to open or close ports that are open by default on the Grid Network, Admin Network, and Client Network. You can also create a list of privileged IP addresses that can access grid ports that are closed. If you are using a Client Network, you can specify whether a node trusts inbound traffic from the Client Network, and you can configure the access of specific ports on the Client Network.

Limiting the number of ports open to IP addresses outside of your grid to only those that are absolutely necessary enhances the security of your grid. You use the settings on each of the three Firewall control tabs to ensure only the needed ports are open.

For more information about using firewall controls, including examples, see [Manage firewall controls](#).

For more information about external firewalls and network security, see [Control access at external firewall](#).

Access firewall controls

Steps

1. Select **CONFIGURATION > Security > Firewall control**.

The three tabs on this page are described in [Manage firewall controls](#).

2. Select any tab to configure the firewall controls.

You can use these tabs in any order. The configurations you set on one tab don't limit what you can do on the other tabs; however, configuration changes you make on one tab might change the behavior of ports configured on other tabs.

Privileged address list

You use the Privileged address list tab to grant hosts access to ports that are closed by default or closed by settings on the Manage external access tab.

Privileged IP addresses and subnets don't have internal grid access by default. Also, load balancer endpoints and additional ports opened in the Privileged address list tab are accessible even if blocked in the Manage external access tab.



Settings on the Privileged address list tab can't override settings on the Untrusted Client Network tab.

Steps

1. On the Privileged address list tab, enter the address or IP subnet you want to grant access to closed ports.
2. Optionally, select **Add another IP address or subnet in CIDR notation** to add additional privileged clients.



Add as few addresses as possible to the privileged list.

3. Optionally, select **Allow privileged IP addresses to access StorageGRID internal ports**. See [StorageGRID internal ports](#).



This option removes some protections for internal services. Leave it disabled if possible.

4. Select **Save**.

Manage external access

When a port is closed in the Manage external access tab, the port can't be accessed by any non-grid IP address unless you add the IP address to the privileged address list. You can only close ports that are open by default, and you can only open ports that you have closed.



Settings on the Manage external access tab can't override settings on the Untrusted Client Network tab. For example, if a node is untrusted, port SSH/22 is blocked on the Client Network even if it is open on the Manage external access tab. Settings on the Untrusted Client Network tab override closed ports (such as 443, 8443, 9443) on the Client Network.

Steps

1. Select **Manage external access**.
The tab displays a table with all of the external ports (ports that are accessible by non-grid nodes by default) for the nodes in your grid.

2. Configure the ports you want open and closed using the following options:

- Use the toggle beside each port to open or close the selected port.
- Select **Open all displayed ports** to open all ports listed in the table.
- Select **Close all displayed ports** to close all ports listed in the table.



If you close Grid Manager ports 443 or 8443, any users currently connected on a blocked port, including you, will lose access to Grid Manager unless their IP address has been added to the Privileged address list.



Use the scroll bar on the right side of the table to be sure you have viewed all available ports. Use the search field to find the settings for any external port by entering a port number. You can enter a partial port number. For example, if you enter a **2**, all ports that have the string "2" as part of their name are displayed.

3. Select **Save**

Untrusted Client Network

If the Client Network for a node is untrusted, the node only accepts inbound traffic on ports configured as load balancer endpoints and, optionally, additional ports you select on this tab. You can also use this tab to specify the default setting for new nodes added in an expansion.



Existing client connections might fail if load balancer endpoints have not been configured.

The configuration changes you make on the **Untrusted Client Network** tab override the settings on the **Manage external access** tab.

Steps

1. Select **Untrusted Client Network**.
2. In the Set New Node Default section, specify what the default setting should be when new nodes are added to the grid in an expansion procedure.
 - **Trusted** (default): When a node is added in an expansion, its Client Network is trusted.
 - **Untrusted**: When a node is added in an expansion, its Client Network is untrusted.

As required, you can return to this tab to change the setting for a specific new node.



This setting does not affect the existing nodes in your StorageGRID system.

3. Use the following options to select the nodes that should allow client connections only on explicitly configured load balancer endpoints or additional selected ports:
 - Select **Untrust on displayed nodes** to add all nodes displayed in the table to the Untrusted Client Network list.
 - Select **Trust on displayed nodes** to remove all nodes displayed in the table from the Untrusted Client Network list.
 - Use the toggle beside each node to set the Client Network as Trusted or Untrusted for the selected node.

For example, you could select **Untrust on displayed nodes** to add all nodes to the Untrusted Client

Network list and then use the toggle besides an individual node to add that single node to the Trusted Client Network list.



Use the scroll bar on the right side of the table to be sure you have viewed all available nodes. Use the search field to find the settings for any node by entering the node name. You can enter a partial name. For example, if you enter a **GW**, all nodes that have the string "GW" as part of their name are displayed.

4. Select **Save**.

The new firewall settings are immediately applied and enforced. Existing client connections might fail if load balancer endpoints have not been configured.

Manage tenants

What are tenant accounts?

A tenant account allows you to use the Simple Storage Service (S3) REST API to store and retrieve objects in a StorageGRID system.



Swift details have been removed from this version of the doc site. See [StorageGRID 11.8: Manage tenants](#).

As a grid administrator, you create and manage the tenant accounts that S3 clients use to store and retrieve objects.

Each tenant account has federated or local groups, users, S3 buckets, and objects.

Tenant accounts can be used to segregate stored objects by different entities. For example, multiple tenant accounts can be used for either of these use cases:

- **Enterprise use case:** If you are administering a StorageGRID system in an enterprise application, you might want to segregate the grid's object storage by the different departments in your organization. In this case, you could create tenant accounts for the Marketing department, the Customer Support department, the Human Resources department, and so on.



If you use the S3 client protocol, you can use S3 buckets and bucket policies to segregate objects between the departments in an enterprise. You don't need to use tenant accounts. See instructions for implementing [S3 buckets and bucket policies](#) for more information.

- **Service provider use case:** If you are administering a StorageGRID system as a service provider, you can segregate the grid's object storage by the different entities that will lease the storage on your grid. In this case, you would create tenant accounts for Company A, Company B, Company C, and so on.

For more information, see [Use a tenant account](#).

How do I create a tenant account?

Use the Grid manager to create a tenant account. When you create a tenant account, you specify the following information:

- Basic information including the tenant name, client type (S3) and optional storage quota.

- Permissions for the tenant account, such as whether the tenant account can use S3 platform services, configure its own identity source, use S3 Select, or use a grid federation connection.
- The initial root access for the tenant, based on whether the StorageGRID system uses local groups and users, identity federation, or single sign-on (SSO).

In addition, you can enable the S3 Object Lock setting for the StorageGRID system if S3 tenant accounts need to comply with regulatory requirements. When S3 Object Lock is enabled, all S3 tenant accounts can create and manage compliant buckets.

What is Tenant Manager used for?

After you create the tenant account, tenant users can sign in to the Tenant Manager to perform tasks such as the following:

- Set up identity federation (unless the identity source is shared with the grid)
- Manage groups and users
- Use grid federation for account clone and cross-grid replication
- Manage S3 access keys
- Create and manage S3 buckets
- Use S3 platform services
- Use S3 Select
- Monitor storage usage



While S3 tenant users can create and manage S3 access key and buckets with the Tenant Manager, they must use an S3 client application to ingest and manage objects. See [Use S3 REST API](#) for details.

Create a tenant account

You must create at least one tenant account to control access to the storage in your StorageGRID system.

The steps for creating a tenant account vary based on whether [identity federation](#) and [single sign-on](#) are configured and whether the Grid Manager account you use to create the tenant account belongs to an admin group with the Root access permission.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access or Tenant accounts permission](#).
- If the tenant account will use the identity source that was configured for the Grid Manager, and you want to grant Root access permission for the tenant account to a federated group, you have imported that federated group into the Grid Manager. You don't need to assign any Grid Manager permissions to this admin group. See [Manage admin groups](#).
- If you want to allow an S3 tenant to clone account data and replicate bucket objects to another grid using a grid federation connection:
 - You have [configured the grid federation connection](#).
 - The status of the connection is **Connected**.

- You have Root access permission.
- You have reviewed the considerations for [managing the permitted tenants for grid federation](#).
- If the tenant account will use the identity source that was configured for Grid Manager, you have imported the same federated group into Grid Manager on both grids.

When you create the tenant, you will select this group to have the initial Root access permission for both the source and destination tenant accounts.



If this admin group doesn't exist on both grids before you create the tenant, the tenant isn't replicated to the destination.

Access the wizard

Steps

1. Select **TENANTS**.
2. Select **Create**.

Enter details

Steps

1. Enter details for the tenant.

Field	Description
Name	A name for the tenant account. Tenant names don't need to be unique. When the tenant account is created, it receives a unique, 20-digit account ID.
Description (optional)	A description to help identify the tenant. If you are creating a tenant that will use a grid federation connection, optionally, use this field to help identify which is the source tenant and which is the destination tenant. For example, this description for a tenant created on Grid 1 will also appear for the tenant replicated to Grid 2: "This tenant was created on Grid 1."
Client type	The type of client protocol this tenant will use, either S3 or Swift. Note: Support for Swift client applications has been deprecated and will be removed in a future release.
Storage quota (optional)	If you want this tenant to have a storage quota, a numerical value for the quota and the units.

2. Select **Continue**.

Select permissions

Steps

1. Optionally, select the basic permissions you want this tenant to have.



Some of these permissions have additional requirements. For details, select the help icon for each permission.

Permission	If selected...
Allow platform services	The tenant can use S3 platform services such as CloudMirror. See Manage platform services for S3 tenant accounts .
Use own identity source	The tenant can configure and manage its own identity source for federated groups and users. This option is disabled if you have configured SSO for your StorageGRID system.
Allow S3 Select	The tenant can issue S3 SelectObjectContent API requests to filter and retrieve object data. See Manage S3 Select for tenant accounts. Important: SelectObjectContent requests can decrease load-balancer performance for all S3 clients and all tenants. Enable this feature only when required and only for trusted tenants.

2. Optionally, select the advanced permissions you want this tenant to have.

Permission	If selected...
Grid federation connection	The tenant can use a grid federation connection, which: • Causes this tenant and all tenant groups and users added to the account to be cloned from this grid (the <i>source grid</i>) to the other grid in the selected connection (the <i>destination grid</i>). • Allows this tenant to configure cross-grid replication between corresponding buckets on each grid. See Manage the permitted tenants for grid federation.
S3 Object Lock	Allow the tenant to use specific features of S3 Object Lock: • Set maximum retention period defines how long new objects added to this bucket should be retained, starting from the time they are ingested. • Allow compliance mode prevents users from overwriting or deleting protected object versions during the retention period.

3. Select **Continue**.

Define root access and create tenant

Steps

1. Define root access for the tenant account, based on whether your StorageGRID system uses identity federation, single sign-on (SSO), or both.

Option	Do this
If identity federation is not enabled	Specify the password to use when signing into the tenant as the local root user.
If identity federation is enabled	1. Select an existing federated group to have Root access permission for the tenant. 2. Optionally, specify the password to use when signing in to the tenant as the local root user.
If both identity federation and single sign-on (SSO) are enabled	Select an existing federated group to have Root access permission for the tenant. No local users can sign in.

2. Select **Create tenant**.

A success message appears, and the new tenant is listed on the Tenants page. To learn how to view tenant details and monitor tenant activity, see [Monitor tenant activity](#).



Applying tenant settings across the grid could take 15 minutes or longer based on network connectivity, node status, and Cassandra operations.

3. If you selected the **Use grid federation connection** permission for the tenant:

- a. Confirm that an identical tenant was replicated to the other grid in the connection. The tenants on both grids will have the same 20-digit account ID, name, description, quota, and permissions.



If you see the error message "Tenant created without a clone," refer to the instructions in [Troubleshoot grid federation errors](#).

- b. If you provided a local root user password when defining root access, [change the password for the local root user](#) for the replicated tenant.



A local root user can't sign in to Tenant Manager on the destination grid until the password is changed.

Sign in to tenant (optional)

As required, you can sign in to the new tenant now to complete the configuration, or you can sign in to the tenant later. The sign-in steps depend on whether you are signed in to the Grid Manager using the default port (443) or a restricted port. See [Control access at external firewall](#).

Sign in now

If you are using...	Do this...
Port 443 and you set a password for the local root user	1. Select Sign in as root. When you sign in, links appear for configuring buckets, identity federation, groups, and users. 2. Select the links to configure the tenant account. Each link opens the corresponding page in the Tenant Manager. To complete the page, see the instructions for using tenant accounts.
Port 443 and you did not set a password for the local root user	Select Sign in , and enter the credentials for a user in the Root access federated group.
A restricted port	1. Select Finish 2. Select Restricted in the Tenant table to learn more about accessing this tenant account. The URL for the Tenant Manager has this format: <code>https://FQDN_or_Admin_Node_IP:port/?accountId=20-digit-account-id/</code> ◦ <i>FQDN_or_Admin_Node_IP</i> is a fully qualified domain name or the IP address of an Admin Node ◦ <i>port</i> is the tenant-only port ◦ <i>20-digit-account-id</i> is the tenant's unique account ID

Sign in later

If you are using...	Do one of these...
Port 443	• From the Grid Manager, select TENANTS, and select Sign in to the right of the tenant name. • Enter the tenant's URL in a web browser: <code>https://FQDN_or_Admin_Node_IP/?accountId=20-digit-account-id/</code> ◦ <i>FQDN_or_Admin_Node_IP</i> is a fully qualified domain name or the IP address of an Admin Node ◦ <i>20-digit-account-id</i> is the tenant's unique account ID

If you are using...	Do one of these...
A restricted port	- From the Grid Manager, select TENANTS, and select Restricted. - Enter the tenant's URL in a web browser: <pre>https://FQDN_or_Admin_Node_IP:port/?accountId=20-digit-account-id</pre> <i>FQDN_or_Admin_Node_IP</i> is a fully qualified domain name or the IP address of an Admin Node <i>port</i> is the tenant-only restricted port <i>20-digit-account-id</i> is the tenant's unique account ID

Configure the tenant

Follow the instructions in [Use a tenant account](#) to manage tenant groups and users, S3 access keys, buckets, platform services, and account clone and cross-grid replication.

Edit tenant account

You can edit a tenant account to change the display name, storage quota, or tenant permissions.



If a tenant has the **Use grid federation connection** permission, you can edit tenant details from either grid in the connection. However, any changes you make on one grid in the connection will not be copied to the other grid. If you want to keep the tenant details exactly in sync between grids, make the same edits on both grids. See [Manage the permitted tenants for grid federation connection](#).

Before you begin

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



Applying tenant settings across the grid could take 15 minutes or longer based on network connectivity, node status, and Cassandra operations.

Steps

- Select **TENANTS**.

Tenants

View information for each tenant account. Depending on the timing of ingests, network connectivity, and node status, the usage data shown might be out of date. To view more recent values, select the tenant name.

Create

Export to CSV

Actions

Search tenants by name or ID

Displaying 5 results

☐	Name	Logical space used	Quota utilization	Quota	Object count	Sign in/Copy URL
☐	Tenant 01	2.00 GB	10%	20.00 GB	100	→ 📄
☐	Tenant 02	85.00 GB	85%	100.00 GB	500	→ 📄
☐	Tenant 03	500.00 TB	50%	1.00 PB	10,000	→ 📄
☐	Tenant 04	475.00 TB	95%	500.00 TB	50,000	→ 📄
☐	Tenant 05	5.00 GB	—	—	500	→ 📄

2. Locate the tenant account you want to edit.

Use the search box to search for a tenant by name or tenant ID.

3. Select the tenant. You can do either of the following:

- Select the checkbox for the tenant, and select **Actions > Edit**.
- Select the tenant name to display the details page, and select **Edit**.

4. Optionally, change the values for these fields:

- **Name**
- **Description**
- **Storage quota**

5. Select **Continue**.

6. Select or clear the permissions for the tenant account.

- If you disable **Platform services** for a tenant who is already using them, the services that they have configured for their S3 buckets will stop working. No error message is sent to the tenant. For example, if the tenant has configured CloudMirror replication for an S3 bucket, they can still store objects in the bucket, but copies of those objects will no longer be made in the external S3 bucket that they have configured as an endpoint. See [Manage platform services for S3 tenant accounts](#).
- Change the setting of **Use own identity source** to determine whether the tenant account will use its own identity source or the identity source that was configured for the Grid Manager.

If **Use own identity source** is:

- Disabled and selected, the tenant has already enabled its own identity source. A tenant must disable its identity source before it can use the identity source that was configured for the Grid Manager.
- Disabled and not selected, SSO is enabled for the StorageGRID system. The tenant must use the identity source that was configured for the Grid Manager.
- Select or clear the **Allow S3 Select** permission as needed. See [Manage S3 Select for tenant accounts](#).

- To remove the **Use grid federation connection** permission:
 - a. Select the **Grid federation** tab.
 - b. Select **Remove permission**.
- To add the **Use grid federation connection** permission:
 - a. Select the **Grid federation** tab.
 - b. Select the **Use grid federation connection** checkbox.
 - c. Optionally, select **Clone existing local users and groups** to clone them to the remote grid. If you want, you can stop the cloning in progress or retry cloning if some local users or groups failed to be cloned after the last clone operation was completed.
- To set a maximum retention period or allow compliance mode:



S3 Object Lock must be enabled on the grid before you can use these settings.

- a. Select the **S3 Object Lock** tab.
- b. For **Set maximum retention period**, enter a value and select the time period from the pull-down.
- c. For **Allow compliance mode**, select the checkbox.

Change password for tenant's local root user

You might need to change the password for a tenant's local root user if the root user is locked out of the account.

Before you begin

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

About this task

If single sign-on (SSO) is enabled for your StorageGRID system, the local root user can't sign in to the tenant account. To perform root user tasks, users must belong to a federated group that has the Root access permission for the tenant.

Steps

1. Select **TENANTS**.

Tenants

View information for each tenant account. Depending on the timing of ingests, network connectivity, and node status, the usage data shown might be out of date. To view more recent values, select the tenant name.

[Create](#)
[Export to CSV](#)
[Actions](#)

Displaying 5 results

☐	Name	Logical space used	Quota utilization	Quota	Object count	Sign in/Copy URL
☐	Tenant 01	2.00 GB	10%	20.00 GB	100	→ 📄
☐	Tenant 02	85.00 GB	85%	100.00 GB	500	→ 📄
☐	Tenant 03	500.00 TB	50%	1.00 PB	10,000	→ 📄
☐	Tenant 04	475.00 TB	95%	500.00 TB	50,000	→ 📄
☐	Tenant 05	5.00 GB	—	—	500	→ 📄

- Select the tenant account. You can do either of the following:
 - Select the checkbox for the tenant, and select **Actions** > **Change root password**.
 - Select the tenant's name to display the details page, and select **Actions** > **Change root password**.
- Enter the new password for the tenant account.
- Select **Save**.

Delete tenant account

You can delete a tenant account if you want to permanently remove the tenant's access to the system.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- You have removed all S3 buckets and objects associated with the tenant account.
- If the tenant is permitted to use a grid federation connection, you have reviewed the considerations for [deleting a tenant with the Use grid federation connection permission](#).

Steps

- Select **TENANTS**.
- Locate the tenant account or accounts you want to delete.

Use the search box to search for a tenant by name or tenant ID.
- To delete multiple tenants, select the checkboxes and select **Actions** > **Delete**.
- To delete a single tenant, do either of the following:
 - Select the checkbox, and select **Actions** > **Delete**.
 - Select the tenant name to display the details page, and then select **Actions** > **Delete**.

5. Select **Yes**.

Manage platform services

What are platform services?

Platform services include CloudMirror replication, event notifications, and the search integration service.

If you enable platform services for S3 tenant accounts, you must configure your grid so that tenants can access the external resources necessary to use these services.

CloudMirror replication

The StorageGRID CloudMirror replication service is used to mirror specific objects from a StorageGRID bucket to a specified external destination.

For example, you might use CloudMirror replication to mirror specific customer records into Amazon S3 and then leverage AWS services to perform analytics on your data.



CloudMirror replication has some important similarities and differences with the cross-grid replication feature. To learn more, see [Compare cross-grid replication and CloudMirror replication](#).



CloudMirror replication is not supported if the source bucket has S3 Object Lock enabled.

Notifications

Per-bucket event notifications are used to send notifications about specific actions performed on objects to a specified external Kafka cluster or Amazon Simple Notification Service.

For example, you could configure alerts to be sent to administrators about each object added to a bucket, where the objects represent log files associated with a critical system event.



Although event notification can be configured on a bucket with S3 Object Lock enabled, the S3 Object Lock metadata (including Retain Until Date and Legal Hold status) of the objects will not be included in the notification messages.

Search integration service

The search integration service is used to send S3 object metadata to a specified Elasticsearch index where the metadata can be searched or analyzed using the external service.

For example, you could configure your buckets to send S3 object metadata to a remote Elasticsearch service. You could then use Elasticsearch to perform searches across buckets, and perform sophisticated analyses of patterns present in your object metadata.



Although Elasticsearch integration can be configured on a bucket with S3 Object Lock enabled, the S3 Object Lock metadata (including Retain Until Date and Legal Hold status) of the objects will not be included in the notification messages.

Platform services give tenants the ability to use external storage resources, notification services, and search or

analysis services with their data. Because the target location for platform services is typically external to your StorageGRID deployment, you must decide if you want to permit tenants to use these services. If you do, you must enable the use of platform services when you create or edit tenant accounts. You must also configure your network such that the platform services messages that tenants generate can reach their destinations.

Recommendations for using platform services

Before using platform services, be aware of the following recommendations:

- If an S3 bucket in the StorageGRID system has both versioning and CloudMirror replication enabled, you should also enable S3 bucket versioning for the destination endpoint. This allows CloudMirror replication to generate similar object versions on the endpoint.
- You should not use more than 100 active tenants with S3 requests requiring CloudMirror replication, notifications, and search integration. Having more than 100 active tenants can result in slower S3 client performance.
- Requests to an endpoint that can't be completed will be queued to a maximum of 500,000 requests. This limit is equally shared among active tenants. New tenants are allowed to temporarily exceed this 500,000 limit so that newly created tenants aren't unfairly penalized.

Related information

- [Manage platform services](#)
- [Configure Storage proxy settings](#)
- [Monitor StorageGRID](#)

Network and ports for platform services

If you allow an S3 tenant to use platform services, you must configure networking for the grid to ensure that platform services messages can be delivered to their destinations.

You can enable platform services for an S3 tenant account when you create or update the tenant account. If platform services are enabled, the tenant can create endpoints that serve as a destination for CloudMirror replication, event notifications, or search integration messages from its S3 buckets. These platform services messages are sent from Storage Nodes that run the ADC service to the destination endpoints.

For example, tenants might configure the following types of destination endpoints:

- A locally-hosted Elasticsearch cluster
- A local application that supports receiving Amazon Simple Notification Service messages
- A locally-hosted Kafka cluster
- A locally-hosted S3 bucket on the same or another instance of StorageGRID
- An external endpoint, such as an endpoint on Amazon Web Services.

To ensure that platform services messages can be delivered, you must configure the network or networks containing the ADC Storage Nodes. You must ensure that the following ports can be used to send platform services messages to the destination endpoints.

By default, platform services messages are sent on the following ports:

- **80**: For endpoint URIs that begin with http (most endpoints)
- **443**: For endpoint URIs that begin with https (most endpoints)

- **9092:** For endpoint URIs that begin with http or https (Kafka endpoints only)

Tenants can specify a different port when they create or edit an endpoint.



If a StorageGRID deployment is used as the destination for CloudMirror replication, replication messages might be received on a port other than 80 or 443. Ensure that the port being used for S3 by the destination StorageGRID deployment is specified in the endpoint.

If you use a non-transparent proxy server, you must also [configure storage proxy settings](#) to allow messages to be sent to external endpoints, such as an endpoint on the internet.


Related information

[Use a tenant account](#)


Per-site delivery of platform services messages

All platform services operations are performed on a per-site basis.

That is, if a tenant uses a client to perform an S3 API Create operation on an object by connecting to a Gateway Node at Data Center Site 1, the notification about that action is triggered and sent from Data Center Site 1.



If the client subsequently performs an S3 API Delete operation on that same object from Data Center Site 2, the notification about the delete action is triggered and sent from Data Center Site 2.



Make sure that the networking at each site is configured such that platform services messages can be delivered to their destinations.

Troubleshoot platform services

The endpoints used in platform services are created and maintained by tenant users in the Tenant Manager; however, if a tenant has issues configuring or using platform services, you might be able to use the Grid Manager to help resolve the issue.

Issues with new endpoints

Before a tenant can use platform services, they must create one or more endpoints using the Tenant Manager. Each endpoint represents an external destination for one platform service, such as a StorageGRID S3 bucket, an Amazon Web Services bucket, an Amazon Simple Notification Service topic, a Kafka topic, or an Elasticsearch cluster hosted locally or on AWS. Each endpoint includes both the location of the external resource and the credentials needed to access that resource.

When a tenant creates an endpoint, the StorageGRID system validates that the endpoint exists and that it can

be reached using the credentials that were specified. The connection to the endpoint is validated from one node at each site.

If endpoint validation fails, an error message explains why endpoint validation failed. The tenant user should resolve the issue, then try creating the endpoint again.



Endpoint creation will fail if platform services aren't enabled for the tenant account.

Issues with existing endpoints

If an error occurs when StorageGRID tries to reach an existing endpoint, a message is displayed on the dashboard in the Tenant Manager.

One or more endpoints have experienced an error and might not be functioning properly. Go to the [Endpoints](#) page to view the error details. The last error occurred 2 hours ago.

Tenant users can go to the Endpoints page to review the most recent error message for each endpoint and to determine how long ago the error occurred. The **Last error** column displays the most recent error message for each endpoint and indicates how long ago the error occurred. Errors that include the icon occurred within the past 7 days.

Platform services endpoints

A platform services endpoint stores the information StorageGRID needs to use an external resource as a target for a platform service (CloudMirror replication, notifications, or search integration). You must configure an endpoint for each platform service you plan to use.

One or more endpoints have experienced an error. Select the endpoint for more details about the error. Meanwhile, the platform service request will be retried automatically.

5 endpoints Create endpoint

Delete endpoint

☐	Display name	Last error	Type	URI	URN
☐	my-endpoint-2	2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es::mydomain/sveloso/_doc
☐	my-endpoint-3	3 days ago	Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1
☐	my-endpoint-5	12 days ago	Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example3
☐	my-endpoint-4		Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example2
☐	my-endpoint-1		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket1



Some error messages in the **Last error** column might include a logID in parentheses. A grid administrator or technical support can use this ID to locate more detailed information about the error in the bycast.log.

Issues related to proxy servers

If you have configured a [storage proxy](#) between Storage Nodes and platform service endpoints, errors might occur if your proxy service does not allow messages from StorageGRID. To resolve these issues, check the settings of your proxy server to ensure that platform service-related messages aren't blocked.

Determine if an error has occurred

If any endpoint errors have occurred within the past 7 days, the dashboard in the Tenant Manager displays an alert message. You can go the Endpoints page to see more details about the error.

Client operations fail

Some platform services issues might cause client operations on the S3 bucket to fail. For example, S3 client operations will fail if the internal Replicated State Machine (RSM) service stops, or if there are too many platform services messages queued for delivery.

To check the status of services:

1. Select **SUPPORT > Tools > Grid topology**.
2. Select **site > Storage Node > SSM > Services**.

Recoverable and unrecoverable endpoint errors

After endpoints have been created, platform service request errors can occur for various reasons. Some errors are recoverable with user intervention. For example, recoverable errors might occur for the following reasons:

- The user's credentials have been deleted or have expired.
- The destination bucket does not exist.
- The notification can't be delivered.

If StorageGRID encounters a recoverable error, the platform service request will be retried until it succeeds.

Other errors are unrecoverable. For example, an unrecoverable error occurs if the endpoint is deleted.

If StorageGRID encounters an unrecoverable endpoint error:

- In the Grid Manager, go to **Support > Tools > Metrics > Grafana > Platform Services Overview** to view error details.
- In the Tenant Manager, go to **STORAGE (S3) > Platform Services Endpoints** to view the error details.
- Check the `/var/local/log/bycast-err.log` for related errors. Storage Nodes that have the ADC service contain this log file.

Platform services messages can't be delivered

If the destination encounters an issue that prevents it from accepting platform services messages, the client operation on the bucket succeeds, but the platform services message is not delivered. For example, this error might happen if credentials are updated on the destination such that StorageGRID can no longer authenticate to the destination service.

Check for related alerts.

Slower performance for platform service requests

StorageGRID software might throttle incoming S3 requests for a bucket if the rate at which the requests are being sent exceeds the rate at which the destination endpoint can receive the requests. Throttling only occurs when there is a backlog of requests waiting to be sent to the destination endpoint.



The only visible effect is that the incoming S3 requests will take longer to execute. If you start to detect significantly slower performance, you should reduce the ingest rate or use an endpoint with higher capacity. If the backlog of requests continues to grow, client S3 operations (such as PUT requests) will eventually fail.

CloudMirror requests are more likely to be affected by the performance of the destination endpoint because these requests typically involve more data transfer than search integration or event notification requests.

Platform service requests fail

To view the request failure rate for platform services:

1. Select **NODES**.
2. Select **site > Platform Services**.
3. View the Request error rate chart.

[1 hour](#) [1 day](#) [1 week](#) [1 month](#) [Custom](#)

Platform services unavailable alert

The **Platform services unavailable** alert indicates that no platform service operations can be performed at a site because too few Storage Nodes with the RSM service are running or available.

The RSM service ensures platform service requests are sent to their respective endpoints.

To resolve this alert, determine which Storage Nodes at the site include the RSM service. (The RSM service is present on Storage Nodes that also include the ADC service.) Then, ensure that a simple majority of those Storage Nodes are running and available.



If more than one Storage Node that contains the RSM service fails at a site, you lose any pending platform service requests for that site.

Additional troubleshooting guidance for platform services endpoints

For additional information see [Use a tenant account > Troubleshoot platform services endpoints](#).

Related information

[Troubleshoot StorageGRID system](#)

Manage S3 Select for tenant accounts

You can allow certain S3 tenants to use S3 Select to issue `SelectObjectContent` requests on individual objects.

S3 Select provides an efficient way to search through large amounts of data without having to deploy a database and associated resources to enable searches. It also reduces the cost and latency of retrieving data.

What is S3 Select?

S3 Select allows S3 clients to use `SelectObjectContent` requests to filter and retrieve only the data needed from an object. The StorageGRID implementation of S3 Select includes a subset of S3 Select commands and features.

Considerations and requirements for using S3 Select

Grid administration requirements

The grid administrator must grant tenants S3 Select ability. Select **Allow S3 Select** when [creating a tenant](#) or [editing a tenant](#).

Object format requirements

The object you want to query must be in one of the following formats:

- **CSV.** Can be used as is or compressed into GZIP or BZIP2 archives.
- **Parquet.** Additional requirements for Parquet objects:
 - S3 Select supports only columnar compression using GZIP or Snappy. S3 Select doesn't support whole-object compression for Parquet objects.
 - S3 Select doesn't support Parquet output. You must specify the output format as CSV or JSON.
 - The maximum uncompressed row group size is 512 MB.
 - You must use the data types specified in the object's schema.
 - You can't use INTERVAL, JSON, LIST, TIME, or UUID logical types.

Endpoint requirements

The `SelectObjectContent` request must be sent to a [StorageGRID load balancer endpoint](#).

The Admin and Gateway Nodes used by the endpoint must be one of the following:

- A services appliance node
- A VMware-based software node
- A bare metal node running a kernel with cgroup v2 enabled

General considerations

Queries can't be sent directly to Storage Nodes.



SelectObjectContent requests can decrease load-balancer performance for all S3 clients and all tenants. Enable this feature only when required and only for trusted tenants.

See the [instructions for using S3 Select](#).

To view [Grafana charts](#) for S3 Select operations over time, select **SUPPORT > Tools > Metrics** in the Grid Manager.

Configure client connections

Configure S3 client connections

As a grid administrator, you manage the configuration options that control how S3 client applications connect to your StorageGRID system to store and retrieve data.



Swift details have been removed from this version of the doc site. See [StorageGRID 11.8: Configure S3 and Swift client connections](#).

Configuration tasks

1. Perform prerequisite tasks in StorageGRID, based on how the client application will connect to StorageGRID.

Required tasks

You must obtain:

- IP addresses
- Domain names
- SSL certificate

Optional tasks

Optionally, configure:

- Identity federation
- SSO

2. Use StorageGRID to obtain the values the application needs to connect to the grid. You can either use the S3 setup wizard or configure each StorageGRID entity manually. +

Use S3 setup wizard

Follow the steps in the S3 setup wizard.

Configure manually

- a. Create high availability group
- b. Create load balancer endpoint
- c. Create tenant account
- d. Create bucket and access keys
- e. Configure ILM rule and policy

3. Use the S3 application to complete the connection to StorageGRID. Create DNS entries to associate IP addresses to any domain names you plan to use.

As needed, perform additional application setup.

4. Perform ongoing tasks in the application and in StorageGRID to manage and monitor object storage over time.

Information needed to attach StorageGRID to a client application

Before you can attach StorageGRID to an S3 client application, you must perform configuration steps in StorageGRID and obtain certain value.

What values do I need?

The following table shows the values you must configure in StorageGRID and where those values are used by the S3 application and the DNS server.

Value	Where value is configured	Where value is used
Virtual IP (VIP) addresses	StorageGRID > HA group	DNS entry
Port	StorageGRID > Load balancer endpoint	Client application
SSL certificate	StorageGRID > Load balancer endpoint	Client application
Server name (FQDN)	StorageGRID > Load balancer endpoint	• Client application • DNS entry
S3 access key ID and secret access key	StorageGRID > Tenant and bucket	Client application
Bucket/Container name	StorageGRID > Tenant and bucket	Client application

How do I get these values?

Depending on your requirements, you can do either of the following to obtain the information you need:

- **Use the [S3 setup wizard](#).** The S3 setup wizard helps you to quickly configure the required values in StorageGRID and outputs one or two files that you can use when you configure the S3 application. The wizard guides you through the required steps and helps to make sure your settings conform to StorageGRID best practices.



If you are configuring an S3 application, using the S3 setup wizard is recommended unless you know you have special requirements or your implementation will require significant customization.

- **Use the [FabricPool setup wizard](#).** Similar to the S3 setup wizard, the FabricPool setup wizard helps you to quickly configure required values and outputs a file that you can use when you configure a FabricPool cloud tier in ONTAP.



If you plan to use StorageGRID as the object storage system for a FabricPool cloud tier, using the FabricPool setup wizard is recommended unless you know you have special requirements or your implementation will require significant customization.

- **Configure items manually.** If you are connecting to an S3 application and prefer not to use the S3 setup wizard, you can obtain the required values by performing the configuration manually. Follow these steps:
 1. Configure the high availability (HA) group you want to use for the S3 application. See [Configure high availability groups](#).
 2. Create the load balancer endpoint that the S3 application will use. See [Configure load balancer endpoints](#).
 3. Create the tenant account that the S3 application will use. See [Create a tenant account](#).
 4. For an S3 tenant, sign in to the tenant account, and generate an access key ID and secret access key for each user that will access the application. See [Create your own access keys](#).
 5. Create one or more S3 buckets within the tenant account. For S3, see [Create S3 bucket](#).
 6. To add specific placement instructions for the objects belonging to the new tenant or bucket/container, create a new ILM rule and activate a new ILM policy to use that rule. See [Create ILM rule](#) and [Create ILM policy](#).

Security for S3 clients

StorageGRID tenant accounts use S3 client applications to save object data to StorageGRID. You should review the security measures implemented for client applications.

Summary

The following list summarizes how security is implemented for the S3 REST API:

Connection security

TLS

Server authentication

X.509 server certificate signed by system CA or custom server certificate supplied by administrator

Client authentication

S3 account access key ID and secret access key

Client authorization

Bucket ownership and all applicable access control policies

How StorageGRID provides security for client applications

S3 client applications can connect to the Load Balancer service on Gateway Nodes or Admin Nodes or directly to Storage Nodes.

- Clients that connect to the Load Balancer service can use HTTPS or HTTP, based on how you [configure the load balancer endpoint](#).

HTTPS provides secure, TLS-encrypted communication and is recommended. You must attach a security certificate to the endpoint.

HTTP provides less secure, unencrypted communication and should only be used for non-production or test grids.

- Clients that connect to Storage Nodes can also use HTTPS or HTTP.

HTTPS is the default and is recommended.

HTTP provides less secure, unencrypted communication but can be optionally [enabled](#) for non-production or test grids.

- Communications between StorageGRID and the client are encrypted using TLS.
- Communications between the Load Balancer service and Storage Nodes within the grid are encrypted whether the load balancer endpoint is configured to accept HTTP or HTTPS connections.
- Clients must supply [HTTP authentication headers](#) to StorageGRID to perform REST API operations.

Security certificates and client applications

In all cases, client applications can make TLS connections using either a custom server certificate uploaded by the grid administrator or a certificate generated by the StorageGRID system:

- When client applications connect to the Load Balancer service, they use the certificate that was configured for the load balancer endpoint. Each load balancer endpoint has its own certificate—either a custom server certificate uploaded by the grid administrator or a certificate that the grid administrator generated in StorageGRID when configuring the endpoint.

See [Considerations for load balancing](#).

- When client applications connect directly to a Storage Node, they use either the system-generated server certificates that were generated for Storage Nodes when the StorageGRID system was installed (which are signed by the system certificate authority), or a single custom server certificate that is supplied for the grid by a grid administrator. See [add a custom S3 API certificate](#).

Clients should be configured to trust the certificate authority that signed whichever certificate they use to establish TLS connections.

Supported hashing and encryption algorithms for TLS libraries

The StorageGRID system supports a set of cipher suites that client applications can use when establishing a TLS session. To configure ciphers, go to **CONFIGURATION > Security > Security settings** and select **TLS and SSH policies**.

Supported versions of TLS

StorageGRID supports TLS 1.2 and TLS 1.3.



SSLv3 and TLS 1.1 (or earlier versions) are no longer supported.

Use S3 setup wizard

Use S3 setup wizard: Considerations and requirements

You can use the S3 setup wizard to configure StorageGRID as the object storage system for an S3 application.

When to use the S3 setup wizard

The S3 setup wizard guides you through each step of configuring StorageGRID for use with an S3 application. As part of completing the wizard, you download files that you can use to enter values into the S3 application. Use the wizard to configure your system more quickly and to make sure your settings conform to StorageGRID best practices.

If you have the [Root access permission](#), you can complete the S3 setup wizard when you start using the StorageGRID Grid Manager, or you can access and complete the wizard at any later time. Depending on your requirements, you can also configure some or all of the required items manually and then use the wizard to assemble the values that an S3 application needs.

Before using the wizard

Before using the wizard, confirm you have completed these prerequisites.

Obtain IP addresses and set up VLAN interfaces

If you will configure a high availability (HA) group, you know which nodes the S3 application will connect to and which StorageGRID network will be used. You also know which values to enter for the subnet CIDR, gateway IP address, and virtual IP (VIP) addresses.

If you plan to use a virtual LAN to segregate the traffic from the S3 application, you have already configured the VLAN interface. See [Configure VLAN interfaces](#).

Configure identity federation and SSO

If you plan to use identity federation or single sign-on (SSO) for your StorageGRID system, you have enabled these features. You also know which federated group should have root access for the tenant account that the S3 application will use. See [Use identity federation](#) and [Configure single sign-on](#).

Obtain and configure domain names

You know which fully qualified domain name (FQDN) to use for StorageGRID. Domain name server (DNS) entries will map this FQDN to the virtual IP (VIP) addresses of the HA group that you create using the wizard.

If you plan to use S3 virtual hosted-style requests, you should have [configured S3 endpoint domain names](#). Using virtual hosted-style requests is recommended.

Review load balancer and security certificate requirements

If you plan to use the StorageGRID load balancer, you have reviewed the general considerations for load balancing. You have the certificates you will upload or the values you need to generate a certificate.

If you plan to use an external (third-party) load balancer endpoint, you have the fully qualified domain name (FQDN), port, and certificate for that load balancer.

Configure any grid federation connections

If you want to allow the S3 tenant to clone account data and replicate bucket objects to another grid using a grid federation connection, confirm the following before starting the wizard:

- You have [configured the grid federation connection](#).
- The status of the connection is **Connected**.
- You have Root access permission.

Access and complete the S3 setup wizard

You can use the S3 setup wizard to configure StorageGRID for use with an S3 application. The setup wizard provides the values the application needs to access a StorageGRID bucket and to save objects.


Before you begin

- You have the [Root access permission](#).
- You have reviewed the [considerations and requirements](#) for using the wizard.

Access the wizard

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. If the **FabricPool and S3 setup wizard** banner appears on the dashboard, select the link in the banner. If the banner no longer appears, select the help icon from the header bar in the Grid Manager and select **FabricPool and S3 setup wizard**.



3. In the S3 application section of the FabricPool and S3 setup wizard page, select **Configure now**.

Step 1 of 6: Configure HA group

An HA group is a collection of nodes that each contain the StorageGRID Load Balancer service. An HA group can contain Gateway Nodes, Admin Nodes, or both.

You can use an HA group to help keep the S3 data connections available. If the active interface in the HA group fails, a backup interface can manage the workload with little impact to S3 operations.

For details about this task, see [Manage high availability groups](#).

Steps

1. If you plan to use an external load balancer, you don't need to create an HA group. Select **Skip this step** and go to [Step 2 of 6: Configure load balancer endpoint](#).
2. To use the StorageGRID load balancer, you can create a new HA group or use an existing HA group.

Create HA group

- a. To create a new HA group, select **Create HA group**.
- b. For the **Enter details** step, complete the following fields.

Field	Description
HA group name	A unique display name for this HA group.
Description (optional)	The description of this HA group.

- c. For the **Add interfaces** step, select the node interfaces you want to use in this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

You can select one or more nodes, but you can select only one interface for each node.

- d. For the **Prioritize interfaces** step, determine the Primary interface and any backup interfaces for this HA group.

Drag rows to change the values in the **Priority order** column.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

If the HA group includes more than one interface and the active interface fails, the virtual IP (VIP) addresses move to the first backup interface in the priority order. If that interface fails, the VIP addresses move to the next backup interface, and so on. When failures are resolved, the VIP addresses move back to the highest priority interface available.

- e. For the **Enter IP addresses** step, complete the following fields.

Field	Description
Subnet CIDR	The address of the VIP subnet in CIDR notation — an IPv4 address followed by a slash and the subnet length (0-32). The network address must not have any host bits set. For example, 192.16.0.0/22.
Gateway IP address (optional)	If the S3 IP addresses used to access StorageGRID aren't on the same subnet as the StorageGRID VIP addresses, enter the StorageGRID VIP local gateway IP address. The local gateway IP address must be within the VIP subnet.
Virtual IP address	Enter at least one and no more than ten VIP addresses for the active interface in the HA group. All VIP addresses must be within the VIP subnet. At least one address must be IPv4. Optionally, you can specify additional IPv4 and IPv6 addresses.

- f. Select **Create HA group** and then select **Finish** to return to the S3 setup wizard.
- g. Select **Continue** to go to the load balancer step.

Use existing HA group

- a. To use an existing HA group, select the HA group name from the **Select an HA group**.
- b. Select **Continue** to go to the load balancer step.

Step 2 of 6: Configure load balancer endpoint

StorageGRID uses a load balancer to manage the workload from client applications. Load balancing maximizes speed and connection capacity across multiple Storage Nodes.

You can use the StorageGRID Load Balancer service, which exists on all Gateway and Admin Nodes, or you can connect to an external (third-party) load balancer. Using the StorageGRID load balancer is recommended.

For details about this task, see [Considerations for load balancing](#).

To use the StorageGRID Load Balancer service, select the **StorageGRID load balancer** tab and then create or select the load balancer endpoint you want to use. To use an external load balancer, select the **External load balancer** tab and provide details about the system you have already configured.

Create endpoint

Steps

1. To create a load balancer endpoint, select **Create endpoint**.
2. For the **Enter endpoint details** step, complete the following fields.

Field	Description
Name	A descriptive name for the endpoint.
Port	The StorageGRID port you want to use for load balancing. This field defaults to 10433 for the first endpoint you create, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes, because these ports are reserved on Admin Nodes. Note: Ports used by other grid services aren't permitted. See the Network port reference.
Client type	Must be S3 .
Network protocol	Select HTTPS. Note: Communicating with StorageGRID without TLS encryption is supported but not recommended.

3. For the **Select binding mode** step, specify the binding mode. The binding mode controls how the endpoint is accessed using any IP address or using specific IP addresses and network interfaces.

Mode	Description
Global (default)	Clients can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN. Use the Global setting (default) unless you need to restrict the accessibility of this endpoint.
Virtual IPs of HA groups	Clients must use a virtual IP address (or corresponding FQDN) of an HA group to access this endpoint. Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints don't overlap.
Node interfaces	Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces to access this endpoint.

Mode	Description
Node type	Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node to access this endpoint.

4. For the Tenant access step, select one of the following:

Field	Description
Allow all tenants (default)	All tenant accounts can use this endpoint to access their buckets.
Allow selected tenants	Only the selected tenant accounts can use this endpoint to access their buckets.
Block selected tenants	The selected tenant accounts can't use this endpoint to access their buckets. All other tenants can use this endpoint.

5. For the **Attach certificate** step, select one of the following:

Field	Description
Upload certificate (recommended)	Use this option to upload a CA-signed server certificate, certificate private key, and optional CA bundle.
Generate certificate	Use this option to generate a self-signed certificate. See Configure load balancer endpoints for details of what to enter.
Use StorageGRID S3 certificate	Use this option only if you have already uploaded or generated a custom version of the StorageGRID global certificate. See Configure S3 API certificates for details.

6. Select **Finish** to return to the S3 setup wizard.
7. Select **Continue** to go to the tenant and bucket step.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

Use existing load balancer endpoint

Steps

1. To use an existing endpoint, select its name from the **Select a load balancer endpoint**.
2. Select **Continue** to go to the tenant and bucket step.

Use external load balancer

Steps

1. To use an external load balancer, complete the following fields.

Field	Description
FQDN	The fully qualified domain name (FQDN) of the external load balancer.
Port	The port number that the S3 application will use to connect to the external load balancer.
Certificate	Copy the server certificate for the external load balancer and paste it into this field.

2. Select **Continue** to go to the tenant and bucket step.

Step 3 of 6: Create tenant and bucket

A tenant is an entity that can use S3 applications to store and retrieve objects in StorageGRID. Each tenant has its own users, access keys, buckets, objects, and a specific set of capabilities.

A bucket is a container used to store a tenant's objects and object metadata. Although tenants might have many buckets, the wizard helps you to create a tenant and a bucket in the quickest and easiest way. If you need to add buckets or set options later, you can use the Tenant Manager.

For details about this task, see [Create tenant account](#) and [Create S3 bucket](#).

Steps

1. Enter a name for the tenant account.

Tenant names don't need to be unique. When the tenant account is created, it receives a unique, numeric account ID.

2. Define root access for the tenant account, based on whether your StorageGRID system uses [identity federation](#), [single sign-on \(SSO\)](#), or both.

Option	Do this
If identity federation is not enabled	Specify the password to use when signing into the tenant as the local root user.
If identity federation is enabled	1. Select an existing federated group to have Root access permission for the tenant. 2. Optionally, specify the password to use when signing in to the tenant as the local root user.
If both identity federation and single sign-on (SSO) are enabled	Select an existing federated group to have Root access permission for the tenant. No local users can sign in.

3. If you want the wizard to create the access key ID and secret access key for the root user, select **Create root user S3 access key automatically**.

Select this option if the only user for the tenant will be the root user. If other users will use this tenant, [use Tenant Manager](#) to configure keys and permissions.

4. If you want to create a bucket for this tenant now, select **Create bucket for this tenant**.



If S3 Object Lock is enabled for the grid, the bucket created in this step doesn't have S3 Object Lock enabled. If you need to use an S3 Object Lock bucket for this S3 application, don't select to create a bucket now. Instead, use Tenant Manager to [create the bucket](#) later.

- a. Enter the name of the bucket that the S3 application will use. For example, `s3-bucket`.

You can't change the bucket name after creating the bucket.

- b. Select the **Region** for this bucket.

Use the default region (`us-east-1`) unless you expect to use ILM in the future to filter objects based on the bucket's region.

5. Select **Create and continue**.

Step 4 of 6: Download data

In the download data step, you can download one or two files to save the details of what you just configured.

Steps

1. If you selected **Create root user S3 access key automatically**, do one or both of the following:
 - Select **Download access keys** to download a `.csv` file containing the tenant account name, access key ID, and secret access key.
 - Select the copy icon () to copy the access key ID and secret access key to the clipboard.
2. Select **Download configuration values** to download a `.txt` file containing the settings for the load balancer endpoint, tenant, bucket, and the root user.
3. Save this information to a secure location.



Don't close this page until you have copied both access keys. The keys will not be available after you close this page. Make sure to save this information in a secure location because it can be used to obtain data from your StorageGRID system.

4. If prompted, select the checkbox to confirm that you have downloaded or copied the keys.
5. Select **Continue** to go to the ILM rule and policy step.

Step 5 of 6: Review ILM rule and ILM policy for S3

Information lifecycle management (ILM) rules control the placement, duration, and ingest behavior of all objects in your StorageGRID system. The ILM policy included with StorageGRID makes two replicated copies of all objects. This policy is in effect until you activate at least one new policy.

Steps

1. Review the information provided on the page.
2. If you want to add specific instructions for the objects belonging to the new tenant or bucket, create a new rule and a new policy. See [Create ILM rule](#) and [Use ILM policies](#).

3. Select **I have reviewed these steps and understand what I need to do**.
4. Select the checkbox to indicate that you understand what to do next.
5. Select **Continue** to go to **Summary**.

Step 6 of 6: Review summary

Steps

1. Review the summary.
2. Make note of the details in the next steps, which describe the additional configuration that might be needed before you connect to the S3 client. For example, selecting **Sign in as root** takes you to the Tenant Manager, where you can add tenant users, create additional buckets, and update bucket settings.
3. Select **Finish**.
4. Configure the application using the file you downloaded from StorageGRID or the values you obtained manually.

Manage HA groups

What are high availability (HA) groups?

High availability (HA) groups provide highly available data connections for S3 clients and highly available connections to the Grid Manager and the Tenant Manager.

You can group the network interfaces of multiple Admin and Gateway Nodes into a high availability (HA) group. If the active interface in the HA group fails, a backup interface can manage the workload.

Each HA group provides access to the shared services on the selected nodes.

- HA groups that include Gateway Nodes, Admin Nodes, or both provide highly available data connections for S3 clients.
- HA groups that include only Admin Nodes provide highly available connections to the Grid Manager and the Tenant Manager.
- An HA group that includes only services appliances and VMware-based software nodes can provide highly available connections for [S3 tenants that use S3 Select](#).
HA groups are recommended when using S3 Select, but not required.

How do you create an HA group?

1. You select a network interface for one or more Admin Nodes or Gateway Nodes. You can use a Grid Network (eth0) interface, Client Network (eth2) interface, VLAN interface, or an access interface you have added to the node.



You can't add an interface to an HA group if it has a DHCP-assigned IP address.

2. You specify one interface to be the Primary interface. The Primary interface is the active interface unless a failure occurs.
3. You determine the priority order for any Backup interfaces.
4. You assign one to 10 virtual IP (VIP) addresses to the group. Clients applications can use any of these VIP addresses to connect to StorageGRID.

For instructions, see [Configure high availability groups](#).

What is the active interface?

During normal operation, all of the VIP addresses for the HA group are added to the Primary interface, which is the first interface in the priority order. As long as the Primary interface remains available, it is used when clients connect to any VIP address for the group. That is, during normal operation, the Primary interface is the "active" interface for the group.

Similarly, during normal operation, any lower priority interfaces for the HA group act as "backup" interfaces. These backup interfaces aren't used unless the Primary (currently active) interface becomes unavailable.

View the current HA group status of a node

To see if a node is assigned to an HA group and determine its current status, select **NODES > node**.

If the **Overview** tab includes an entry for **HA groups**, the node is assigned to the HA groups listed. The value after the group name is the current status of the node in the HA group:

- **Active:** The HA group is currently being hosted on this node.
- **Backup:** The HA group is not currently using this node; this is a backup interface.
- **Stopped:** The HA group can't be hosted on this node because the High Availability (keepalived) service has been stopped manually.
- **Fault:** The HA group can't be hosted on this node because of one or more of the following:
 - The Load Balancer (nginx-gw) service is not running on the node.
 - The node's eth0 or VIP interface is down.
 - The node is down.

In this example, the primary Admin Node has been added to two HA groups. This node is currently the active interface for the Admin clients group and a backup interface for the FabricPool clients group.

DC1-ADM1 (Primary Admin Node)

Overview
Hardware
Network
Storage
Load balancer
Tasks

Node information

Name: DC1-ADM1
Type: Primary Admin Node
ID: ce00d9c8-8a79-4742-bdef-c9c658db5315
Connection state: Connected
Software version: 11.6.0 (build 20211207.1804.614bc17)
HA groups: Admin clients (Active)
FabricPool clients (Backup)
IP addresses: 172.16.1.225 - eth0 (Grid Network)
10.224.1.225 - eth1 (Admin Network)
47.47.0.2, 47.47.1.225 - eth2 (Client Network)
Show additional IP addresses

What happens when the active interface fails?

The interface that currently hosts the VIP addresses is the active interface. If the HA group includes more than one interface and the active interface fails, the VIP addresses move to the first available backup interface in the priority order. If that interface fails, the VIP addresses move to the next available backup interface, and so on.

Failover can be triggered for any of these reasons:

- The node on which the interface is configured goes down.
- The node on which the interface is configured loses connectivity to all other nodes for at least 2 minutes.
- The active interface goes down.
- The Load Balancer service stops.
- The High Availability service stops.



Failover might not be triggered by network failures external to the node that hosts the active interface. Similarly, failover is not triggered by the services for the Grid Manager or the Tenant Manager.

The failover process generally takes only a few seconds and is fast enough that client applications should experience little impact and can rely on normal retry behaviors to continue operation.

When failure is resolved and a higher priority interface becomes available again, the VIP addresses are automatically moved to the highest priority interface that is available.

How are HA groups used?

You can use high availability (HA) groups to provide highly available connections to StorageGRID for object data and for administrative use.

- An HA group can provide highly available administrative connections to the Grid Manager or the Tenant Manager.
- An HA group can provide highly available data connections for S3 clients.
- An HA group that contains only one interface allows you to provide many VIP addresses and to explicitly set IPv6 addresses.

An HA group can provide high availability only if all nodes included in the group provide the same services. When you create an HA group, add interfaces from the types of nodes that provide the services you require.

- **Admin Nodes:** Include the Load Balancer service and enable access to the Grid Manager or the Tenant Manager.
- **Gateway Nodes:** Include the Load Balancer service.

Purpose of HA group	Add nodes of this type to the HA group
Access to Grid Manager	• Primary Admin Node (Primary) • Non-primary Admin Nodes Note: The primary Admin Node must be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node.
Access to Tenant Manager only	• Primary or non-primary Admin Nodes
S3 client access — Load Balancer service	• Admin Nodes • Gateway Nodes
S3 client access for S3 Select	• Services appliances • VMware-based software nodes Note: HA groups are recommended when using S3 Select, but not required.

Limitations of using HA groups with Grid Manager or Tenant Manager

If a Grid Manager or Tenant Manager service fails, HA group failover is not triggered.


If you are signed in to the Grid Manager or the Tenant Manager when failover occurs, you are signed out and must sign in again to resume your task.

Some maintenance procedures can't be performed when the primary Admin Node is unavailable. During failover, you can use the Grid Manager to monitor your StorageGRID system.

Configuration options for HA groups

The following diagrams provide examples of different ways you can configure HA groups. Each option has advantages and disadvantages.

In the diagrams, blue indicates the primary interface in the HA group and yellow indicates the backup interface in the HA group.



The table summarizes the benefits of each HA configuration shown in the diagram.

Configuration	Advantages	Disadvantages
Active-Backup HA	• Managed by StorageGRID with no external dependencies. • Fast failover.	• Only one node in an HA group is active. At least one node per HA group will be idle.

Configuration	Advantages	Disadvantages
DNS Round Robin	• Increased aggregate throughput. • No idle hosts.	• Slow failover, which could depend on client behavior. • Requires configuration of hardware outside of StorageGRID. • Needs a customer-implemented health check.
Active-Active HA	• Traffic is distributed across multiple HA groups. • High aggregate throughput that scales with the number of HA groups. • Fast failover.	• More complex to configure. • Requires configuration of hardware outside of StorageGRID. • Needs a customer-implemented health check.

Configure high availability groups

You can configure high availability (HA) groups to provide highly available access to the services on Admin Nodes or Gateway Nodes.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).
- If you plan to use a VLAN interface in an HA group, you have created the VLAN interface. See [Configure VLAN interfaces](#).
- If you plan to use an access interface for a node in an HA group, you have created the interface:
 - **Red Hat Enterprise Linux (before installing the node):** [Create node configuration files](#)
 - **Ubuntu or Debian (before installing the node):** [Create node configuration files](#)
 - **Linux (after installing the node):** [Linux: Add trunk or access interfaces to a node](#)
 - **VMware (after installing the node):** [VMware: Add trunk or access interfaces to a node](#)

Create a high availability group

When you create a high availability group, you select one or more interfaces and organize them in priority order. Then, you assign one or more VIP addresses to the group.

An interface must be for a Gateway Node or an Admin Node to be included in an HA group. An HA group can only use one interface for any given node; however, other interfaces for the same node can be used in other HA groups.

Access the wizard

Steps

1. Select **CONFIGURATION > Network > High availability groups**.
2. Select **Create**.

Enter details for the HA group

Steps

1. Provide a unique name for the HA group.
2. Optionally, enter a description for the HA group.
3. Select **Continue**.

Add interfaces to the HA group

Steps

1. Select one or more interfaces to add to this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

Add interfaces to the HA group

Select one or more interfaces for this HA group. You can select only one interface for each node.



Total interface count: 4

	Node 	Interface  	Site  	IPv4 subnet 	Node type  
☐	DC1-ADM1-104-96	eth0 	DC1	10.96.104.0/22	Primary Admin Node
☐	DC1-ADM1-104-96	eth2 	DC1	—	Primary Admin Node
☐	DC2-ADM1-104-103	eth0 	DC2	10.96.104.0/22	Admin Node
☐	DC2-ADM1-104-103	eth2 	DC2	—	Admin Node

0 interfaces selected




After creating a VLAN interface, wait up to 5 minutes for the new interface to appear in the table.

Guidelines for selecting interfaces

- You must select at least one interface.
- You can select only one interface for a node.
- If the HA group is for HA protection of Admin Node services, which include the Grid Manager and the Tenant Manager, select interfaces on Admin Nodes only.
- If the HA group is for HA protection of S3 client traffic, select interfaces on Admin Nodes, Gateway Nodes, or both.
- If you select interfaces on different types of nodes, an informational note appears. You are reminded that if a failover occurs, services provided by the previously active node might not be available on the newly active node. For example, a backup Gateway Node can't provide HA protection of Admin Node services. Similarly, a backup Admin Node can't perform all of the maintenance procedures that the primary Admin Node can provide.

- If you can't select an interface, its checkbox is disabled. The tool tip provides more information.



- You can't select an interface if its subnet value or gateway conflicts with another selected interface.
- You can't select a configured interface if it does not have a static IP address.

2. Select **Continue**.

Determine the priority order

If the HA group includes more than one interface, you can determine which is the Primary interface and which are the Backup (failover) interfaces. If the Primary interface fails, the VIP addresses move to the highest priority interface that is available. If that interface fails, the VIP addresses move to the next highest priority interface that is available, and so on.

Steps

1. Drag rows in the **Priority order** column to determine the Primary interface and any Backup interfaces.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

Determine the priority order

Determine the primary interface and the backup (failover) interfaces for this HA group. Drag and drop rows or select the arrows.

Priority order ?	Node	Interface ?	Node type ?
1 (Primary interface)	DC1-ADM1-104-96	eth2	Primary Admin Node
2	DC2-ADM1-104-103	eth2	Admin Node



If the HA group provides access to the Grid Manager, you must select an interface on the primary Admin Node to be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node.

2. Select **Continue**.

Enter IP addresses

Steps

1. In the **Subnet CIDR** field, specify the VIP subnet in CIDR notation—an IPv4 address followed by a slash and the subnet length (0-32).

The network address must not have any host bits set. For example, 192.16.0.0/22.



If you use a 32-bit prefix, the VIP network address also serves as the gateway address and the VIP address.

Enter details for the HA group

Subnet CIDR

Specify the subnet in CIDR notation. The optional gateway IP and all VIPs must be in this subnet.

IPv4 address followed by a slash and the subnet length (0-32)

Gateway IP address (optional)

Optionally specify the IP address of the gateway, which must be in the subnet. If the subnet address length is 32, the gateway IP address is automatically set to the subnet IP.

Virtual IP address

Specify at least 1 and no more than 10 virtual IPs for the HA group. All virtual IPs must be in the same subnet. If the subnet length is 32, only one VIP is allowed, which is automatically set to the subnet/gateway IP.

[Add another IP address](#)

2. Optionally, if any S3 administrative or tenant clients will access these VIP addresses from a different subnet, enter the **Gateway IP address**. The gateway address must be within the VIP subnet.

Client and admin users will use this gateway to access the virtual IP addresses.

3. Enter at least one and no more than ten VIP addresses for the active interface in the HA group. All VIP addresses must be within the VIP subnet and all will be active at the same time on the active interface.

You must provide at least one IPv4 address. Optionally, you can specify additional IPv4 and IPv6 addresses.

4. Select **Create HA group** and select **Finish**.

The HA Group is created, and you can now use the configured virtual IP addresses.

Next steps

If you will use this HA group for load balancing, create a load balancer endpoint to determine the port and network protocol and to attach any required certificates. See [Configure load balancer endpoints](#).

Edit a high availability group

You can edit a high availability (HA) group to change its name and description, add or remove interfaces, change the priority order, or add or update virtual IP addresses.

For example, you might need to edit an HA group if you want to remove the node associated with a selected interface in a site or node decommission procedure.

Steps

1. Select **CONFIGURATION > Network > High availability groups**.

The High availability groups page shows all existing HA groups.

2. Select the checkbox for the HA group you want to edit.
3. Do one of the following, based on what you want to update:
 - Select **Actions > Edit virtual IP address** to add or remove VIP addresses.
 - Select **Actions > Edit HA group** to update the group's name or description, add or remove interfaces, change the priority order, or add or remove VIP addresses.
4. If you selected **Edit virtual IP address**:
 - a. Update the virtual IP addresses for the HA group.
 - b. Select **Save**.
 - c. Select **Finish**.
5. If you selected **Edit HA group**:
 - a. Optionally, update the group's name or description.
 - b. Optionally, select or clear the checkboxes to add or remove interfaces.



If the HA group provides access to the Grid Manager, you must select an interface on the primary Admin Node to be the Primary interface. Some maintenance procedures can only be performed from the primary Admin Node

- c. Optionally, drag rows to change the priority order of the Primary interface and any Backup interfaces for this HA group.
- d. Optionally, update the virtual IP addresses.
- e. Select **Save** and then select **Finish**.

Remove a high availability group

You can remove one or more high availability (HA) groups at a time.



You can't remove an HA group if it is bound to a load balancer endpoint. To delete an HA group, you must remove it from any load balancer endpoints that use it.

To prevent client disruptions, update any affected S3 client applications before you remove an HA group. Update each client to connect using another IP address, for example, the virtual IP address of a different HA group or the IP address that was configured for an interface during installation.

Steps

1. Select **CONFIGURATION > Network > High availability groups**.

2. Review the **Load balancer endpoints** column for each HA group you want to remove. If any load balancer endpoints are listed:
 - a. Go to **CONFIGURATION > Network > Load balancer endpoints**.
 - b. Select the checkbox for the endpoint.
 - c. Select **Actions > Edit endpoint binding mode**.
 - d. Update the binding mode to remove the HA group.
 - e. Select **Save changes**.
3. If no load balancer endpoints are listed, select the checkbox for each HA group you want to remove.
4. Select **Actions > Remove HA group**.
5. Review the message and select **Delete HA group** to confirm your selection.

All HA groups you selected are removed. A green success banner appears on the High availability groups page.

Manage load balancing

Considerations for load balancing

You can use load balancing to handle ingest and retrieval workloads from S3 clients.

What is load balancing?

When a client application saves or retrieves data from a StorageGRID system, StorageGRID uses a load balancer to manage the ingest and retrieval workload. Load balancing maximizes speed and connection capacity by distributing the workload across multiple Storage Nodes.

The StorageGRID Load Balancer service is installed on all Admin Nodes and all Gateway Nodes and provides Layer 7 load balancing. It performs Transport Layer Security (TLS) termination of client requests, inspects the requests, and establishes new secure connections to the Storage Nodes.

The Load Balancer service on each node operates independently when forwarding client traffic to the Storage Nodes. Through a weighting process, the Load Balancer service routes more requests to Storage Nodes with higher CPU availability.



Although the StorageGRID Load Balancer service is the recommended load balancing mechanism, you might want to integrate a third-party load balancer instead. For information, contact your NetApp account representative or refer to [TR-4626: StorageGRID third-party and global load balancers](#).

How many load balancing nodes do I need?

As a general best practice, each site in your StorageGRID system should include two or more nodes with the Load Balancer service. For example, a site might include two Gateway Nodes or both an Admin Node and a Gateway Node. Make sure that there is adequate networking, hardware, or virtualization infrastructure for each load-balancing node, whether you are using services appliances, bare metal nodes, or virtual machine (VM) based nodes.

What is a load balancer endpoint?

A load balancer endpoint defines the port and the network protocol (HTTPS or HTTP) that incoming and outgoing client application requests will use to access those nodes that contain the Load Balancer service. The endpoint also defines the client type (S3), the binding mode, and optionally a list of allowed or blocked tenants.

To create a load balancer endpoint, either select **CONFIGURATION > Network > Load balancer endpoints** or complete the FabricPool and S3 setup wizard. For instructions:

- [Configure load balancer endpoints](#)
- [Use the S3 setup wizard](#)
- [Use the FabricPool setup wizard](#)

Considerations for the port

The port for a load balancer endpoint defaults to 10433 for the first endpoint you create, but you can specify any unused external port between 1 and 65535. If you use port 80 or 443, the endpoint will use the Load Balancer service on Gateway Nodes only. These ports are reserved on Admin Nodes. If you use the same port for more than one endpoint, you must specify a different binding mode for each endpoint.

Ports used by other grid services aren't permitted. See the [Network port reference](#).

Considerations for the network protocol

In most cases, the connections between client applications and StorageGRID should use Transport Layer Security (TLS) encryption. Connecting to StorageGRID without TLS encryption is supported but not recommended, especially in production environments. When you select the network protocol for the StorageGRID load balancer endpoint, you should select **HTTPS**.

Considerations for load balancer endpoint certificates

If you select **HTTPS** as the network protocol for the load balancer endpoint, you must provide a security certificate. You can use any of these three options when you create the load balancer endpoint:

- **Upload a signed certificate (recommended).** This certificate can be signed by either a publicly trusted or a private certificate authority (CA). Using a publicly trusted CA server certificate to secure the connection is the best practice. In contrast to generated certificates, certificates signed by a CA can be rotated nondisruptively, which can help avoid expiration issues.

You must obtain the following files before you create the load balancer endpoint:

- The custom server certificate file.
- The custom server certificate private key file.
- Optionally, a CA bundle of the certificates from each intermediate issuing certificate authority.
- **Generate a self-signed certificate.**
- **Use the global StorageGRID S3 certificate.** You must upload or generate a custom version of this certificate before you can select it for the load balancer endpoint. See [Configure S3 API certificates](#).

What values do I need?

To create the certificate, you must know all of the domain names and IP addresses that S3 client applications will use to access the endpoint.

The **Subject DN** (Distinguished Name) entry for the certificate must include the fully qualified domain name that the client application will use for StorageGRID. For example:

```
Subject DN:  
/C=Country/ST=State/O=Company, Inc./CN=s3.storagegrid.example.com
```

As required, the certificate can use wildcards to represent the fully qualified domain names of all Admin Nodes and Gateway Nodes running the Load Balancer service. For example, `*.storagegrid.example.com` uses the `*` wildcard to represent `adm1.storagegrid.example.com` and `gn1.storagegrid.example.com`.

If you plan to use S3 virtual hosted-style requests, the certificate must also include an **Alternative Name** entry for each [S3 endpoint domain name](#) you have configured, including any wildcard names. For example:

```
Alternative Name: DNS:*.s3.storagegrid.example.com
```



If you use wildcards for domain names, review the [Hardening guidelines for server certificates](#).

You must also define a DNS entry for each name in the security certificate.

How do I manage expiring certificates?



If the certificate used to secure the connection between the S3 application and StorageGRID expires, the application might temporarily lose access to StorageGRID.

To avoid certificate expiration issues, follow these best practices:

- Carefully monitor any alerts that warn of approaching certificate expiration dates, such as the **Expiration of load balancer endpoint certificate** and **Expiration of global server certificate for S3 API** alerts.
- Always keep the StorageGRID and S3 application's versions of the certificate in sync. If you replace or renew the certificate used for a load balancer endpoint, you must replace or renew the equivalent certificate used by the S3 application.
- Use a publicly signed CA certificate. If you use a certificate signed by a CA, you can replace soon-to-expire certificates nondisruptively.
- If you have generated a self-signed StorageGRID certificate and that certificate is about to expire, you must manually replace the certificate in both StorageGRID and in the S3 application before the existing certificate expires.

Considerations for the binding mode

The binding mode lets you control which IP addresses can be used to access a load balancer endpoint. If an endpoint uses a binding mode, client applications can only access the endpoint if they use an allowed IP address or its corresponding fully qualified domain name (FQDN). Client applications using any other IP address or FQDN can't access the endpoint.

You can specify any of the following binding modes:

- **Global** (default): Client applications can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN. Use

this setting unless you need to restrict the accessibility of an endpoint.

- **Virtual IPs of HA groups.** Client applications must use a virtual IP address (or corresponding FQDN) of an HA group.
- **Node interfaces.** Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces.
- **Node type.** Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node.

Considerations for tenant access

Tenant access is an optional security feature that lets you control which StorageGRID tenant accounts can use a load balancer endpoint to access their buckets. You can allow all tenants to access an endpoint (default), or you can specify a list of the allowed or blocked tenants for each endpoint.

You can use this feature to provide better security isolation between tenants and their endpoints. For example, you might use this feature to ensure that the top-secret or highly classified materials owned by one tenant remain completely inaccessible to other tenants.



For the purpose of access control, the tenant is determined from the access keys used in the client request, if no access keys are provided as part of the request (such as with anonymous access) the bucket owner is used to determine the tenant.

Tenant access example

To understand how this security feature works, consider the following example:

1. You have created two load balancer endpoints, as follows:
 - **Public** endpoint: Uses port 10443 and allows access to all tenants.
 - **Top secret** endpoint: Uses port 10444 and allows access to the **Top secret** tenant only. All other tenants are blocked from accessing this endpoint.
2. The `top-secret.pdf` is in a bucket owned by the **Top secret** tenant.

To access the `top-secret.pdf`, a user in the **Top secret** tenant can issue a GET request to `https://w.x.y.z:10444/top-secret.pdf`. Because this tenant is allowed to use the 10444 endpoint, the user can access the object. However, if a user belonging to any other tenant issues the same request to the same URL, they receive an immediate Access Denied message. Access is denied even if the credentials and signature are valid.

CPU availability

The Load Balancer service on each Admin Node and Gateway Node operates independently when forwarding S3 traffic to the Storage Nodes. Through a weighting process, the Load Balancer service routes more requests to Storage Nodes with higher CPU availability. Node CPU load information is updated every few minutes, but weighting might be updated more frequently. All Storage Nodes are assigned a minimal base weight value, even if a node reports 100% utilization or fails to report its utilization.

In some cases, information about CPU availability is limited to the site where the Load Balancer service is located.

Configure load balancer endpoints

Load balancer endpoints determine the ports and network protocols S3 clients can use when connecting to the StorageGRID load balancer on Gateway and Admin Nodes. You can also use endpoints to access the Grid Manager, Tenant Manager, or both.



Swift details have been removed from this version of the doc site. See [Configure S3 and Swift client connections](#).

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).
- You have reviewed the [considerations for load balancing](#).
- If you previously remapped a port you intend to use for the load balancer endpoint, you have [removed the port remap](#).
- You have created any high availability (HA) groups you plan to use. HA groups are recommended, but not required. See [Manage high availability groups](#).
- If the load balancer endpoint will be used by [S3 tenants for S3 Select](#), it must not use the IP addresses or FQDNs of any bare-metal nodes. Only services appliances and VMware-based software nodes are allowed for the load balancer endpoints used for S3 Select.
- You have configured any VLAN interfaces you plan to use. See [Configure VLAN interfaces](#).
- If you are creating an HTTPS endpoint (recommended), you have the information for the server certificate.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

- To upload a certificate, you need the server certificate, the certificate private key, and optionally, a CA bundle.
- To generate a certificate, you need all of the domain names and IP addresses that S3 clients will use to access the endpoint. You must also know the subject (Distinguished Name).
- If you want to use the StorageGRID S3 API certificate (which can also be used for connections directly to Storage Nodes), you have already replaced the default certificate with a custom certificate signed by an external certificate authority. See [Configure S3 API certificates](#).

Create a load balancer endpoint

Each S3 client load balancer endpoint specifies a port, a client type (S3), and a network protocol (HTTP or HTTPS). Management interface load balancer endpoints specifies a port, interface type, and untrusted Client Network.

Access the wizard

Steps

1. Select **CONFIGURATION > Network > Load balancer endpoints**.
2. To create an endpoint for an S3 or Swift client, select the **S3 or Swift client** tab.
3. To create an endpoint for access to the Grid Manager, Tenant Manager, or both, select the **Management interface** tab.

4. Select **Create**.

Enter endpoint details

Steps

1. Select the appropriate instructions to enter details for the type of endpoint you want to create.

S3 or Swift client

Field	Description
Name	A descriptive name for the endpoint, which will appear in the table on the Load balancer endpoints page.
Port	The StorageGRID port you want to use for load balancing. This field defaults to 10433 for the first endpoint you create, but you can enter any unused external port from 1 to 65535. If you enter 80 or 8443, the endpoint is configured only on Gateway Nodes, unless you have freed up port 8443. Then you can use port 8443 as an S3 endpoint, and the port will be configured on both Gateway and Admin Nodes.
Client type	The type of client application that will use this endpoint, either S3 or Swift .
Network protocol	The network protocol that clients will use when connecting to this endpoint. • Select HTTPS for secure, TLS encrypted communication (recommended). You must attach a security certificate before you can save the endpoint. • Select HTTP for less secure, unencrypted communication. Use HTTP only for a non-production grid.

Management interface

Field	Description
Name	A descriptive name for the endpoint, which will appear in the table on the Load balancer endpoints page.
Port	The StorageGRID port you want to use to access the Grid Manager, Tenant Manager, or both. • Grid Manager: 8443 • Tenant Manager: 9443 • Both Grid Manager and Tenant Manager: 443 Note: You can use these preset ports or other available ports.
Interface type	Select the radio button for the StorageGRID interface you will access using this endpoint.
Untrusted Client Network	Select Yes if this endpoint should be accessible to untrusted Client Networks. Otherwise, select No. When you select Yes, the port is open on all untrusted Client Networks. Note: You can only configure a port to be open or closed to untrusted Client Networks when you are creating the load balancer endpoint.

2. Select **Continue**.

Select a binding mode

Steps

1. Select a binding mode for the endpoint to control how the endpoint is accessed using any IP address or using specific IP addresses and network interfaces.

Some binding modes are available for either client endpoints or management interface endpoints. All modes for both endpoint types are listed here.

Mode	Description
Global (default for client endpoints)	Clients can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN. Use the Global setting unless you need to restrict the accessibility of this endpoint.
Virtual IPs of HA groups	Clients must use a virtual IP address (or corresponding FQDN) of an HA group to access this endpoint. Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints don't overlap.
Node interfaces	Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces to access this endpoint.
Node type (client endpoints only)	Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node to access this endpoint.
All Admin Nodes (default for management interface endpoints)	Clients must use the IP address (or corresponding FQDN) of any Admin Node to access this endpoint.

If more than one endpoint uses the same port, StorageGRID uses this priority order to decide which endpoint to use: **Virtual IPs of HA groups** > **Node interfaces** > **Node type** > **Global**.

If you are creating management interface endpoints, only Admin Nodes are allowed.

2. If you selected **Virtual IPs of HA groups**, select one or more HA groups.

If you are creating management interface endpoints, select VIPs associated only with Admin Nodes.

3. If you selected **Node interfaces**, select one or more node interfaces for each Admin Node or Gateway Node that you want to associate with this endpoint.
4. If you selected **Node type**, select either Admin Nodes, which includes both the primary Admin Node and any non-primary Admin Nodes, or Gateway Nodes.

Control tenant access



A management interface endpoint can control tenant access only when the endpoint has the [interface type of Tenant Manager](#).

Steps

1. For the **Tenant access** step, select one of the following:

Field	Description
Allow all tenants (default)	All tenant accounts can use this endpoint to access their buckets. You must select this option if you have not yet created any tenant accounts. After you add tenant accounts, you can edit the load balancer endpoint to allow or block specific accounts.
Allow selected tenants	Only the selected tenant accounts can use this endpoint to access their buckets.
Block selected tenants	The selected tenant accounts can't use this endpoint to access their buckets. All other tenants can use this endpoint.

2. If you are creating an **HTTP** endpoint, you don't need to attach a certificate. Select **Create** to add the new load balancer endpoint. Then, go to [After you finish](#). Otherwise, select **Continue** to attach the certificate.

Attach certificate

Steps

1. If you are creating an **HTTPS** endpoint, select the type of security certificate you want to attach to the endpoint.

The certificate secures the connections between S3 clients and the Load Balancer service on Admin Node or Gateway Nodes.

- **Upload certificate.** Select this option if you have custom certificates to upload.
- **Generate certificate.** Select this option if you have the values needed to generate a custom certificate.
- **Use StorageGRID S3 certificate.** Select this option if you want to use the global S3 API certificate, which can also be used for connections directly to Storage Nodes.

You can't select this option unless you have replaced the default S3 API certificate, which is signed by the grid CA, with a custom certificate signed by an external certificate authority. See [Configure S3 API certificates](#).

- **Use management interface certificate.** Select this option if you want to use the global management interface certificate, which can also be used for direct connections to Admin Nodes.
2. If you aren't using the StorageGRID S3 certificate, upload or generate the certificate.

Upload certificate

a. Select **Upload certificate**.

b. Upload the required server certificate files:

- **Server certificate**: The custom server certificate file in PEM encoding.
- **Certificate private key**: The custom server certificate private key file (.key).



EC private keys must be 224 bits or larger. RSA private keys must be 2048 bits or larger.

- **CA bundle**: A single optional file containing the certificates from each intermediate issuing certificate authority (CA). The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.

c. Expand **Certificate details** to see the metadata for each certificate you uploaded. If you uploaded an optional CA bundle, each certificate displays on its own tab.

- Select **Download certificate** to save the certificate file or select **Download CA bundle** to save the certificate bundle.

Specify the certificate file name and download location. Save the file with the extension .pem.

For example: storagegrid_certificate.pem

- Select **Copy certificate PEM** or **Copy CA bundle PEM** to copy the certificate contents for pasting elsewhere.

d. Select **Create**.

The load balancer endpoint is created. The custom certificate is used for all subsequent new connections between S3 clients or the management interface and the endpoint.

Generate certificate

a. Select **Generate certificate**.

b. Specify the certificate information:

Field	Description
Domain name	One or more fully qualified domain names to include in the certificate. Use an * as a wildcard to represent multiple domain names.
IP	One or more IP addresses to include in the certificate.
Subject (optional)	X.509 subject or distinguished name (DN) of the certificate owner. If no value is entered in this field, the generated certificate uses the first domain name or IP address as the subject common name (CN).
Days valid	Number of days after creation that the certificate expires.

Field	Description
Add key usage extensions	If selected (default and recommended), key usage and extended key usage extensions are added to the generated certificate. These extensions define the purpose of the key contained in the certificate. Note: Leave this checkbox selected unless you experience connection problems with older clients when certificates include these extensions.

c. Select **Generate**.

d. Select **Certificate details** to see the metadata for the generated certificate.

- Select **Download certificate** to save the certificate file.

Specify the certificate file name and download location. Save the file with the extension `.pem`.

For example: `storagegrid_certificate.pem`

- Select **Copy certificate PEM** to copy the certificate contents for pasting elsewhere.

e. Select **Create**.

The load balancer endpoint is created. The custom certificate is used for all subsequent new connections between S3 clients or the management interface and this endpoint.

After you finish

Steps

1. If you use a DNS, ensure that the DNS includes a record to associate the StorageGRID fully qualified domain name (FQDN) to each IP address that clients will use to make connections.

The IP address you enter in the DNS record depends on whether you are using an HA group of load-balancing nodes:

- If you have configured an HA group, clients will connect to the virtual IP addresses of that HA group.
- If you aren't using an HA group, clients will connect to the StorageGRID Load Balancer service using the IP address of a Gateway Node or Admin Node.

You must also ensure that the DNS record references all required endpoint domain names, including any wildcard names.

2. Provide S3 clients with the information needed to connect to the endpoint:

- Port number
- Fully qualified domain name or IP address
- Any required certificate details

View and edit load balancer endpoints

You can view details for existing load balancer endpoints, including the certificate metadata for a secured endpoint. You can change certain settings for an endpoint.

- To view basic information for all load balancer endpoints, review the tables on the Load balancer endpoints page.
- To view all details about a specific endpoint, including certificate metadata, select the endpoint's name in the table. The information shown varies depending on the endpoint type and how it's configured.

S3 load balancer endpoint

Port:	10443
Client type:	S3
Network protocol:	HTTPS
Binding mode:	Global
Endpoint ID:	3d02c126-9437-478c-8b24-08384401d3cb

[Remove](#)

Binding mode

Certificate

Tenant access (2 allowed)

You can select a different binding mode or change IP addresses for the current binding mode.

[Edit binding mode](#)

Binding mode: Global

 This endpoint uses the Global binding mode. Unless there are one or more overriding endpoints for the same port, clients can access this endpoint using the IP address of any Gateway Node, any Admin Node, or the virtual IP of any HA group on any network.

- To edit an endpoint, use the **Actions** menu on the Load balancer endpoints page.



If you lose access to Grid Manager while editing the port of a management interface endpoint, update the URL and port to regain access.



After editing an endpoint, you might need to wait up to 15 minutes for your changes to be applied to all nodes.

Task	Actions menu	Details page
Edit endpoint name	a. Select the checkbox for the endpoint. b. Select Actions > Edit endpoint name. c. Enter the new name. d. Select Save.	a. Select the endpoint name to display the details. b. Select the edit icon . c. Enter the new name. d. Select Save.
Edit endpoint port	a. Select the checkbox for the endpoint. b. Select Actions > Edit endpoint port c. Enter a valid port number. d. Select Save.	n/a
Edit endpoint binding mode	a. Select the checkbox for the endpoint. b. Select Actions > Edit endpoint binding mode. c. Update the binding mode as required. d. Select Save changes.	a. Select the endpoint name to display the details. b. Select Edit binding mode. c. Update the binding mode as required. d. Select Save changes.
Edit endpoint certificate	a. Select the checkbox for the endpoint. b. Select Actions > Edit endpoint certificate. c. Upload or generate a new custom certificate or begin using the global S3 certificate, as required. d. Select Save changes.	a. Select the endpoint name to display the details. b. Select the Certificate tab. c. Select Edit certificate. d. Upload or generate a new custom certificate or begin using the global S3 certificate, as required. e. Select Save changes.
Edit tenant access	a. Select the checkbox for the endpoint. b. Select Actions > Edit tenant access. c. Choose a different access option, select or remove tenants from the list, or do both. d. Select Save changes.	a. Select the endpoint name to display the details. b. Select the Tenant access tab. c. Select Edit tenant access. d. Choose a different access option, select or remove tenants from the list, or do both. e. Select Save changes.

Remove load balancer endpoints

You can remove one or more endpoints using the **Actions** menu, or you can remove a single endpoint from the details page.



To prevent client disruptions, update any affected S3 client applications before you remove a load balancer endpoint. Update each client to connect using a port assigned to another load balancer endpoint. Be sure to update any required certificate information as well.



If you lose access to Grid Manager while removing a management interface endpoint, update the URL.

- To remove one or more endpoints:
 - a. From the Load balancer page, select the checkbox for each endpoint you want to remove.
 - b. Select **Actions** > **Remove**.
 - c. Select **OK**.
- To remove one endpoint from the details page:
 - a. From the Load balancer page, select the endpoint name.
 - b. Select **Remove** on the details page.
 - c. Select **OK**.

Configure S3 endpoint domain names

To support S3 virtual-hosted-style requests, you must use the Grid Manager to configure the list of S3 endpoint domain names that S3 clients connect to.



Using an IP address for an endpoint domain name is unsupported. Future releases will prevent this configuration.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- You have confirmed that a grid upgrade is not in progress.



Don't make any changes to the domain name configuration when a grid upgrade is in progress.

About this task

To enable clients to use S3 endpoint domain names, you must do all of the following:

- Use the Grid Manager to add the S3 endpoint domain names to the StorageGRID system.
- Ensure that the [certificate the client uses for HTTPS connections to StorageGRID](#) is signed for all domain names that the client requires.

For example, if the endpoint is `s3.company.com`, you must ensure that the certificate used for HTTPS connections includes the `s3.company.com` endpoint and the endpoint's wildcard Subject Alternative Name (SAN): `*.s3.company.com`.

- Configure the DNS server used by the client. Include DNS records for the IP addresses that clients use to make connections, and ensure that the records reference all required S3 endpoint domain names, including any wildcard names.



Clients can connect to StorageGRID using the IP address of a Gateway Node, an Admin Node, or a Storage Node, or by connecting to the virtual IP address of a high availability group. You should understand how client applications connect to the grid so you include the correct IP addresses in the DNS records.

Clients that use HTTPS connections (recommended) to the grid can use either of these certificates:

- Clients that connect to a load balancer endpoint can use a custom certificate for that endpoint. Each load balancer endpoint can be configured to recognize different S3 endpoint domain names.
- Clients that connect to a load balancer endpoint or directly to a Storage Node can customize the global S3 API certificate to include all required S3 endpoint domain names.



If you don't add S3 endpoint domain names and the list is empty, support for S3 virtual-hosted-style requests is disabled.

Add an S3 endpoint domain name

Steps

1. Select **CONFIGURATION > Network > S3 endpoint domain names**.
2. Enter the domain name in the **Domain name 1** field. Select **Add another domain name** to add more domain names.
3. Select **Save**.
4. Ensure that the server certificates that clients use match the required S3 endpoint domain names.
 - If clients connect to a load balancer endpoint that uses its own certificate, [update the certificate associated with the endpoint](#).
 - If clients connect to a load balancer endpoint that uses the global S3 API certificate or directly to Storage Nodes, [update the global S3 API certificate](#).
5. Add the DNS records required to ensure that endpoint domain name requests can be resolved.

Result

Now, when clients use the endpoint *bucket.s3.company.com*, the DNS server resolves to the correct endpoint and the certificate authenticates the endpoint as expected.

Rename an S3 endpoint domain name

If you change a name used by S3 applications, virtual-hosted-style requests will fail.

Steps

1. Select **CONFIGURATION > Network > S3 endpoint domain names**.
2. Select the domain name field you want to edit and make the necessary changes.
3. Select **Save**.
4. Select **Yes** to confirm your change.

Delete an S3 endpoint domain name

If you remove a name used by S3 applications, virtual-hosted-style requests will fail.

Steps

1. Select **CONFIGURATION > Network > S3 endpoint domain names**.
2. Select the delete icon  next to the domain name.
3. Select **Yes** to confirm the deletion.

Related information

- [Use S3 REST API](#)
- [View IP addresses](#)
- [Configure high availability groups](#)

Summary: IP addresses and ports for client connections

To store or retrieve objects, S3 client applications connect to the Load Balancer service, which is included on all Admin Nodes and Gateway Nodes, or to the Local Distribution Router (LDR) service, which is included on all Storage Nodes.

Client applications can connect to StorageGRID using the IP address of a grid node and the port number of the service on that node. Optionally, you can create high availability (HA) groups of load-balancing nodes to provide highly available connections that use virtual IP (VIP) addresses. If you want to connect to StorageGRID using a fully qualified domain name (FQDN) instead of an IP or VIP address, you can configure DNS entries.

This table summarizes the different ways that clients can connect to StorageGRID and the IP addresses and ports that are used for each type of connection. If you have already created load balancer endpoints and high availability (HA) groups, see [Where to find IP addresses](#) to locate these values in the Grid Manager.

Where connection is made	Service that client connects to	IP address	Port
HA group	Load Balancer	Virtual IP address of an HA group	Port assigned to the load balancer endpoint
Admin Node	Load Balancer	IP address of the Admin Node	Port assigned to the load balancer endpoint
Gateway Node	Load Balancer	IP address of the Gateway Node	Port assigned to the load balancer endpoint
Storage Node	LDR	IP address of Storage Node	Default S3 ports: • HTTPS: 18082 • HTTP: 18084

Example URLs

To connect a client application to the Load Balancer endpoint of an HA group of Gateway Nodes, use a URL structured as shown below:

```
https://VIP-of-HA-group:LB-endpoint-port
```


For example, if the virtual IP address of the HA group is 192.0.2.5 and the port number of the load balancer endpoint is 10443, then an application could use the following URL to connect to StorageGRID:

```
https://192.0.2.5:10443
```

Where to find IP addresses

1. Sign in to the Grid Manager using a [supported web browser](#).
2. To find the IP address of a grid node:
 - a. Select **NODES**.
 - b. Select the Admin Node, Gateway Node, or Storage Node to which you want to connect.
 - c. Select the **Overview** tab.
 - d. In the Node Information section, note the IP addresses for the node.
 - e. Select **Show more** to view IPv6 addresses and interface mappings.

You can establish connections from client applications to any of the IP addresses in the list:

- **eth0**: Grid Network
- **eth1**: Admin Network (optional)
- **eth2**: Client Network (optional)



If you are viewing an Admin Node or a Gateway Node and it is the active node in a high availability group, the virtual IP address of the HA group is shown on eth2.

3. To find the virtual IP address of a high availability group:
 - a. Select **CONFIGURATION > Network > High availability groups**.
 - b. In the table, note the virtual IP address of the HA group.
4. To find the port number of a Load Balancer endpoint:
 - a. Select **CONFIGURATION > Network > Load balancer endpoints**.
 - b. Note the port number for the endpoint you want to use.



If the port number is 80 or 443, the endpoint is configured only on Gateway Nodes, because those ports are reserved on Admin Nodes. All other ports are configured on both Gateway Nodes and Admin Nodes.

- c. Select the name of the endpoint from the table.
- d. Confirm that the **Client type** (S3) matches the client application that will use the endpoint.

Manage networks and connections

Configure network settings

You can configure various network settings from the Grid Manager to fine tune the operation of your StorageGRID system.

Configure VLAN interfaces

You can [create virtual LAN \(VLAN\) interfaces](#) to isolate and partition traffic for security, flexibility, and performance. Each VLAN interface is associated with one or more parent interfaces on Admin Nodes and Gateway Nodes. You can use VLAN interfaces in HA groups and in load balancer endpoints to segregate client or admin traffic by application or tenant.

Traffic classification policies

You can use [traffic classification policies](#) to identify and handle different types of network traffic, including traffic related to specific buckets, tenants, client subnets, or load balancer endpoints. These policies can assist with traffic limiting and monitoring.

Guidelines for StorageGRID networks

You can use the Grid Manager to configure and manage StorageGRID networks and connections.

See [Configure S3 client connections](#) to learn how to connect S3 clients.

Default StorageGRID networks

By default, StorageGRID supports three network interfaces per grid node, allowing you to configure the networking for each individual grid node to match your security and access requirements.

For more information about network topology, see [Networking guidelines](#).

Grid Network

Required. The Grid Network is used for all internal StorageGRID traffic. It provides connectivity between all nodes in the grid, across all sites and subnets.

Admin Network

Optional. The Admin Network is typically used for system administration and maintenance. It can also be used for client protocol access. The Admin Network is typically a private network and does not need to be routable between sites.

Client Network

Optional. The Client Network is an open network typically used to provide access to S3 client applications, so the Grid Network can be isolated and secured. The Client Network can communicate with any subnet reachable through the local gateway.

Guidelines

- Each StorageGRID node requires a dedicated network interface, IP address, subnet mask, and gateway for each network it is assigned to.
- A grid node can't have more than one interface on a network.
- A single gateway, per network, per grid node is supported, and it must be on the same subnet as the node. You can implement more complex routing in the gateway, if required.
- On each node, each network maps to a specific network interface.

Network	Interface name
Grid	eth0
Admin (optional)	eth1
Client (optional)	eth2

- If the node is connected to a StorageGRID appliance, specific ports are used for each network. For details, see the installation instructions for your appliance.
- The default route is generated automatically, per node. If eth2 is enabled, then 0.0.0.0/0 uses the Client Network on eth2. If eth2 is not enabled, then 0.0.0.0/0 uses the Grid Network on eth0.
- The Client Network does not become operational until the grid node has joined the grid
- The Admin Network can be configured during grid node deployment to allow access to the installation user interface before the grid is fully installed.

Optional interfaces

Optionally, you can add extra interfaces to a node. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use [VLAN interfaces](#) to segregate the traffic belonging to different applications or tenants. Or, you might want to add an access interface to use in a [high availability \(HA\) group](#).

To add trunk or access interfaces, see the following:

- **VMware (after installing the node):** [VMware: Add trunk or access interfaces to a node](#)
 - **Red Hat Enterprise Linux (before installing the node):** [Create node configuration files](#)
 - **Ubuntu or Debian (before installing the node):** [Create node configuration files](#)
 - **RHEL, Ubuntu, or Debian (after installing the node):** [Linux: Add trunk or access interfaces to a node](#)

View IP addresses

You can view the IP address for each grid node in your StorageGRID system. You can then use this IP address to log in to the grid node at the command line and perform various maintenance procedures.

Before you begin

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

About this task

For information about changing IP addresses, see [Configure IP addresses](#).

Steps

1. Select **NODES > *grid node* > Overview**.
2. Select **Show more** to the right of the IP Addresses title.

The IP addresses for that grid node are listed in a table.

[Overview](#) [Hardware](#) [Network](#) [Storage](#) [Objects](#) [ILM](#) [Tasks](#)Node information [?](#)

Name: DC2-SGA-010-096-106-021

Type: Storage Node

ID: f0890e03-4c72-401f-ae92-245511a38e51

Connection state:  Connected

Storage used: Object data  7% [?](#)
Object metadata  5% [?](#)

Software version: 11.6.0 (build 20210915.1941.afce2d9)

IP addresses: 10.96.106.21 - eth0 (Grid Network)

[Hide additional IP addresses](#) [^](#)

Interface ^	IP address ^
eth0 (Grid Network)	10.96.106.21
eth0 (Grid Network)	fe80::2a0:98ff:fe64:6582
hic2	10.96.106.21
hic4	10.96.106.21
mtc2	169.254.0.1

Alerts

Alert name ^	Severity ? ^	Time triggered ^	Current values
ILM placement unachievable 🔗	 Major	2 hours ago ?	
A placement instruction in an ILM rule cannot be achieved for certain objects.			

Configure VLAN interfaces

You can create virtual LAN (VLAN) interfaces on Admin Nodes and Gateway Nodes and use them in HA groups and load balancer endpoints to isolate and partition traffic for security, flexibility, and performance.

Considerations for VLAN interfaces

- You create a VLAN interface by entering a VLAN ID and choosing a parent interface on one or more nodes.
- A parent interface must be configured as a trunk interface at the switch.
- A parent interface can be the Grid Network (eth0), the Client Network (eth2), or an additional trunk interface for the VM or bare-metal host (for example, ens256).

- For each VLAN interface, you can select only one parent interface for a given node. For example, you can't use both the Grid Network interface and the Client Network interface on the same Gateway Node as the parent interface for the same VLAN.
- If the VLAN interface is for Admin Node traffic, which includes traffic related to the Grid Manager and the Tenant Manager, select interfaces on Admin Nodes only.
- If the VLAN interface is for S3 client traffic, select interfaces on either Admin Nodes or Gateway Nodes.
- If you need to add trunk interfaces, see the following for details:
 - **VMware (after installing the node):** [VMware: Add trunk or access interfaces to a node](#)
 - **RHEL (before installing the node):** [Create node configuration files](#)
 - **Ubuntu or Debian (before installing the node):** [Create node configuration files](#)
 - **RHEL, Ubuntu, or Debian (after installing the node):** [Linux: Add trunk or access interfaces to a node](#)

Create a VLAN interface

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).
- A trunk interface has been configured in the network and attached to the VM or Linux node. You know the name of the trunk interface.
- You know the ID of the VLAN you are configuring.

About this task

Your network administrator might have configured one or more trunk interfaces and one or more VLANs to segregate the client or admin traffic belonging to different applications or tenants. Each VLAN is identified by a numeric ID or tag. For example, your network might use VLAN 100 for FabricPool traffic and VLAN 200 for an archive application.

You can use the Grid Manager to create VLAN interfaces that allow clients to access StorageGRID on a specific VLAN. When you create VLAN interfaces, you specify the VLAN ID and select parent (trunk) interfaces on one or more nodes.

Access the wizard

Steps

1. Select **CONFIGURATION > Network > VLAN interfaces**.
2. Select **Create**.

Enter details for the VLAN interfaces

Steps

1. Specify the ID of the VLAN in your network. You can enter any value between 1 and 4094.

VLAN IDs don't need to be unique. For example, you might use VLAN ID 200 for admin traffic at one site and the same VLAN ID for client traffic at another site. You can create separate VLAN interfaces with different sets of parent interfaces at each site. However, two VLAN interfaces with the same ID can't share the same interface on a node.

If you specify an ID that has already been used, a message appears.

2. Optionally, enter a short description for the VLAN interface.
3. Select **Continue**.

Choose parent interfaces

The table lists the available interfaces for all Admin Nodes and Gateway Nodes at each site in your grid. Admin Network (eth1) interfaces can't be used as parent interfaces and aren't shown.

Steps

1. Select one or more parent interfaces to attach this VLAN to.

For example, you might want to attach a VLAN to the Client Network (eth2) interface for a Gateway Node and an Admin Node.

Parent interfaces

Select one or more parent interfaces for this VLAN interface. You can only select one parent interface on each node for each VLAN interface.

	Site ?	Node name ?	Interface ?	Description ?	Node type ?	Attached VLANs ?
☐	Data Center 2	DC2-ADM1	eth0	Grid Network	Non-primary Admin	—
☒	Data Center 2	DC2-ADM1	eth2	Client Network	Non-primary Admin	—
☐	Data Center 1	DC1-G1	eth0	Grid Network	Gateway	—
☒	Data Center 1	DC1-G1	eth2	Client Network	Gateway	—
☐	Data Center 1	DC1-ADM1	eth0	Grid Network	Primary Admin	—

2 interfaces are selected.

[Previous](#)
[Continue](#)

2. Select **Continue**.

Confirm the settings

Steps

1. Review the configuration and make any changes.
 - If you need to change the VLAN ID or description, select **Enter VLAN details** at the top of the page.
 - If you need to change a parent interface, select **Choose parent interfaces** at the top of the page or select **Previous**.
 - If you need to remove a parent interface, select the trash can .
2. Select **Save**.
3. Wait up to 5 minutes for the new interface to appear as a selection on the High availability groups page and to be listed in the **Network interfaces** table for the node (**NODES** > **parent interface node** > **Network**).

Edit a VLAN interface

When you edit a VLAN interface, you can make the following types of changes:

- Change the VLAN ID or description.
- Add or remove parent interfaces.

For example, you might want to remove a parent interface from a VLAN interface if you plan to decommission the associated node.

Note the following:

- You can't change a VLAN ID if the VLAN interface is used in an HA group.
- You can't remove a parent interface if that parent interface is used in an HA group.

For example, suppose VLAN 200 is attached to parent interfaces on Nodes A and B. If an HA group uses the VLAN 200 interface for Node A and the eth2 interface for Node B, you can remove the unused parent interface for Node B, but you can't remove the used parent interface for Node A.

Steps

1. Select **CONFIGURATION > Network > VLAN interfaces**.
2. Select the checkbox for the VLAN interface you want to edit. Then, select **Actions > Edit**.
3. Optionally, update the VLAN ID or the description. Then, select **Continue**.

You can't update a VLAN ID if the VLAN is used in an HA group.

4. Optionally, select or clear the checkboxes to add parent interfaces or to remove unused interfaces. Then, select **Continue**.
5. Review the configuration and make any changes.
6. Select **Save**.

Remove a VLAN interface

You can remove one or more VLAN interfaces.

You can't remove a VLAN interface if it is currently used in an HA group. You must remove the VLAN interface from the HA group before you can remove it.

To avoid any disruptions in client traffic, consider doing one of the following:

- Add a new VLAN interface to the HA group before removing this VLAN interface.
- Create a new HA group that does not use this VLAN interface.
- If the VLAN interface you want to remove is currently the active interface, edit the HA group. Move the VLAN interface you want to remove to the bottom of the priority list. Wait until communication is established on the new primary interface and then remove the old interface from the HA group. Finally, delete the VLAN interface on that node.

Steps

1. Select **CONFIGURATION > Network > VLAN interfaces**.
2. Select the checkbox for each VLAN interface you want to remove. Then, select **Actions > Delete**.

3. Select **Yes** to confirm your selection.

All VLAN interfaces you selected are removed. A green success banner appears on the VLAN interfaces page.

Manage traffic classification policies

What are traffic classification policies?

Traffic classification policies allow you to identify and monitor different types of network traffic. These policies can assist with traffic limiting and monitoring to enhance your quality-of-service (QoS) offerings.

Traffic classification policies are applied to endpoints on the StorageGRID Load Balancer service for Gateway Nodes and Admin Nodes. To create traffic classification policies, you must have already created load balancer endpoints.

Matching rules

Each traffic classification policy contains one or more matching rules to identify the network traffic related to one or more of the following entities:

- Buckets
- Subnet
- Tenant
- Load balancer endpoints

StorageGRID monitors traffic that matches any rule within the policy according to the objectives of the rule. Any traffic that matches any rule for a policy is handled by that policy. Conversely, you can set rules to match all traffic except a specified entity.

Traffic limiting

Optionally, you can add the following limit types to a policy:

- Aggregate bandwidth
- Per-request bandwidth
- Concurrent requests
- Request rate

Limit values are enforced on a per load balancer basis. If traffic is distributed simultaneously across multiple load balancers, the total maximum rates are a multiple of the rate limits you specify.



You can create policies to limit aggregate bandwidth or to limit per-request bandwidth. However, StorageGRID can't limit both types of bandwidth at the same time. Aggregate bandwidth limits might impose an additional minor performance impact on non-limited traffic.

For aggregate or per-request bandwidth limits, the requests stream in or out at the rate you set. StorageGRID can only enforce one speed, so the most specific policy match, by matcher type, is the one enforced. The bandwidth consumed by the request does not count against other less specific matching policies containing

aggregate bandwidth limit policies. For all other limit types, client requests are delayed by 250 milliseconds and receive a 503 Slow Down response for requests that exceed any matching policy limit.

In the Grid Manager, you can view traffic charts and verify that the policies are enforcing the traffic limits you expect.

Use traffic classification policies with SLAs

You can use traffic classification policies in conjunction with capacity limits and data protection to enforce service-level agreements (SLAs) that provide specifics for capacity, data protection, and performance.

The following example shows three tiers of an SLA. You can create traffic classification policies to achieve the performance objectives of each SLA tier.

Service Level Tier	Capacity	Data Protection	Maximum performance allowed	Cost
Gold	1 PB storage allowed	3 copy ILM rule	25 K requests/sec 5 GB/sec (40 Gbps) bandwidth	\$\$\$ per month
Silver	250 TB storage allowed	2 copy ILM rule	10 K requests/sec 1.25 GB/sec (10 Gbps) bandwidth	\$\$ per month
Bronze	100 TB storage allowed	2 copy ILM rule	5 K requests/sec 1 GB/sec (8 Gbps) bandwidth	\$ per month

Create traffic classification policies

You can create traffic classification policies if you want to monitor, and optionally limit network traffic by bucket, bucket regex, CIDR, load balancer endpoint, or tenant. Optionally, you can set limits for a policy based on bandwidth, the number of concurrent requests, or the request rate.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).
- You have created any load balancer endpoints you want to match.
- You have created any tenants you want to match.

Steps

1. Select **CONFIGURATION > Network > Traffic classification**.
2. Select **Create**.

3. Enter a name and a description (optional) for the policy and select **Continue**.

For example, describe what this traffic classification policy applies to and what it will limit.

4. Select **Add rule** and specify the following details to create one or more matching rules for the policy. Any policy that you create should have at least one matching rule. Select **Continue**.

Field	Description
Type	Select the types of traffic that the matching rule applies to. Traffic types are bucket, bucket regex, CIDR, load balancer endpoint, and tenant.
Match value	Enter the value that matches the selected Type. • Bucket: Enter one or more bucket names. • Bucket regex: Enter one or more regular expressions used to match a set of bucket names. The regular expression is unanchored. Use the ^ anchor to match at the beginning of the bucket name, and use the \$ anchor to match at the end of the name. Regular expression matching supports a subset of PCRE (Perl compatible regular expression) syntax. • CIDR: Enter one or more IPv4 subnets, in CIDR notation, that matches the desired subnet. • Load balancer endpoint: Select an endpoint name. These are the load balancer endpoints you defined on the Configure load balancer endpoints. • Tenant: Tenant matching uses the access key ID. If the request does not contain an access key ID (for example, anonymous access), then the ownership of the bucket accessed is used to determine the tenant.
Inverse match	If you want to match all network traffic <i>except</i> traffic consistent with the Type and Match Value just defined, select the Inverse match checkbox. Otherwise, leave the checkbox cleared. For example, if you want this policy to apply to all but one of the load balancer endpoints, specify the load balancer endpoint to be excluded, and select Inverse match. For a policy containing multiple matchers where at least one is an inverse matcher, be careful not to create a policy that matches all requests.

5. Optionally, select **Add a limit** and select the following details to add one or more limits to control the network traffic matched by a rule.



StorageGRID collects metrics even if you don't add any limits, so you can understand traffic trends.

Field	Description
Type	The type of limit you want to apply to the network traffic matched by the rule. For example, you can limit bandwidth or request rate. Note: You can create policies to limit aggregate bandwidth or to limit per-request bandwidth. However, StorageGRID can't limit both types of bandwidth at the same time. When aggregate bandwidth is in use, per-request bandwidth is unavailable. Conversely, when per-request bandwidth is in use, aggregate bandwidth is unavailable. Aggregate bandwidth limits might impose an additional minor performance impact on non-limited traffic. For bandwidth limits, StorageGRID applies the policy that best matches the type of limit set. For example, if you have a policy that limits traffic in only one direction, then traffic in the opposite direction will be unlimited, even if there is traffic that matches additional policies that have bandwidth limits. StorageGRID implements "best" matches for bandwidth limits in the following order: • Exact IP address (/32 mask) • Exact bucket name • Bucket regex • Tenant • Endpoint • Non-exact CIDR matches (not /32) • Inverse matches
Applies to	Whether this limit applies to client read requests (GET or HEAD) or write requests (PUT, POST, or DELETE).
Value	The value that network traffic will be limited to, based on the Unit you select. For example, enter 10 and select MiB/s to prevent the network traffic matched by this rule from exceeding 10 MiB/s. Note: Depending on the units setting, the available units will be either binary (for example, GiB) or decimal (for example, GB). To change the units setting, select the user drop-down in the upper right of the Grid Manager, then select User Preferences.
Unit	The unit that describes the value you entered.

For example, if you want to create a 40 GB/s bandwidth limit for an SLA tier, create two Aggregate bandwidth limits: GET/HEAD at 40 GB/s and PUT/POST/DELETE at 40 GB/s.

6. Select **Continue**.
7. Read and review the Traffic classification policy. Use the **Previous** button to go back and make changes as required. When you are satisfied with the policy, select **Save and continue**.

S3 client traffic is now handled according to the traffic classification policy.

After you finish

View [network traffic metrics](#) to verify that the policies are enforcing the traffic limits you expect.

Edit traffic classification policy

You can edit a traffic classification policy to change its name or description, or to create, edit, or delete any rules or limits for the policy.

Before you begin

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

Steps

1. Select **CONFIGURATION > Network > Traffic classification**.

The Traffic classification policies page appears and the existing policies are listed in a table.

2. Edit the policy using the Actions menu or the details page. See [create traffic classification policies](#) for what to enter.

Actions menu

- a. Select the checkbox for the policy.
- b. Select **Actions > Edit**.

Details page

- a. Select the policy name.
- b. Select the **Edit** button beside the policy name.

3. For the Enter policy name step, optionally edit the policy name or description, and select **Continue**.
4. For the Add matching rules step, optionally add a rule or edit the **Type** and **Match value** of the existing rule, and select **Continue**.
5. For the Set limits step, optionally add, edit, or delete a limit, and select **Continue**.
6. Review the updated policy, and select **Save and continue**.

The changes you made to the policy are saved, and network traffic is now handled according to the traffic classification policies. You can view traffic charts and verify that the policies are enforcing the traffic limits you expect.

Delete a traffic classification policy

You can delete a traffic classification policy if you no longer need it. Make sure you delete the right policy because a policy can't be retrieved when deleted.

Before you begin

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

Steps

1. Select **CONFIGURATION > Network > Traffic classification**.

The Traffic classification policies page appears with the existing policies listed in a table.

2. Delete the policy using the Actions menu or the details page.

Actions menu

- a. Select the checkbox for the policy.
- b. Select **Actions > Remove**.

Policy details page

- a. Select the policy name.
- b. Select the **Remove** button beside the policy name.

3. Select **Yes** to confirm that you want to delete the policy.

The policy is deleted.

View network traffic metrics

You can monitor network traffic by viewing the graphs that are available from the Traffic classification policies page.

Before you begin

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

About this task

For any existing traffic classification policy, you can view metrics for the load balancer service to determine if the policy is successfully limiting traffic across the network. The data in the graphs can help you determine if you need to adjust the policy.

Even if no limits are set for a traffic classification policy, metrics are collected and the graphs provide useful information for understanding traffic trends.

Steps

1. Select **CONFIGURATION > Network > Traffic classification**.

The Traffic classification policies page appears, and the existing policies are listed in the table.

2. Select the traffic classification policy name for which you want to view metrics.
3. Select the **Metrics** tab.

The traffic classification policy graphs appear. The graphs display metrics only for the traffic that matches the selected policy.

The following graphs are included on the page.

- Request rate: This graph provides the amount of bandwidth matching this policy handled by all load balancers. Received data includes request headers for all requests and body data size for responses that have body data. Sent includes response headers for all requests and response body data size for requests that include body data in the response.



When requests are complete, this chart only shows bandwidth usage. For slow or large object requests the actual instantaneous bandwidth might differ from the values reported in this graph.

- Error response rate: This graph provides an approximate rate at which requests matching this policy are returning errors (HTTP status code ≥ 400) to clients.
 - Average request duration (non-error): This graph provides an average duration of successful requests matching this policy.
 - Policy bandwidth usage: This graph provides the amount of bandwidth matching this policy handled by all load balancers. Received data includes request headers for all requests and body data size for responses that have body data. Sent includes response headers for all requests and response body data size for requests that include body data in the response.
4. Position the cursor over a line graph to see a pop-up of values on a specific part of the graph.
 5. Select **Grafana dashboard** right below the Metrics title to view all the graphs for a policy. In addition to the four graphs from the **Metrics** tab, you can view two more graphs:
 - Write request rate by object size: The rate for PUT/POST/DELETE requests matching this policy. Positioning on an individual cell shows per second rates. Rates shown in the hover view are truncated to integer counts and might report 0 when there are non-zero requests in the bucket.
 - Read request rate by object size: The rate for GET/HEAD requests matching this policy. Positioning on an individual cell shows per second rates. Rates shown in the hover view are truncated to integer counts and might report 0 when there are non-zero requests in the bucket.
 6. Alternatively, access the graphs from the **SUPPORT** menu.
 - a. Select **SUPPORT > Tools > Metrics**.
 - b. Select **Traffic Classification Policy** from the **Grafana** section.
 - c. Select the policy from the menu on the upper left of the page.
 - d. Position the cursor over a graph to see a pop-up that shows the date and time of the sample, object sizes that are aggregated into the count, and the number of requests per second during that time period.

Traffic classification policies are identified by their ID. Policy IDs are listed on the Traffic classification policies page.
 7. Analyze the graphs to determine how often the policy is limiting traffic and whether you need to adjust the policy.

Supported ciphers for outgoing TLS connections

The StorageGRID system supports a limited set of cipher suites for Transport Layer Security (TLS) connections to the external systems used for identity federation and Cloud Storage Pools.

Supported versions of TLS

StorageGRID supports TLS 1.2 and TLS 1.3 for connections to external systems used for identity federation and Cloud Storage Pools.

The TLS ciphers that are supported for use with external systems have been selected to ensure compatibility with a range of external systems. The list is larger than the list of ciphers that are supported for use with S3 client applications. To configure ciphers, go to **CONFIGURATION > Security > Security settings** and select **TLS and SSH policies**.



TLS configuration options such as protocol versions, ciphers, key exchange algorithms, and MAC algorithms aren't configurable in StorageGRID. Contact your NetApp account representative if you have specific requests about these settings.

Benefits of active, idle, and concurrent HTTP connections

How you configure HTTP connections can impact the performance of the StorageGRID system. Configurations differ depending on whether the HTTP connection is active or idle or you have concurrent multiple connections.

You can identify the performance benefits for the following types of HTTP connections:

- Idle HTTP connections
- Active HTTP connections
- Concurrent HTTP connections

Benefits of keeping idle HTTP connections open

You should keep HTTP connections open even when client applications are idle to allow client applications to perform subsequent transactions over the open connection. Based on system measurements and integration experience, you should keep an idle HTTP connection open for a maximum of 10 minutes. StorageGRID might automatically close an HTTP connection that is kept open and idle for longer than 10 minutes.

Open and idle HTTP connections provide the following benefits:

- Reduced latency from the time that the StorageGRID system determines it has to perform an HTTP transaction to the time that the StorageGRID system can perform the transaction

Reduced latency is the main advantage, especially for the amount of time required to establish TCP/IP and TLS connections.

- Increased data transfer rate by priming the TCP/IP slow-start algorithm with previously performed transfers
- Instantaneous notification of several classes of fault conditions that interrupt connectivity between the client application and the StorageGRID system

Determining how long to keep an idle connection open is a trade-off between the benefits of slow start that is associated with the existing connection and the ideal allocation of the connection to internal system resources.

Benefits of active HTTP connections

For connections directly to Storage Nodes, you should limit the duration of an active HTTP connection to a maximum of 10 minutes, even if the HTTP connection continuously performs transactions.

Determining the maximum duration that a connection should be held open is a trade-off between the benefits of connection persistence and the ideal allocation of the connection to internal system resources.

For client connections to Storage Nodes, limiting active HTTP connections provides the following benefits:

- Enables optimal load balancing across the StorageGRID system.

Over time, an HTTP connection might no longer be optimal as load balancing requirements change. The system performs its best load balancing when client applications establish a separate HTTP connection for each transaction, but this negates the much more valuable gains associated with persistent connections.

- Allows client applications to direct HTTP transactions to LDR services that have available space.
- Allows maintenance procedures to start.

Some maintenance procedures start only after all the in-progress HTTP connections are complete.

For client connections to the Load Balancer service, limiting the duration of open connections can be useful for allowing some maintenance procedures to start promptly. If the duration of client connections is not limited, it might take several minutes for active connections to be automatically terminated.

Benefits of concurrent HTTP connections

You should keep multiple TCP/IP connections to the StorageGRID system open to allow parallelism, which increases performance. The optimal number of parallel connections depends on a variety of factors.

Concurrent HTTP connections provide the following benefits:

- Reduced latency

Transactions can start immediately instead of waiting for other transactions to be completed.

- Increased throughput

The StorageGRID system can perform parallel transactions and increase aggregate transaction throughput.

Client applications should establish multiple HTTP connections. When a client application has to perform a transaction, it can select and immediately use any established connection that is not currently processing a transaction.

Each StorageGRID system's topology has different peak throughput for concurrent transactions and connections before performance begins to degrade. Peak throughput depends on factors such as computing resources, network resources, storage resources, and WAN links. The number of servers and services and the number of applications that the StorageGRID system supports are also factors.

StorageGRID systems often support multiple client applications. You should keep this in mind when you determine the maximum number of concurrent connections used by a client application. If the client application consists of multiple software entities that each establish connections to the StorageGRID system, you should add up all the connections across the entities. You might have to adjust the maximum number of concurrent connections in the following situations:

- The StorageGRID system's topology affects the maximum number of concurrent transactions and connections that the system can support.

- Client applications that interact with the StorageGRID system over a network with limited bandwidth might have to reduce the degree of concurrency to ensure that individual transactions are completed in a reasonable time.
- When many client applications share the StorageGRID system, you might have to reduce the degree of concurrency to avoid exceeding the limits of the system.

Separation of HTTP connection pools for read and write operations

You can use separate pools of HTTP connections for read and write operations and control how much of a pool to use for each. Separate pools of HTTP connections enable you to better control transactions and balance loads.

Client applications can create loads that are retrieve-dominant (read) or store-dominant (write). With separate pools of HTTP connections for read and write transactions, you can adjust how much of each pool to dedicate for read or write transactions.


Manage link costs

Link costs let you prioritize which data center site provides a requested service when two or more data center sites exist. You can adjust link costs to reflect latency between sites.

What are link costs?

- Link costs are used to prioritize which object copy is used to fulfill object retrievals.
- Link costs are used by the Grid Management API and the Tenant Management API to determine which internal StorageGRID services to use.
- Link costs are used by the Load Balancer service on Admin Nodes and Gateway Nodes to direct client connections. See [Considerations for load balancing](#).

The diagram shows a three site grid that has link costs configured between sites:



- The Load Balancer service on Admin Nodes and Gateway Nodes equally distributes client connections to all Storage Nodes at the same data center site and to any data center sites with a link cost of 0.

In the example, a Gateway Node at data center site 1 (DC1) equally distributes client connections to Storage Nodes at DC1 and to Storage Nodes at DC2. A Gateway Node at DC3 sends client connections only to Storage Nodes at DC3.

- When retrieving an object that exists as multiple replicated copies, StorageGRID retrieves the copy at the data center that has the lowest link cost.

In the example, if a client application at DC2 retrieves an object that is stored both at DC1 and DC3, the object is retrieved from DC1, because the link cost from DC1 to DC2 is 0, which is lower than the link cost from DC3 to DC2 (25).

Link costs are arbitrary relative numbers with no specific unit of measure. For example, a link cost of 50 is used less preferentially than a link cost of 25. The table shows commonly used link costs.

Link	Link cost	Notes
Between physical data center sites	25 (default)	Data centers connected by a WAN link.
Between logical data center sites at the same physical location	0	Logical data centers in the same physical building or campus connected by a LAN.

Update link costs

You can update the link costs between data center sites to reflect latency between sites.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Grid topology page configuration permission](#).

Steps

1. Select **SUPPORT > Other > Link cost**.



Link Cost

Updated: 2023-02-15 18:09:28 MST

Site Names (1 - 3 of 3)

Site ID	Site Name	Actions
10	Data Center 1	
20	Data Center 2	
30	Data Center 3	

Show Records Per Page

Previous
« 1 » Next

Link Costs

Link Source	Link Destination			Actions
	10	20	30	
Data Center 1	0	25	25	



2. Select a site under **Link Source** and enter a cost value between 0 and 100 under **Link Destination**.

You can't change the link cost if the source is the same as the destination.

To cancel changes, select  **Revert**.

3. Select **Apply Changes**.

Use AutoSupport

What is AutoSupport?

The AutoSupport feature enables StorageGRID to send health and status packages to NetApp technical support.

Using AutoSupport can significantly speed up problem determination and resolution. Technical support can also monitor the storage needs of your system and help you determine if you need to add new nodes or sites. Optionally, you can configure AutoSupport packages to be sent to one additional destination.

StorageGRID has two types of AutoSupport:

- **StorageGRID AutoSupport** reports StorageGRID software issues. Enabled by default when you first install StorageGRID. You can [change the default AutoSupport configuration](#) if needed.



If StorageGRID AutoSupport is not enabled, a message appears on the Grid Manager dashboard. The message includes a link to the AutoSupport configuration page. If you close the message, it will not appear again until your browser cache is cleared, even if AutoSupport remains disabled.

- **Appliance hardware AutoSupport** reports StorageGRID appliance issues. You must [configure hardware AutoSupport on each appliance](#).

What is Active IQ?

Active IQ is a cloud-based digital advisor that leverages predictive analytics and community wisdom from NetApp's installed base. Its continuous risk assessments, predictive alerts, prescriptive guidance, and automated actions help you prevent problems before they occur, leading to improved system health and higher system availability.

If you want to use the Active IQ dashboards and functionality on the NetApp Support Site, you must enable AutoSupport.

[Active IQ Digital Advisor Documentation](#)

Information included in AutoSupport package

An AutoSupport package contains the following files and details.

File name	Fields	Description
AUTOSUPPORT-HISTORY.XML	AutoSupport Sequence Number Destination for this AutoSupport Status of Delivery Delivery Attempts AutoSupport Subject Delivery URI Last error AutoSupport PUT Filename Time of Generation Autosupport Compressed Size Autosupport Decompressed Size Total Collection Time (ms)	AutoSupport history file.
AUTOSUPPORT.XML	Node Protocol to contact support Support URL for HTTP/HTTPS Support Address AutoSupport OnDemand State AutoSupport OnDemand Server URL AutoSupport OnDemand Polling Interval	AutoSupport status file. Provides details of protocol used, technical support URL and address, polling interval, and OnDemand AutoSupport if enabled or disabled.

File name	Fields	Description
BUCKETS.XML	Bucket ID Account ID Build Version Location Constraint Configuration Compliance Enabled Compliance Configuration S3 Object Lock Enabled S3 Object Lock Configuration Consistency Configuration CORS Enabled CORS Configuration Last Access Time Enabled Policy Enabled Policy Configuration Notifications Enabled Notifications Configuration Cloud Mirror Enabled Cloud Mirror Configuration Search Enabled Search Configuration Bucket Tagging Enabled Bucket Tagging Configuration Versioning Configuration	Provides configuration details and statistics at the bucket level. Example of bucket configurations include platform services, compliance, and bucket consistency.
GRID-CONFIGURATIONS.XML	Attribute ID Attribute Name Value Index Table ID Table Name	Grid-wide configuration information file. Contains information about grid certificates, metadata reserved space, grid-wide configuration settings (compliance, S3 Object Lock, object compression, alerts, syslog, and ILM configuration), erasure-coding profile details, DNS name, and NMS name .
GRID-SPEC.XML	Grid specifications, raw XML	Used for configuring and deploying StorageGRID. Contains grid specifications, NTP server IP, DNS server IP, network topology, and hardware profiles of the nodes.
GRID-TASKS.XML	Node Service Path Attribute ID Attribute name Value Index Table ID Table name	Grid tasks (maintenance procedures) status file. Provides details of the grid's active, terminated, completed, failed, and pending tasks.

File name	Fields	Description
GRID.JSON	Grid Revision Software Version Description License Passwords DNS NTP Sites Nodes	Grid information.
ILM-CONFIGURATION.XML	Attribute ID Attribute Name Value Index Table ID Table Name	List of attributes for ILM configurations.
ILM-STATUS.XML	Node Service path Attribute ID Attribute name Value Index Table ID Table name	ILM metrics information file. Contains ILM evaluation rates for each node and grid-wide metrics.
ILM.XML	ILM raw XML	ILM active policy file. Contains details about the active ILM policies, such as storage pool ID, ingest behavior, filters, rules, and description.
LOG.TGZ	<i>n/a</i>	Downloadable log file. Contains <code>bycast-err.log</code> and <code>servermanager.log</code> from each node.
MANIFEST.XML	Collection order AutoSupport content filename for this data Description of this data item Number of bytes collected Time spent collecting Status of this data item Description of the error AutoSupport content type for this data +	Contains AutoSupport metadata and brief descriptions of all AutoSupport files.

File name	Fields	Description
NMS-ENTITIES.XML	Attribute index Entity OID Node ID Device model ID Device model version Entity name	Group and service entities in the NMS tree . Provides grid topology details. The node can be determined based on the services running on the node.
OBJECTS-STATUS.XML	Node Service path Attribute ID Attribute name Value Index Table ID Table name	Object status, including background scan status, active transfer, transfer rate, total transfers, delete rate, corrupted fragments, lost objects, missing objects, repair attempted, scan rate, estimated scan period, and repair completion status.
SERVER-STATUS.XML	Node Service path Attribute ID Attribute name Value Index Table ID Table name	Server configurations. Contains these details for each node: platform type, operating system, installed memory, available memory, storage connectivity, storage appliance chassis serial number, storage controller failed drive count, compute controller chassis temperature, compute hardware, compute controller serial number, power supply, drive size, and drive type.
SERVICE-STATUS.XML	Node Service path Attribute ID Attribute name Value Index Table ID Table name	Service node information file. Contains details such as allocated table space, free table space, Reaper metrics of the database, segment repair duration, repair job duration, auto job restarts, and auto job termination.
STORAGE-GRADES.XML	Storage grade ID Storage grade name Storage node ID Storage node path	Storage grade definitions file for each Storage Node.
SUMMARY-ATTRIBUTES.XML	Group OID Group Path Summary attribute ID Summary attribute name Value Index Table ID Table name	High-level system status data that summarizes StorageGRID usage information. Provides details such as name of grid, names of sites, number of Storage Nodes per grid and per site, license type, license capacity and usage, software support terms, and details of S3 operations.

File name	Fields	Description
SYSTEM-ALERTS.XML	Name Severity Node name Alert Status Site name Alert triggered time Alert resolved time Rule ID Node ID Site ID Silenced Other annotations Other labels	Current system alerts that indicate potential problems in the StorageGRID system.
USERAGENTS.XML	User agent Number of days Total HTTP requests Total bytes ingested Total bytes retrieved PUT requests GET requests DELETE requests HEAD requests POST requests OPTIONS requests Average request time (ms) Average PUT request time (ms) Average GET request time (ms) Average DELETE request time (ms) Average HEAD request time (ms) Average POST request time (ms) Average OPTIONS request time (ms)	Statistics based on the application user agents. For example, the number of PUT/GET/DELETE/HEAD operations per user agent and total bytes size of each operation.
X-HEADER-DATA	X-Netapp-asup-generated-on X-Netapp-asup-hostname X-Netapp-asup-os-version X-Netapp-asup-serial-num X-Netapp-asup-subject X-Netapp-asup-system-id X-Netapp-asup-model-name +	AutoSupport header data.

Configure AutoSupport

By default, the StorageGRID AutoSupport feature is enabled when you first install StorageGRID. However, you must configure hardware AutoSupport on each appliance. As needed, you can change the AutoSupport configuration.

If you want to change the configuration of StorageGRID AutoSupport, make your changes only on the primary Admin Node. You must [configure hardware AutoSupport](#) on each appliance.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).
- If you will use HTTPS for sending AutoSupport packages, you have provided outbound internet access to the primary Admin Node, either directly or [using a proxy server](#) (inbound connections not required).
- If HTTP is selected on the StorageGRID AutoSupport page, you have [configured a proxy server](#) to forward AutoSupport packages as HTTPS. NetApp's AutoSupport servers will reject packages sent using HTTP.
- If you will use SMTP as the protocol for AutoSupport packages, you have configured an SMTP mail server.

About this task

You can use any combination of the following options to send AutoSupport packages to technical support:

- **Weekly:** Automatically send AutoSupport packages once per week. Default setting: Enabled.
- **Event-triggered:** Automatically send AutoSupport packages every hour or when significant system events occur. Default setting: Enabled.
- **On Demand:** Allow technical support to request that your StorageGRID system send AutoSupport packages automatically, which is useful when they are actively working on an issue (requires HTTPS AutoSupport transmission protocol). Default setting: Disabled.
- **User-triggered:** Manually send AutoSupport packages at any time.

Specify the protocol for AutoSupport packages

You can use any of the following protocols for sending AutoSupport packages:

- **HTTPS:** This is the default and recommended setting for new installations. This protocol uses port 443. If you want to [enable the AutoSupport on Demand feature](#), you must use HTTPS.
- **HTTP:** If you select HTTP, you must configure a proxy server to forward AutoSupport packages as HTTPS. NetApp's AutoSupport servers reject packages sent using HTTP. This protocol uses port 80.
- **SMTP:** Use this option if you want AutoSupport packages to be emailed.

The protocol you set is used for sending all types of AutoSupport packages.

Steps

1. Select **SUPPORT > Tools > AutoSupport > Settings**.
2. Select the protocol you want to use to send AutoSupport packages.
3. If you selected **HTTPS**, select whether to use a NetApp support certificate (TLS certificate) to secure the connection to the technical support server.
 - **Verify certificate** (default): Ensures that the transmission of AutoSupport packages is secure. The NetApp support certificate is already installed with the StorageGRID software.

- **Do not verify certificate:** Select this option only when you have a good reason not to use certificate validation, such as when there is a temporary problem with a certificate.

4. Select **Save**. All weekly, user-triggered, and event-triggered packages are sent using the selected protocol.

Disable weekly AutoSupport

By default, the StorageGRID system is configured to send an AutoSupport package to technical support once a week.

To determine when the weekly AutoSupport package will be sent, go to the **AutoSupport > Results** tab. In the **Weekly AutoSupport** section, look at the value for **Next Scheduled Time**.

You can disable the automatic sending of weekly AutoSupport packages at any time.

Steps

1. Select **SUPPORT > Tools > AutoSupport > Settings**.
2. Clear the **Enable Weekly AutoSupport** checkbox.
3. Select **Save**.

Disable event-triggered AutoSupport

By default, the StorageGRID system is configured to send an AutoSupport package to technical support every hour.

You can disable event-triggered AutoSupport at any time.

Steps

1. Select **SUPPORT > Tools > AutoSupport > Settings**.
2. Clear the **Enable Event-Triggered AutoSupport** checkbox.
3. Select **Save**.

Enable AutoSupport on Demand

AutoSupport on Demand can assist in solving issues that technical support is actively working on.

By default, AutoSupport on Demand is disabled. Enabling this feature allows technical support to request that your StorageGRID system send AutoSupport packages automatically. Technical support can also set the polling time interval for AutoSupport on Demand queries.

Technical support can't enable or disable AutoSupport on Demand.

Steps

1. Select **SUPPORT > Tools > AutoSupport > Settings**.
2. Select the **HTTPS** for the protocol.
3. Select the **Enable Weekly AutoSupport** checkbox.
4. Select the **Enable AutoSupport on Demand** checkbox.
5. Select **Save**.

AutoSupport on Demand is enabled, and technical support can send AutoSupport on Demand requests to StorageGRID.

Disable checks for software updates

By default, StorageGRID contacts NetApp to determine if software updates are available for your system. If a StorageGRID hotfix or new version is available, the new version is shown on the StorageGRID Upgrade page.

As required, you can optionally disable the check for software updates. For example, if your system does not have WAN access, you should disable the check to avoid download errors.

Steps

1. Select **SUPPORT > Tools > AutoSupport > Settings**.
2. Clear the **Check for software updates** checkbox.
3. Select **Save**.

Add an additional AutoSupport destination

When you enable AutoSupport, health and status packages are sent to technical support. You can specify one additional destination for all AutoSupport packages.

To verify or change the protocol used to send AutoSupport packages, see the instructions to [specify the protocol for AutoSupport packages](#).



You can't use the SMTP protocol to send AutoSupport packages to an additional destination.

Steps

1. Select **SUPPORT > Tools > AutoSupport > Settings**.
2. Select **Enable Additional AutoSupport Destination**.
3. Specify the following:

Hostname

The server hostname or IP address of an additional AutoSupport destination server.



You can enter only one additional destination.

Port

The port used to connect to an additional AutoSupport destination server. The default is port 80 for HTTP or port 443 for HTTPS.

Certificate validation

Whether a TLS certificate is used to secure the connection to the additional destination.

- Select **Verify certificate** to use certificate validation.
- Select **Do not verify certificate** to send your AutoSupport packages without certificate validation.

Select this choice only when you have a good reason not to use certificate validation, such as when there is a temporary problem with a certificate.

4. If you selected **Verify certificate**, do the following:
 - a. Browse to the location of the CA certificate.
 - b. Upload the CA certificate file.

The CA certificate metadata appears.

5. Select **Save**.

All future weekly, event-triggered, and user-triggered AutoSupport packages will be sent to the additional destination.

Configure AutoSupport for appliances

AutoSupport for appliances reports StorageGRID hardware issues, and StorageGRID AutoSupport reports StorageGRID software issues, with one exception: for the SGF6112, StorageGRID AutoSupport reports both hardware and software issues. You must configure AutoSupport on each appliance except the SGF6112, which does not require additional configuration. AutoSupport is implemented differently for services appliances and storage appliances.

You use SANtricity to enable AutoSupport for each storage appliance. You can configure SANtricity AutoSupport during initial appliance setup or after an appliance has been installed:

- For SG6000 and SG5700 appliances, [configure AutoSupport in SANtricity System Manager](#)

AutoSupport packages from E-Series appliances can be included in StorageGRID AutoSupport if you configure AutoSupport delivery by proxy in [SANtricity System Manager](#).

StorageGRID AutoSupport does not report hardware issues, such as DIMM or host interface card (HIC) faults. However, some component failures might trigger [hardware alerts](#). For StorageGRID appliances with a baseboard management controller (BMC) you can configure email and SNMP traps to report hardware failures:

- [Set up email notifications for BMC alerts](#)
- [Configure SNMP settings for BMC](#)

Related information

[NetApp Support](#)

Manually trigger an AutoSupport package

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

Before you begin

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

Steps

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

StorageGRID attempts to send an AutoSupport package to the NetApp Support Site. If the attempt is successful, the **Most Recent Result** and **Last Successful Time** values on the **Results** tab are updated. If there is a problem, the **Most Recent Result** value updates to "Failed," and StorageGRID does not try to send the AutoSupport package again.



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

Troubleshoot AutoSupport packages

If an attempt to send an AutoSupport package fails, the StorageGRID system takes different actions depending on the type of AutoSupport package. You can check the status of AutoSupport packages by selecting **SUPPORT > Tools > AutoSupport > Results**.

When the AutoSupport package fails to send, "Failed" appears on the **Results** tab of the **AutoSupport** page.



If you configured a proxy server to forward AutoSupport packages to NetApp, you should [verify that the proxy server configuration settings are correct](#).

Weekly AutoSupport package failure

If a weekly AutoSupport package fails to send, the StorageGRID system takes the following actions:

1. Updates the Most Recent Result attribute to Retrying.
2. Attempts to resend the AutoSupport package 15 times every four minutes for one hour.
3. After one hour of send failures, updates the Most Recent Result attribute to Failed.
4. Attempts to send an AutoSupport package again at the next scheduled time.
5. Maintains the regular AutoSupport schedule if the package fails because the NMS service is unavailable, and if a package is sent before seven days pass.
6. When the NMS service is available again, sends an AutoSupport package immediately if a package has not been sent for seven days or more.

User-triggered or event-triggered AutoSupport package failure

If a user-triggered or an event-triggered AutoSupport package fails to send, the StorageGRID system takes the following actions:

1. Displays an error message if the error is known. For example, if a user selects the SMTP protocol without providing correct email configuration settings, the following error is displayed: AutoSupport packages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.
2. Does not attempt to send the package again.
3. Logs the error in `nms.log`.

If a failure occurs and SMTP is the selected protocol, verify that the StorageGRID system's email server is correctly configured and that your email server is running (**SUPPORT > Alarms (legacy) > Legacy Email Setup**). The following error message might appear on the AutoSupport page: AutoSupport packages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.

Learn how to [configure email server settings](#).

Correct an AutoSupport package failure

If a failure occurs and SMTP is the selected protocol, verify that the StorageGRID system's email server is correctly configured and that your email server is running. The following error message might appear on the AutoSupport page: AutoSupport packages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.

Send E-Series AutoSupport packages through StorageGRID

You can send E-Series SANtricity System Manager AutoSupport packages to technical support through a StorageGRID Admin Node rather than the storage appliance management port.

See [E-Series hardware AutoSupport](#) for more information about using AutoSupport with E-Series appliances.

Before you begin

- You are signed into the Grid Manager using a [supported web browser](#).
- You have the [Storage appliance administrator or Root access permission](#).
- You have configured SANtricity AutoSupport:
 - For SG6000 and SG5700 appliances, [configure AutoSupport in SANtricity System Manager](#)



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager.

About this task

E-Series AutoSupport packages contain details of the storage hardware and are more specific than other AutoSupport packages sent by the StorageGRID system.

You can configure a special proxy server address in SANtricity System Manager to transmit AutoSupport packages through a StorageGRID Admin Node without the use of the appliance's management port. AutoSupport packages transmitted in this way are sent by the [preferred sender Admin Node](#), and they use any [admin proxy settings](#) that have been configured in the Grid Manager.




This procedure is only for configuring a StorageGRID proxy server for E-Series AutoSupport packages. For additional details on E-Series AutoSupport configuration, see the [NetApp E-Series and SANtricity Documentation](#).

Steps

1. In the Grid Manager, select **NODES**.
2. From the list of nodes on the left, select the storage appliance node you want to configure.
3. Select **SANtricity System Manager**.

The SANtricity System Manager home page appears.



4. Select **SUPPORT** > **Support center** > **AutoSupport**.

The AutoSupport operations page appears.

Technical Support

Chassis serial number: 031517000693

 [NetApp My Support](#)

US/Canada 888.463.8277

[Other Contacts](#)

Support Resources

Diagnostics

AutoSupport

AutoSupport operations

AutoSupport status: Enabled 

[Enable/Disable AutoSupport Features](#)
AutoSupport proactively monitors the health of your storage array and automatically sends support data ("dispatches") to the support team.

[Configure AutoSupport Delivery Method](#)
Connect to the support team via HTTPS, HTTP or Mail (SMTP) server delivery methods.

[Schedule AutoSupport Dispatches](#)
AutoSupport dispatches are sent daily at 03:06 PM UTC and weekly at 07:39 AM UTC on Thursday.

[Send AutoSupport Dispatch](#)
Automatically sends the support team a dispatch to troubleshoot system issues without waiting for periodic dispatches.

[View AutoSupport Log](#)
The AutoSupport log provides information about status, dispatch history, and errors encountered during delivery of AutoSupport dispatches.

[Enable AutoSupport Maintenance Window](#)
Enable AutoSupport Maintenance window to allow maintenance activities to be performed on the storage array without generating support cases.

[Disable AutoSupport Maintenance Window](#)
Disable AutoSupport Maintenance window to allow the storage array to generate support cases on component failures and other destructive actions.

5. Select **Configure AutoSupport Delivery Method**.

The Configure AutoSupport Delivery Method page appears.

6. Select **HTTPS** for the delivery method.



The certificate that enables HTTPS is pre-installed.

7. Select **via Proxy server**.

8. Enter `tunnel-host` for the **Host address**.

`tunnel-host` is the special address to use an Admin Node to send E-Series AutoSupport packages.

9. Enter `10225` for the **Port number**.

`10225` is the port number on the StorageGRID proxy server that receives AutoSupport packages from the E-Series controller in the appliance.

10. Select **Test Configuration** to test the routing and configuration of your AutoSupport proxy server.

If correct, a message in a green banner appears: "Your AutoSupport configuration has been verified."

If the test fails, an error message appears in a red banner. Check your StorageGRID DNS settings and

networking, ensure the [preferred sender Admin Node](#) can connect to the NetApp Support Site, and try the test again.

11. Select **Save**.

The configuration is saved, and a confirmation message appears: "AutoSupport delivery method has been configured."

Manage Storage Nodes

Manage Storage Nodes

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

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

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. However, writes to other volumes can continue 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** applies only 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 **SUPPORT > Other > Storage watermarks**.
2. On the Storage watermarks page, look at the Use optimized values checkbox.
 - If the checkbox is selected, 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. Do not update these values. Optionally, you can [View optimized storage watermarks](#).

- If the Use optimized values checkbox is unselected, 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.

When you specify custom watermark settings, you must enter values greater than 0.

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, 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 11.9	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 settings**.
2. On the Storage settings page, expand the **Metadata reserved space** section.

For StorageGRID 11.8 or higher, the Metadata reserved space value must be at least 100 GB and no more than 1 PB.

The default setting for a new StorageGRID 11.6 or higher installation in which each Storage Node has 128 GB or more of RAM is 8,000 GB (8 TB).

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


Size of volume 0 for the node	Actual reserved space for metadata
500 GB or more or Metadata-only Storage Nodes	The smaller of these values: • Volume 0 • Metadata reserved space setting Note: Only one rangedb is required for metadata-only Storage Nodes.

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 or later. 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 a grid](#) by adding new Storage Nodes.

Increase Metadata Reserved Space setting

You might be able to increase the Metadata Reserved Space system setting if your Storage Nodes meet specific requirements for RAM and available space.

What you'll need

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

- You have the [Root Access permission or the Grid Topology Page Configuration and Other Grid Configuration permissions](#).



The Grid topology page has been deprecated and will be removed in a future release.

About this task

You might be able to manually increase the system-wide Metadata Reserved Space setting up to 8 TB.

You can only increase the value of the system-wide Metadata Reserved Space setting if both of these statements are true:

- The Storage Nodes at any site in your system each have 128 GB or more RAM.
- The Storage Nodes at any site in your system each have sufficient available space on storage volume 0.

Be aware that if you increase this setting, you will simultaneously reduce the space available for object storage on storage volume 0 of all Storage Nodes. For this reason, you might prefer to set the Metadata Reserved Space to a value smaller than 8 TB, based on your expected object metadata requirements.



In general, it is better to use a higher value instead of a lower value. If the Metadata Reserved Space setting is too large, you can decrease it later. In contrast, if you increase the value later, the system might need to move object data to free up space.

For a detailed explanation of how the Metadata Reserved Space setting affects the allowed space for object metadata storage on a particular Storage Node, see [Manage object metadata storage](#).

Steps

1. Determine the current Metadata Reserved Space setting.
 - a. Select **CONFIGURATION > System > Storage options**.
 - b. In the Storage watermarks section, note the value of **Metadata Reserved Space**.
2. Ensure you have enough available space on storage volume 0 of each Storage Node to increase this value.
 - a. Select **NODES**.
 - b. Select the first Storage Node in the grid.
 - c. Select the Storage tab.
 - d. In the Volumes section, locate the **/var/local/rangedb/0** entry.
 - e. Confirm that the Available value is equal to or greater than difference between the new value you want to use and the current Metadata Reserved Space value.

For example, if the Metadata Reserved Space setting is currently 4 TB and you want to increase it to 6 TB, the Available value must be 2 TB or greater.
 - f. Repeat these steps for all Storage Nodes.
 - If one or more Storage Nodes do not have enough available space, the Metadata Reserved Space value cannot be increased. Do not continue with this procedure.
 - If each Storage Node has enough available space on volume 0, go to the next step.
3. Ensure you have at least 128 GB of RAM on each Storage Node.
 - a. Select **NODES**.

- b. Select the first Storage Node in the grid.
- c. Select the **Hardware** tab.
- d. Hover your cursor over the Memory Usage chart. Ensure that **Total Memory** is at least 128 GB.
- e. Repeat these steps for all Storage Nodes.
 - If one or more Storage Nodes do not have enough available total memory, the Metadata Reserved Space value cannot be increased. Do not continue with this procedure.
 - If each Storage Node has at least 128 GB of total memory, go to the next step.

4. Update the Metadata Reserved Space setting.

- a. Select **CONFIGURATION > System > Storage options**.
- b. Select the Configuration tab.
- c. In the Storage watermarks section, select **Metadata Reserved Space**.
- d. Enter the new value.

For example, to enter 8 TB, which is the maximum supported value, enter **8000000000000** (8, followed by 12 zeros)

Storage Options

Overview

Configuration

Configure Storage Options
Updated: 2021-12-10 13:48:23 MST

Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1000000000

Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark Override	0
Storage Volume Soft Read-Only Watermark Override	0
Storage Volume Hard Read-Only Watermark Override	0
Metadata Reserved Space	8000000000000

Apply Changes

- e. Select **Apply Changes**.

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 client applications should avoid performing GetObject 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. GetObject 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 > Storage settings > Object compression**.
2. Select the **Compress stored objects** checkbox.
3. Select **Save**.

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 or SG6160, 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](#).

Manage Admin Nodes

Use multiple Admin Nodes

A StorageGRID system can include multiple Admin Nodes to enable you to continuously monitor and configure your StorageGRID system even if one Admin Node fails.

If an Admin Node becomes unavailable, attribute processing continues, alerts are still triggered, and email notifications and AutoSupport packages are still sent. However, having multiple Admin Nodes does not provide failover protection except for notifications and AutoSupport packages.



There are two options for continuing to view and configure the StorageGRID system if an Admin Node fails:

- Web clients can reconnect to any other available Admin Node.
- If a system administrator has configured a high availability group of Admin Nodes, web clients can continue to access the Grid Manager or the Tenant Manager using the virtual IP address of the HA group. See



When using an HA group, access is interrupted if the active Admin Node fails. Users must sign in again after the virtual IP address of the HA group fails over to another Admin Node in the group.

Some maintenance tasks can only be performed using the primary Admin Node. If the primary Admin Node fails, it must be recovered before the StorageGRID system is fully functional again.

Identify the primary Admin Node

The primary Admin Node provides more functionality than non-primary Admin Nodes. For example, some maintenance procedures must be performed using the primary Admin Node.

For more information about Admin Nodes, see [What is an Admin Node](#).

Before you begin

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

Steps

1. Select **NODES**.
2. Enter **primary** in the search box.

In the search results, identify the node with "Primary Admin Node" displayed in the Type column. One primary Admin Node should be listed.

View notification status and queues

The Network Management System (NMS) service on Admin Nodes sends notifications to the mail server. You can view the current status of the NMS service and the size of its notifications queue on the Interface Engine page.

To access the Interface Engine page, select **SUPPORT > Tools > Grid topology**. Then select **site > Admin Node > NMS > Interface Engine**.

Overview
Alarms
Reports
Configuration

Main


Overview: NMS (170-176) - Interface Engine
Updated: 2009-03-09 10:12:17 PDT

NMS Interface Engine Status:	Connected	
Connected Services:	15	

E-mail Notification Events

E-mail Notifications Status:	No Errors	
E-mail Notifications Queued:	0	

Database Connection Pool

Maximum Supported Capacity:	100	
Remaining Capacity:	95 %	
Active Connections:	5	

Notifications are processed through the email notifications queue and are sent to the mail server one after another in the order they are triggered. If there is a problem (for example, a network connection error) and the mail server is unavailable when the attempt is made to send the notification, a best effort attempt to resend the notification to the mail server continues for a period of 60 seconds. If the notification is not sent to the mail server after 60 seconds, the notification is dropped from the notifications queue and an attempt to send the next notification in the queue is made.

Manage objects with ILM

Manage objects with ILM

The information lifecycle management (ILM) rules in an ILM policy instruct StorageGRID how to create and distribute copies of object data and how to manage those copies over time.

About these instructions

Designing and implementing ILM rules and policies requires careful planning. You must understand your operational requirements, the topology of your StorageGRID system, your object protection needs, and the available storage types. Then, you must determine how you want different types of objects to be copied, distributed, and stored.


Use these instructions to:

- Learn about StorageGRID ILM, including [how ILM operates throughout an object's life](#).
- Learn how to configure [storage pools](#), [Cloud Storage Pools](#), and [ILM rules](#).
- Learn how to [create, simulate, and activate an ILM policy](#) that will protect object data across one or more sites.
- Learn how to [manage objects with S3 Object Lock](#), which helps to ensure that objects in specific S3 buckets aren't deleted or overwritten for a specified amount of time.


Learn more

To learn more, review these videos:

- [Video: ILM rules overview.](#)



- [Video: ILM policies overview](#)



ILM and object lifecycle

How ILM operates throughout an object's life

Understanding how StorageGRID uses ILM to manage objects during every stage of their life can help you design a more effective policy.


- **Ingest:** Ingest begins when an S3 client application establishes a connection to save an object to the StorageGRID system, and is complete when StorageGRID returns an "ingest successful" message to the client. Object data is protected during ingest either by applying ILM instructions immediately (synchronous placement) or by creating interim copies and applying ILM later (dual commit), depending on how the ILM requirements were specified.
- **Copy management:** After creating the number and type of object copies that are specified in the ILM's placement instructions, StorageGRID manages object locations and protects objects against loss.
 - **ILM scanning and evaluation:** StorageGRID continuously scans the list of objects stored in the grid and checks if the current copies meet ILM requirements. When different types, numbers, or locations of object copies are required, StorageGRID creates, deletes, or moves copies as needed.
 - **Background verification:** StorageGRID continuously performs background verification to check the integrity of object data. If a problem is found, StorageGRID automatically creates a new object copy or a replacement erasure-coded object fragment in a location that meets current ILM requirements. See [Verify object integrity](#).
- **Object deletion:** Management of an object ends when all copies are removed from the StorageGRID

system. Objects can be removed as a result of a delete request by a client, or as a result of deletion by ILM or deletion caused by the expiration of an S3 bucket lifecycle.



Objects in a bucket that has S3 Object Lock enabled can't be deleted if they are under a legal hold or if a retain-until-date has been specified but not yet met.

The diagram summarizes how ILM operates throughout an object's lifecycle.



How objects are ingested


Ingest options

When you create an ILM rule, you specify one of three options for protecting objects at ingest: Dual commit, Strict, or Balanced.

Depending on your choice, StorageGRID makes interim copies and queues the objects for ILM evaluation later, or it uses synchronous placement and immediately makes copies to meet ILM requirements.

Flowchart of ingest options

The flowchart shows what happens when objects are matched by an ILM rule that uses each of the three ingest options.



Dual commit

When you select the Dual commit option, StorageGRID immediately makes interim object copies on two different Storage Nodes and returns an "ingest successful" message to the client. The object is queued for ILM evaluation, and copies that meet the rule's placement instructions are made later. If the ILM policy can't be processed immediately after the dual commit, site-loss protection could take time to achieve.

Use the Dual commit option in either of these cases:

- You are using multi-site ILM rules and client ingest latency is your primary consideration. When using Dual commit, you must ensure your grid can perform the additional work of creating and removing the dual-commit copies if they don't satisfy ILM. Specifically:
 - The load on the grid must be low enough to prevent an ILM backlog.
 - The grid must have excess hardware resources (IOPS, CPU, memory, network bandwidth, and so on).
- You are using multi-site ILM rules and the WAN connection between the sites usually has high latency or limited bandwidth. In this scenario, using the Dual commit option can help prevent client timeouts. Before choosing the Dual commit option, you should test the client application with realistic workloads.

Balanced (default)

When you select the Balanced option, StorageGRID also uses synchronous placement on ingest and immediately makes all copies specified in the rule's placement instructions. In contrast with the Strict option, if StorageGRID can't immediately make all copies, it uses Dual commit instead. If the ILM policy uses placements on multiple sites and immediate site-loss protection can't be achieved, the **ILM placement unachievable** alert is triggered.

Use the Balanced option to achieve the best combination of data protection, grid performance, and ingest success. Balanced is the default option in the Create ILM rule wizard.

Strict

When you select the Strict option, StorageGRID uses synchronous placement on ingest and immediately makes all object copies specified in the rule's placement instructions. Ingest fails if StorageGRID can't create all copies, for example, because a required storage location is temporarily unavailable. The client must retry the operation.

Use the Strict option if you have an operational or regulatory requirement to immediately store objects only in the locations outlined in the ILM rule. For example, to satisfy a regulatory requirement, you might need to use the Strict option and a Location Constraint advanced filter to guarantee that objects are never stored at certain data centers.

See [Example 5: ILM rules and policy for Strict ingest behavior](#).

Advantages, disadvantages, and limitations of the ingest options

Understanding the advantages and disadvantages of each of the three options for protecting data at ingest (Balanced, Strict, or Dual commit) can help you decide which one to select for an ILM rule.

For an overview of ingest options, see [Ingest options](#).

Advantages of the Balanced and Strict options

When compared to Dual commit, which creates interim copies during ingest, the two synchronous placement options can provide the following advantages:

- **Better data security:** Object data is immediately protected as specified in the ILM rule's placement instructions, which can be configured to protect against a wide variety of failure conditions, including the failure of more than one storage location. Dual commit can only protect against the loss of a single local copy.
- **More efficient grid operation:** Each object is processed only once, as it is ingested. Because the StorageGRID system does not need to track or delete interim copies, there is less processing load and less database space is consumed.
- **(Balanced) Recommended:** The Balanced option provides optimal ILM efficiency. Using the Balanced option is recommended unless Strict ingest behavior is required or the grid meets all of the criteria for using Dual commit.
- **(Strict) Certainty about object locations:** The Strict option guarantees that objects are immediately stored according to the placement instructions in the ILM rule.

Disadvantages of the Balanced and Strict options

When compared to Dual commit, the Balanced and Strict options have some disadvantages:

- **Longer client ingests:** Client ingest latencies might be longer. When you use the Balanced or Strict options, an "ingest successful" message is not returned to the client until all erasure-coded fragments or replicated copies are created and stored. However, object data will most likely reach its final placement much faster.
- **(Strict) Higher rates of ingest failure:** With the Strict option, ingest fails whenever StorageGRID can't immediately make all copies specified in the ILM rule. You might see high rates of ingest failure if a required storage location is temporarily offline or if network issues cause delays in copying objects between sites.
- **(Strict) S3 multipart upload placements might not be as expected in some circumstances:** With Strict, you expect objects either to be placed as described by the ILM rule or for ingest to fail. However, with an S3 multipart upload, ILM is evaluated for each part of the object as it is ingested, and for the object as a whole when the multipart upload completes. In the following circumstances this might result in placements that are different than you expect:
 - **If ILM changes while an S3 multipart upload is in progress:** Because each part is placed according to the rule that is active when the part is ingested, some parts of the object might not meet current ILM requirements when the multipart upload completes. In these cases, ingest of the object does not fail. Instead, any part that is not placed correctly is queued for ILM re-evaluation and is moved to the correct location later.
 - **When ILM rules filter on size:** When evaluating ILM for a part, StorageGRID filters on the size of the part, not the size of the object. This means that parts of an object can be stored in locations that don't meet ILM requirements for the object as a whole. For example, if a rule specifies that all objects 10 GB or larger are stored at DC1 while all smaller objects are stored at DC2, at ingest each 1 GB part of a 10-part multipart upload is stored at DC2. When ILM is evaluated for the object, all parts of the object are moved to DC1.
- **(Strict) Ingest does not fail when object tags or metadata are updated and newly required placements cannot be made:** With Strict, you expect objects either to be placed as described by the ILM rule or for ingest to fail. However, when you update metadata or tags for an object that is already stored in the grid, the object is not re-ingested. This means that any changes to object placement that are triggered by the update aren't made immediately. Placement changes are made when ILM is re-evaluated by normal

background ILM processes. If required placement changes can't be made (for example, because a newly required location is unavailable), the updated object retains its current placement until the placement changes are possible.

Limitations on object placements with the Balanced and Strict options

The Balanced or Strict options can't be used for ILM rules that have any of these placement instructions:

- Placement in a Cloud Storage Pool at day 0.
- Placements in a Cloud Storage Pool when the rule has a User defined creation time as its Reference time.

These restrictions exist because StorageGRID can't synchronously make copies to a Cloud Storage Pool, and a User defined creation time could resolve to the present.

How ILM rules and consistency interact to affect data protection

Both your ILM rule and your choice of consistency affect how objects are protected. These settings can interact.

For example, the ingest behavior selected for an ILM rule affects the initial placement of object copies, while the consistency used when an object is stored affects the initial placement of object metadata. Because StorageGRID requires access to both an object's data and metadata to fulfill client requests, selecting matching levels of protection for the consistency and ingest behavior can provide better initial data protection and more predictable system responses.

Here is a brief summary of the consistency values that are available in StorageGRID:

- **All:** All nodes receive object metadata immediately or the request will fail.
- **Strong-global:** Object metadata is immediately distributed to all sites. Guarantees read-after-write consistency for all client requests across all sites.
- **Strong-site:** Object metadata is immediately distributed to other nodes at the site. Guarantees read-after-write consistency for all client requests within a site.
- **Read-after-new-write:** Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.
- **Available:** Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that don't exist). Not supported for S3 FabricPool buckets.



Before selecting a consistency value, [read the full description of consistency](#). You should understand the benefits and limitations before changing the default value.

Example of how consistency and ILM rules can interact

Suppose you have a two-site grid with the following ILM rule and the following consistency:

- **ILM rule:** Create two object copies, one at the local site and one at a remote site. Use Strict ingest behavior.
- **consistency:** Strong-global (object metadata is immediately distributed to all sites).

When a client stores an object to the grid, StorageGRID makes both object copies and distributes metadata to both sites before returning success to the client.

The object is fully protected against loss at the time of the ingest successful message. For example, if the local site is lost shortly after ingest, copies of both the object data and the object metadata still exist at the remote site. The object is fully retrievable.

If you instead used the same ILM rule and the strong-site consistency, the client might receive a success message after object data is replicated to the remote site but before object metadata is distributed there. In this case, the level of protection of object metadata does not match the level of protection for object data. If the local site is lost shortly after ingest, object metadata is lost. The object can't be retrieved.

The inter-relationship between consistency and ILM rules can be complex. Contact NetApp if you need assistance.

Related information

[Example 5: ILM rules and policy for Strict ingest behavior](#)


How objects are stored (replication or erasure coding)

What is replication?

Replication is one of two methods used by StorageGRID to store object data (erasure coding is the other method). When objects match an ILM rule that uses replication, the system creates exact copies of object data and stores the copies on Storage Nodes.

When you configure an ILM rule to create replicated copies, you specify how many copies should be created, where those copies should be placed, and how long the copies should be stored at each location.

In the following example, the ILM rule specifies that two replicated copies of each object be placed in a storage pool that contains three Storage Nodes.



When StorageGRID matches objects to this rule, it creates two copies of the object, placing each copy on a

different Storage Node in the storage pool. The two copies might be placed on any two of the three available Storage Nodes. In this case, the rule placed object copies on Storage Nodes 2 and 3. Because there are two copies, the object can be retrieved if any of the nodes in the storage pool fails.



StorageGRID can store only one replicated copy of an object on any given Storage Node. If your grid includes three Storage Nodes and you create a 4-copy ILM rule, only three copies will be made—one copy for each Storage Node. The **ILM placement unachievable** alert is triggered to indicate that the ILM rule could not be completely applied.

Related information

- [What is erasure coding](#)
- [What is a storage pool](#)
- [Enable site-loss protection using replication and erasure coding](#)


Why you should not use single-copy replication

When creating an ILM rule to create replicated copies, you should always specify at least two copies for any time period in the placement instructions.




Don't use an ILM rule that creates only one replicated copy for any time period. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.


In the following example, the Make 1 Copy ILM rule specifies that one replicated copy of an object be placed in a storage pool that contains three Storage Nodes. When an object is ingested that matches this rule, StorageGRID places a single copy on only one Storage Node.



When an ILM rule creates only one replicated copy of an object, the object becomes inaccessible when the Storage Node is unavailable. In this example, you will temporarily lose access to object AAA whenever Storage Node 2 is offline, such as during an upgrade or other maintenance procedure. You will lose object AAA entirely if Storage Node 2 fails.



To avoid losing object data, you should always make at least two copies of all objects you want to protect with replication. If two or more copies exist, you can still access the object if one Storage Node fails or goes offline.





What is erasure coding?

Erasure coding is one of two methods StorageGRID uses to store object data (replication is the other method). When objects match an ILM rule that uses erasure coding, those objects are sliced into data fragments, additional parity fragments are computed, and each fragment is stored on a different Storage Node.



When an object is accessed, it is reassembled using the stored fragments. If a data or a parity fragment becomes corrupt or lost, the erasure-coding algorithm can recreate that fragment using a subset of the remaining data and parity fragments.

As you create ILM rules, StorageGRID creates erasure-coding profiles that support those rules. You can view a list of erasure-coding profiles, [rename an erasure-coding profile](#), or [deactivate an erasure-coding profile if it is not currently used in any ILM rules](#).


The following example illustrates the use of an erasure-coding algorithm on an object's data. In this example, the ILM rule uses a 4+2 erasure-coding scheme. Each object is sliced into four equal data fragments, and two parity fragments are computed from the object data. Each of the six fragments is stored on a different node across three data center sites to provide data protection for node failures or site loss.



The 4+2 erasure-coding scheme can be configured in various ways. For example, you can configure a single-site storage pool that contains six Storage Nodes. For [site-loss protection](#), you can use a storage pool containing three sites with three Storage Nodes at each site. An object can be retrieved as long as any four of the six fragments (data or parity) remain available. Up to two fragments can be lost without loss of the object data. If an entire site is lost, the object can still be retrieved or repaired, as long as all of the other fragments remain accessible.



If more than two Storage Nodes are lost, the object is not retrievable.



Related information

- [What is replication](#)
- [What is a storage pool](#)
- [What are erasure-coding schemes](#)
- [Rename an erasure-coding profile](#)
- [Deactivate an erasure-coding profile](#)

What are erasure-coding schemes?

Erasure-coding schemes control how many data fragments and how many parity fragments are created for each object.

When you create or edit an ILM rule, you select an available erasure-coding scheme. StorageGRID automatically creates erasure-coding schemes based on how many Storage Nodes and sites make up the storage pool you plan to use.

Data protection

The StorageGRID system uses the Reed-Solomon erasure-coding algorithm. The algorithm slices an object into k data fragments and computes m parity fragments.

The $k + m = n$ fragments are spread across n Storage Nodes to provide data protection as follows:

- To retrieve or repair an object, k fragments are needed.
- An object can sustain up to m lost or corrupt fragments. The higher the value of m , the higher the failure tolerance.

The best data protection is provided by the erasure-coding scheme with the highest node or volume failure

tolerance within a storage pool.

Storage overhead

The storage overhead of an erasure-coding scheme is calculated by dividing the number of parity fragments (m) by the number of data fragments (k). You can use the storage overhead to calculate how much disk space each erasure-coded object requires:

$$\text{disk space} = \text{object size} + (\text{object size} * \text{storage overhead})$$

For example, if you store a 10 MB object using the 4+2 scheme (which has 50% storage overhead), the object consumes 15 MB of grid storage. If you store the same 10 MB object using the 6+2 scheme (which has 33% storage overhead), the object consumes approximately 13.3 MB.

Select the erasure-coding scheme with the lowest total value of $k+m$ that meets your needs. Erasure-coding schemes with a lower number of fragments are more computationally efficient because:

- Fewer fragments are created and distributed (or retrieved) per object
- They show better performance because the fragment size is larger
- They can require fewer nodes be added in an [expansion when more storage is required](#)

Guidelines for storage pools

When selecting the storage pool to use for a rule that will create an erasure-coded copy, use the following guidelines for storage pools:

- The storage pool must include three or more sites, or exactly one site.



You can't use erasure coding if the storage pool includes two sites.

- [Erasure-coding schemes for storage pools containing three or more sites](#)
- [Erasure-coding schemes for one-site storage pools](#)
- Don't use a storage pool that includes the All Sites site.
- The storage pool should include at least $k+m + 1$ Storage Nodes that can store object data.



Storage Nodes can be configured during installation to contain only object metadata and not object data. For more information, see [Types of Storage Nodes](#).

The minimum number of Storage Nodes required is $k+m$. However, having at least one additional Storage Node can help prevent ingest failures or ILM backlogs if a required Storage Node is temporarily unavailable.

Erasure-coding schemes for storage pools containing three or more sites

The following table describes the erasure-coding schemes currently supported by StorageGRID for storage pools that include three or more sites. All of these schemes provide site-loss protection. One site can be lost, and the object will still be accessible.

For erasure-coding schemes that provide site-loss protection, the recommended number of Storage Nodes in the storage pool exceeds $k+m + 1$ because each site requires a minimum of three Storage Nodes.

Erasure-coding scheme ($k+m$)	Minimum number of deployed sites	Recommended number of Storage Nodes at each site	Total recommended number of Storage Nodes	Site loss protection?	Storage overhead
4+2	3	3	9	Yes	50%
6+2	4	3	12	Yes	33%
8+2	5	3	15	Yes	25%
6+3	3	4	12	Yes	50%
9+3	4	4	16	Yes	33%
2+1	3	3	9	Yes	50%
4+1	5	3	15	Yes	25%
6+1	7	3	21	Yes	17%
7+5	3	5	15	Yes	71%



StorageGRID requires a minimum of three Storage Nodes per site. To use the 7+5 scheme, each site requires a minimum of four Storage Nodes. Using five Storage Nodes per site is recommended.

When selecting an erasure-coding scheme that provides site protection, balance the relative importance of the following factors:

- **Number of fragments:** Performance and expansion flexibility are generally better when the total number of fragments is lower.
- **Fault tolerance:** Fault tolerance is increased by having more parity segments (that is, when m has a higher value.)
- **Network traffic:** When recovering from failures, using a scheme with more fragments (that is, a higher total for $k+m$) creates more network traffic.
- **Storage overhead:** Schemes with higher overhead require more storage space per object.

For example, when deciding between a 4+2 scheme and 6+3 scheme (which both have 50% storage overhead), select the 6+3 scheme if additional fault tolerance is required. Select the 4+2 scheme if network resources are constrained. If all other factors are equal, select 4+2 because it has a lower total number of fragments.



If you are unsure of which scheme to use, select 4+2 or 6+3, or contact technical support.

Erasure-coding schemes for one-site storage pools

A one-site storage pool supports all of the erasure-coding schemes defined for three or more sites, provided that the site has enough Storage Nodes.

The minimum number of Storage Nodes required is $k+m$, but a storage pool with $k+m + 1$ Storage Nodes is recommended. For example, the 2+1 erasure-coding scheme requires a storage pool with a minimum of three Storage Nodes, but four Storage Nodes is recommended.

Erasure-coding scheme ($k+m$)	Minimum number of Storage Nodes	Recommended number of Storage Nodes	Storage overhead
4+2	6	7	50%
6+2	8	9	33%
8+2	10	11	25%
6+3	9	10	50%
9+3	12	13	33%
2+1	3	4	50%
4+1	5	6	25%
6+1	7	8	17%
7+5	12	13	71%

Advantages, disadvantages, and requirements for erasure coding

Before deciding whether to use replication or erasure coding to protect object data from loss, you should understand the advantages, disadvantages, and the requirements for erasure coding.

Advantages of erasure coding

When compared to replication, erasure coding offers improved reliability, availability, and storage efficiency.

- **Reliability:** Reliability is gauged in terms of fault tolerance—that is, the number of simultaneous failures that can be sustained without loss of data. With replication, multiple identical copies are stored on different nodes and across sites. With erasure coding, an object is encoded into data and parity fragments and distributed across many nodes and sites. This dispersal provides both site and node failure protection. When compared to replication, erasure coding provides improved reliability at comparable storage costs.
- **Availability:** Availability can be defined as the ability to retrieve objects if Storage Nodes fail or become inaccessible. When compared to replication, erasure coding provides increased availability at comparable storage costs.
- **Storage efficiency:** For similar levels of availability and reliability, objects protected through erasure

coding consume less disk space than the same objects would if protected through replication. For example, a 10 MB object that is replicated to two sites consumes 20 MB of disk space (two copies), while an object that is erasure-coded across three sites with a 6+3 erasure-coding scheme only consumes 15 MB of disk space.



Disk space for erasure-coded objects is calculated as the object size plus the storage overhead. The storage overhead percentage is the number of parity fragments divided by the number of data fragments.

Disadvantages of erasure coding

When compared to replication, erasure coding has the following disadvantages:

- An increased number of Storage Nodes and sites is recommended, depending on the erasure-coding scheme. In contrast, if you replicate object data, you need only one Storage Node for each copy. See [Erasure-coding schemes for storage pools containing three or more sites](#) and [Erasure-coding schemes for one-site storage pools](#).
- Increased cost and complexity of storage expansions. To expand a deployment that uses replication, you add storage capacity in every location where object copies are made. To expand a deployment that uses erasure coding, you must consider both the erasure-coding scheme in use and how full existing Storage Nodes are. For example, if you wait until existing nodes are 100% full, you must add at least $k+m$ Storage Nodes, but if you expand when existing nodes are 70% full, you can add two nodes per site and still maximize usable storage capacity. For more information, see [Add storage capacity for erasure-coded objects](#).
- There are increased retrieval latencies when you use erasure coding across geographically distributed sites. The object fragments for an object that is erasure-coded and distributed across remote sites take longer to retrieve over WAN connections than an object that is replicated and available locally (the same site to which the client connects).
- When you use erasure coding across geographically distributed sites, there is higher WAN network traffic usage for retrievals and repairs, especially for frequently retrieved objects or for object repairs over WAN network connections.
- When you use erasure coding across sites, the maximum object throughput declines sharply as network latency between sites increases. This decrease is due to the corresponding decrease in TCP network throughput, which affects how quickly the StorageGRID system can store and retrieve object fragments.
- Higher usage of compute resources.

When to use erasure coding

Erasure coding is best suited for the following requirements:

- Objects greater than 1 MB in size.



Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

- Long-term or cold storage for infrequently retrieved content.
- High data availability and reliability.
- Protection against complete site and node failures.

- Storage efficiency.
- Single-site deployments that require efficient data protection with only a single erasure-coded copy rather than multiple replicated copies.
- Multiple-site deployments where the inter-site latency is less than 100 ms.

How object retention is determined

StorageGRID provides options for both grid administrators and individual tenant users to specify how long to store objects. In general, any retention instructions provided by a tenant user take precedence over the retention instructions provided by the grid administrator.

How tenant users control object retention

Tenant users can use these methods to control how long their objects are stored in StorageGRID:

- If the global S3 Object Lock setting is enabled for the grid, S3 tenant users can create buckets with S3 Object Lock enabled and then select a **Default retention period** for each bucket.
- If the global S3 Object Lock setting is enabled for the grid, S3 tenant users can create buckets with S3 Object Lock enabled and then use the S3 REST API to specify retain-until-date and legal hold settings for each object version added to that bucket.
 - An object version that is under a legal hold can't be deleted by any method.
 - Before an object version's retain-until-date is reached, that version can't be deleted by any method.
 - Objects in buckets with S3 Object Lock enabled are retained by ILM "forever." However, after its retain-until-date is reached, an object version can be deleted by a client request or the expiration of the bucket lifecycle. See [Manage objects with S3 Object Lock](#).
- S3 tenant users can add a lifecycle configuration to their buckets that specifies an Expiration action. If a bucket lifecycle exists, StorageGRID stores an object until the date or number of days specified in the Expiration action are met, unless the client deletes the object first. See [Create S3 lifecycle configuration](#).
- An S3 client can issue a delete object request. StorageGRID always prioritizes client delete requests over S3 bucket lifecycle or ILM when determining whether to delete or retain an object.

How grid administrators control object retention

Grid administrators can use these methods to control object retention:

- Set an S3 Object Lock maximum retention period for each tenant. Then, tenant users can set a default retention period for each of their buckets. The maximum retention period is also enforced on any newly ingested objects for that bucket (object's retain-until-date).
- Create ILM placement instructions to control how long objects are stored. When objects are matched by an ILM rule, StorageGRID stores those objects until the last time period in the ILM rule has elapsed. Objects are retained indefinitely if "forever" is specified for the placement instructions.
- Regardless of who controls how long objects are retained, ILM settings control what types of object copies (replicated or erasure-coded) are stored and where the copies are located (Storage Nodes or Cloud Storage Pools).

How S3 bucket lifecycle and ILM interact

When an S3 bucket lifecycle is configured, the lifecycle expiration actions override the ILM policy for objects

that match the lifecycle filter. As a result, an object might be retained on the grid even after any ILM instructions for placing the object have lapsed.

Examples for object retention

To better understand the interactions between S3 Object Lock, bucket lifecycle settings, client delete requests, and ILM, consider the following examples.

Example 1: S3 bucket lifecycle keeps objects longer than ILM

ILM

Store two copies for 1 year (365 days)

Bucket lifecycle

Expire objects in 2 years (730 days)

Result

StorageGRID stores the object for 730 days. StorageGRID uses the bucket lifecycle settings to determine whether to delete or retain an object.



If the bucket lifecycle specifies that objects should be kept longer than specified by ILM, StorageGRID continues to use the ILM placement instructions when determining the number and type of copies to store. In this example, two copies of the object will continue to be stored in StorageGRID from days 366 to 730.

Example 2: S3 bucket lifecycle expires objects before ILM

ILM

Store two copies for 2 years (730 days)

Bucket lifecycle

Expire objects in 1 year (365 days)

Result

StorageGRID deletes both copies of the object after day 365.

Example 3: Client delete overrides bucket lifecycle and ILM

ILM

Store two copies on Storage Nodes "forever"

Bucket lifecycle

Expire objects in 2 years (730 days)

Client delete request

Issued on day 400

Result

StorageGRID deletes both copies of the object on day 400 in response to the client delete request.

Example 4: S3 Object Lock overrides client delete request

S3 Object Lock

Retain-until-date for an object version is 2026-03-31. A legal hold is not in effect.

Compliant ILM rule

Store two copies on Storage Nodes "forever"

Client delete request

Issued on 2024-03-31

Result

StorageGRID will not delete the object version because the retain-until-date is still 2 years away.

How objects are deleted

StorageGRID can delete objects either in direct response to a client request or automatically as a result of the expiration of an S3 bucket lifecycle or the requirements of the ILM policy. Understanding the different ways that objects can be deleted and how StorageGRID handles delete requests can help you manage objects more effectively.

StorageGRID can use one of two methods to delete objects:

- Synchronous deletion: When StorageGRID receives a client delete request, all object copies are removed immediately. The client is informed that deletion was successful after the copies have been removed.
- Objects are queued for deletion: When StorageGRID receives a delete request, the object is queued for deletion and the client is informed immediately that deletion was successful. Object copies are removed later by background ILM processing.

When deleting objects, StorageGRID uses the method that optimizes delete performance, minimizes potential delete backlogs, and frees space most quickly.

The table summarizes when StorageGRID uses each method.

Method of performing deletion	When used
Objects are queued for deletion	When any of the following conditions are true: - Automatic object deletion has been triggered by one of the following events: - The expiration date or number of days in the lifecycle configuration for an S3 bucket is reached. - The last time period specified in an ILM rule elapses. Note: Objects in a bucket that has S3 Object Lock enabled can't be deleted if they are under a legal hold or if a retain-until-date has been specified but not yet met. - An S3 client requests deletion and one or more of these conditions is true: - Copies can't be deleted within 30 seconds because, for example, an object location is temporarily unavailable. - Background deletion queues are idle.
Objects are removed immediately (synchronous deletion)	When an S3 client makes a delete request and all of the following conditions are met: - All copies can be removed within 30 seconds. - Background deletion queues contain objects to process.

When S3 clients make delete requests, StorageGRID begins by adding objects to the delete queue. It then switches to performing synchronous deletion. Making sure that the background deletion queue has objects to process allows StorageGRID to process deletes more efficiently, especially for low concurrency clients, while helping to prevent client delete backlogs.

Time required to delete objects

The way that StorageGRID deletes objects can affect how the system appears to perform:

- When StorageGRID performs synchronous deletion, it can take StorageGRID up to 30 seconds to return a result to the client. This means that deletion can appear to be happening more slowly, even though copies are actually being removed more quickly than they are when StorageGRID queues objects for deletion.
- If you are closely monitoring delete performance during a bulk delete, you might notice that the deletion rate appears to be slow after a certain number of objects have been deleted. This change occurs when StorageGRID shifts from queuing objects for deletion to performing synchronous deletion. The apparent reduction in the deletion rate does not mean that object copies are being removed more slowly. On the contrary, it indicates that on average, space is now being freed more quickly.

If you are deleting large numbers of objects and your priority is to free space quickly, consider using a client request to delete objects rather than deleting them using ILM or other methods. In general, space is freed more quickly when deletion is performed by clients because StorageGRID can use synchronous deletion.

The amount of time required to free space after an object is deleted depends on several factors:

- Whether object copies are synchronously removed or are queued for removal later (for client delete requests).

- Other factors such as the number of objects in the grid or the availability of grid resources when object copies are queued for removal (for both client deletes and other methods).

How S3 versioned objects are deleted

When versioning is enabled for an S3 bucket, StorageGRID follows Amazon S3 behavior when responding to delete requests, whether those requests come from an S3 client, the expiration of an S3 bucket lifecycle, or the requirements of the ILM policy.

When objects are versioned, object delete requests don't delete the current version of the object and don't free space. Instead, an object delete request creates a zero-byte delete marker as the current version of the object, which makes the previous version of the object "noncurrent." An object delete marker becomes an expired object delete marker when it is the current version and there are no noncurrent versions.

Even though the object has not been removed, StorageGRID behaves as though the current version of the object is no longer available. Requests to that object return 404 Not Found. However, because noncurrent object data has not been removed, requests that specify a noncurrent version of the object can succeed.

To free space when deleting versioned objects, or to remove delete markers, use one of the following:

- **S3 client request:** Specify the object version ID in the S3 DELETE Object request (`DELETE /object?versionId=ID`). Keep in mind that this request only removes object copies for the specified version (the other versions are still taking up space).
- **Bucket lifecycle:** Use the `NoncurrentVersionExpiration` action in the bucket lifecycle configuration. When the number of `NoncurrentDays` specified is met, StorageGRID permanently removes all copies of noncurrent object versions. These object versions can't be recovered.

The `NewerNoncurrentVersions` action in the bucket lifecycle configuration specifies the number of noncurrent versions retained in a versioned S3 bucket. If there are more noncurrent versions than `NewerNoncurrentVersions` specifies, StorageGRID removes the older versions when the `NoncurrentDays` value has elapsed. The `NewerNoncurrentVersions` threshold overrides lifecycle rules provided by ILM, meaning that a noncurrent object with a version within the `NewerNoncurrentVersions` threshold is retained if ILM requests its deletion.

To remove expired object delete markers use the `Expiration` action with one of the following tags: `ExpiredObjectDeleteMarker`, `Days`, or `Date`.

- **ILM:** [Clone an active policy](#) and add two ILM rules to the new policy:
 - First rule: Use "Noncurrent time" as the Reference time to match the noncurrent versions of the object. In [Step 1 \(Enter details\) of the Create an ILM rule wizard](#), select **Yes** for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?"
 - Second rule: Use **Ingest time** to match the current version. The "Noncurrent time" rule must appear in the policy above the **Ingest time** rule.

To remove expired object delete markers, use an **Ingest time** rule to match the current delete markers. Delete markers are only removed when a **Time period of Days** has passed and the current delete maker has become expired (there are no non-current versions).

- **Delete objects in bucket:** Use the tenant manager to [delete all object versions](#), including delete markers, from a bucket.

When a versioned object is deleted, StorageGRID creates a zero-byte delete marker as the current version of the object. All objects and delete markers must be removed before a versioned bucket can be deleted.

- Delete markers created in StorageGRID 11.7 or earlier can only be removed through S3 client requests, they are not removed by ILM, bucket lifecycle rules, or Delete objects in bucket operations.
- Delete markers from a bucket that was created in StorageGRID 11.8 or later can be removed by ILM, bucket lifecycle rules, Delete objects in bucket operations, or an explicit S3 client deletion.

Related information

- [Use S3 REST API](#)
- [Example 4: ILM rules and policy for S3 versioned objects](#)

Create and assign storage grades

Storage grades identify the type of storage used by a Storage Node. You can create storage grades if you want ILM rules to place certain objects on certain Storage Nodes.

Before you begin

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

About this task

When you first install StorageGRID, the **Default** storage grade is automatically assigned to every Storage Node in your system. As required, you can optionally define custom storage grades and assign them to different Storage Nodes.

Using custom storage grades allows you to create ILM storage pools that contain only a specific type of Storage Node. For example, you might want certain objects to be stored on your fastest Storage Nodes, such as StorageGRID all-flash storage appliances.



Storage Nodes can be configured during installation to contain only object metadata and not object data. Metadata-only Storage Nodes can't be assigned a storage grade. For more information, see [Types of Storage Nodes](#).

If storage grade is not a concern (for example, all Storage Nodes are identical), you can skip this procedure and use the **includes all storage grades** selection for the storage grade when you [create storage pools](#). Using this selection ensures that the storage pool will include every Storage Node at the site, regardless of its storage grade.



Don't create more storage grades than necessary. For example, don't create a storage grade for each Storage Node. Instead, assign each storage grade to two or more nodes. Storage grades assigned to only one node can cause ILM backlogs if that node becomes unavailable.

Steps

1. Select **ILM > Storage grades**.
2. Define custom storage grades:
 - a. For each custom storage grade you want to add, select **Insert**  to add a row.
 - b. Enter a descriptive label.



Storage Grades

Updated: 2017-05-26 11:22:39 MDT

Storage Grade Definitions

Storage Grade	Label	Actions
0	Default	
1	disk	

Storage Grades

LDR	Storage Grade	Actions
Data Center 1/DC1-S1/LDR	Default	
Data Center 1/DC1-S2/LDR	Default	
Data Center 1/DC1-S3/LDR	Default	
Data Center 2/DC2-S1/LDR	Default	
Data Center 2/DC2-S2/LDR	Default	
Data Center 2/DC2-S3/LDR	Default	
Data Center 3/DC3-S1/LDR	Default	
Data Center 3/DC3-S2/LDR	Default	
Data Center 3/DC3-S3/LDR	Default	

Apply Changes

c. Select **Apply Changes**.

d. Optionally, if you need to modify a saved label, select **Edit** and select **Apply Changes**.



You can't delete storage grades.

3. Assign new storage grades to Storage Nodes:

- Locate the Storage Node in the LDR list, and select its **Edit** icon .
- Select the appropriate storage grade from the list.



LDR	Storage Grade	Actions
Data Center 1/DC1-S1/LDR	Default	
Data Center 1/DC1-S2/LDR	Default disk	
Data Center 1/DC1-S3/LDR	Default	
Data Center 2/DC2-S1/LDR	Default	
Data Center 2/DC2-S2/LDR	Default	
Data Center 2/DC2-S3/LDR	Default	
Data Center 3/DC3-S1/LDR	Default	
Data Center 3/DC3-S2/LDR	Default	
Data Center 3/DC3-S3/LDR	Default	

Apply Changes



Assign a storage grade to a given Storage Node only once. A Storage Node recovered from failure maintains the previously assigned storage grade. Don't change this assignment after the ILM policy is activated. If the assignment is changed, data is stored based on the new storage grade.

c. Select **Apply Changes**.

Use storage pools

What is a storage pool?

A storage pool is a logical grouping of Storage Nodes.

When you install StorageGRID, one storage pool per site is automatically created. You can configure additional storage pools as needed for your storage requirements.



Storage Nodes can be configured during installation to contain object data and object metadata, or only object metadata. Metadata-only Storage Nodes can't be used in storage pools. For more information, see [Types of Storage Nodes](#).

Storage pools have two attributes:

- **Storage grade:** For Storage Nodes, the relative performance of backing storage.
- **Site:** The data center where objects will be stored.

Storage pools are used in ILM rules to determine where object data is stored and the type of storage used. When you configure ILM rules for replication, you select one or more storage pools.

Guidelines for creating storage pools

Configure and use storage pools to protect against data loss by distributing data across multiple sites. Replicated copies and erasure-coded copies require different storage pool

configurations.

See [Examples of enabling site-loss protection using replication and erasure coding](#).

Guidelines for all storage pools

- Keep storage pool configurations as simple as possible. Don't create more storage pools than necessary.
- Create storage pools with as many nodes as possible. Each storage pool should contain two or more nodes. A storage pool with insufficient nodes can cause ILM backlogs if a node becomes unavailable.
- Avoid creating or using storage pools that overlap (contain one or more of the same nodes). If storage pools overlap, more than one copy of object data might be saved on the same node.
- In general, don't use the All Storage Nodes storage pool (StorageGRID 11.6 and earlier) or the All Sites site. These items are automatically updated to include any new sites you add in an expansion, which might not be the behavior you want.

Guidelines for storage pools used for replicated copies

- For site-loss protection using [replication](#), specify one or more site-specific storage pools in the [placement instructions for each ILM rule](#).

One storage pool is automatically created for each site during StorageGRID installation.

Using a storage pool for each site ensures that replicated object copies are placed exactly where you expect (for example, one copy of every object at each site for site-loss protection).

- If you add a site in an expansion, create a new storage pool that contains only the new site. Then, [update ILM rules](#) to control which objects are stored on the new site.
- If the number of copies is less than the number of storage pools, the system distributes the copies to balance disk usage among the pools.
- If the storage pools overlap (contain the same Storage Nodes), all copies of the object might be saved at only one site. You must ensure that the selected storage pools don't contain the same Storage Nodes.

Guidelines for storage pools used for erasure-coded copies

- For site-loss protection using [erasure coding](#), create storage pools that consist of at least three sites. If a storage pool includes only two sites, you can't use that storage pool for erasure coding. No erasure-coding schemes are available for a storage pool that has two sites.
- The number of Storage Nodes and sites contained in the storage pool determine which [erasure-coding schemes](#) are available.
- If possible, a storage pool should include more than the minimum number of Storage Nodes required for the erasure-coding scheme you select. For example, if you use a 6+3 erasure-coding scheme, you must have at least nine Storage Nodes. However, having at least one additional Storage Node per site is recommended.
- Distribute Storage Nodes across sites as evenly as possible. For example, to support a 6+3 erasure-coding scheme, configure a storage pool that includes at least three Storage Nodes at three sites.
- If you have high throughput requirements, using a storage pool that includes multiple sites is not recommended if the network latency between sites is greater than 100 ms. As latency increases, the rate at which StorageGRID can create, place, and retrieve object fragments decreases sharply due to the decrease in TCP network throughput.

The decrease in throughput affects the maximum achievable rates of object ingest and retrieval (when

Balanced or Strict are selected as the ingest behavior) or could lead to ILM queue backlogs (when Dual commit is selected as the ingest behavior). See [ILM rule ingest behavior](#).



If your grid includes only one site, you are prevented from using the All Storage Nodes storage pool (StorageGRID 11.6 and earlier) or the All Sites site in an erasure-coding profile. This behavior prevents the profile from becoming invalid if a second site is added.

Enable site-loss protection

If your StorageGRID deployment includes more than one site, you can use replication and erasure coding with appropriately configured storage pools to enable site-loss protection.

Replication and erasure coding require different storage pool configurations:

- To use replication for site-loss protection, use the site-specific storage pools that are automatically created during StorageGRID installation. Then create ILM rules with [placement instructions](#) that specify multiple storage pools so that one copy of each object will be placed at each site.
- To use erasure coding for site-loss protection, [create storage pools that consist of multiple sites](#). Then create ILM rules that use one storage pool consisting of multiple sites and any available erasure-coding schema.



When configuring your StorageGRID deployment for site-loss protection, you must also take into account the effects of [ingest options](#) and [consistency](#).


Replication example

By default, one storage pool is created for each site during StorageGRID installation. Having storage pools that consist of only one site enables you to configure ILM rules that use replication for site-loss protection. In this example:

- Storage pool 1 contains Site 1
- Storage pool 2 contains Site 2
- The ILM rule contains two placements:
 - Store objects by replicating 1 copy at Site 1
 - Store objects by replicating 1 copy at Site 2

ILM rule placements:

The screenshot shows a configuration interface for ILM rule placements. It consists of two rows of controls. The first row is for Site 1 and the second row is for Site 2. Each row has a label 'Store objects by' followed by a dropdown menu set to 'replicating', a numeric input field set to '1', and a label 'copies at' followed by a site selection box (Site 1 for the first row, Site 2 for the second row). Each site selection box has a close button (X), an edit button (pencil), and a delete button (X).



If one site is lost, copies of the objects are available at the other site.

Erasure coding example

Having storage pools that consist of more than one site per storage pool enables you to configure ILM rules that use erasure coding for site-loss protection. In this example:

- Storage pool 1 contains Sites 1 through 3
- The ILM rule contains one placement: Store objects by erasure coding using a 4+2 EC scheme at Storage pool 1, which contains three sites

ILM rule placements:

Store objects by	erasure coding ▼	using 4+2 EC at Storage pool 1 (3 sites)		


In this example:

- The ILM rule uses a 4+2 erasure-coding scheme.
- Each object is sliced into four equal data fragments, and two parity fragments are computed from the object data.
- Each of the six fragments is stored on a different node across three data center sites to provide data protection for node failures or site loss.




Erasure coding is allowed in storage pools containing any number of sites *except* two sites.

ILM rule using 4+2 erasure-coding scheme:



If one site is lost, data can still be recovered:



Create a storage pool

You create storage pools to determine where the StorageGRID system stores object data and the type of storage used. Each storage pool includes one or more sites and one or more storage grades.



When you install StorageGRID 11.9 on a new grid, storage pools are automatically created for each site. However, if you initially installed StorageGRID 11.6 or earlier, storage pools aren't automatically created for each site.

If you want to create Cloud Storage Pools to store object data outside of your StorageGRID system, see the [information about using Cloud Storage Pools](#).

Before you begin

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

About this task

Storage pools determine where object data is stored. The number of storage pools you need depends on the number of sites in your grid and on the types of copies you want: replicated or erasure-coded.

- For replication and single-site erasure coding, create a storage pool for each site. For example, if you want to store replicated object copies at three sites, create three storage pools.
- For erasure coding at three or more sites, create one storage pool that includes an entry for each site. For example, if you want to erasure code objects across three sites, create one storage pool.



Don't include the All Sites site in a storage pool that will be used in an erasure-coding profile. Instead, add a separate entry to the storage pool for each site that will store erasure-coded data. See [this step](#) for an example.

- If you have more than one storage grade, don't create a storage pool that includes different storage grades at a single site. See the [Guidelines for creating storage pools](#).

Steps

1. Select **ILM > Storage pools**.

The Storage pools tab lists all defined storage pools.



For new installations of StorageGRID 11.6 or earlier, the All Storage Nodes storage pool is automatically updated whenever you add new data center sites. Don't use this pool in ILM rules.

2. To create a new storage pool, select **Create**.
3. Enter a unique name for the storage pool. Use a name that will be easy to identify when you configure erasure-coding profiles and ILM rules.
4. From the **Site** drop-down list, select a site for this storage pool.

When you select a site, the number of Storage Nodes in the table are automatically updated.

In general, don't use the All Sites site in any storage pool. ILM rules that use an All Sites storage pool place objects at any available site, giving you less control of object placement. Also, an All Sites storage pool uses the Storage Nodes at a new site immediately, which might not be the behavior you expect.

5. From the **Storage grade** drop-down list, select the type of storage that will be used if an ILM rule uses this storage pool.

The storage grade, *includes all storage grades*, includes all Storage Nodes at the selected site. If you created additional storage grades for the Storage Nodes in your grid, they are listed in the drop-down.

6. If you want to use the storage pool in a multi-site erasure-coding profile, select **Add more nodes** to add an entry for each site to the storage pool.



You are warned if you add more than one entry with different storage grades for a site.

To remove an entry, select the delete icon .

7. When you are satisfied with your selections, select **Save**.

The new storage pool is added to the list.

View storage pool details

You can view the details of a storage pool to determine where the storage pool is used and to see which nodes and storage grades are included.

Before you begin

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

Steps

1. Select **ILM > Storage pools**.

The Storage pools table includes the following information for each storage pool that includes Storage Nodes:

- **Name:** The unique display name of the storage pool.
- **Node count:** The number of nodes in the storage pool.
- **Storage usage:** The percentage of the total usable space that has been used for object data on this node. This value does not include object metadata.
- **Total capacity:** The size of the storage pool, which equals the total amount of usable space for object data for all nodes in the storage pool.
- **ILM usage:** How the storage pool is currently being used. A storage pool might be unused or it might be used in one or more ILM rules, erasure-coding profiles, or both.

2. To view details for a specific storage pool, select its name.

The details page for the storage pool appears.

3. View the **Nodes** tab to learn about the Storage Nodes included in the storage pool.

The table includes the following information for each node:

- Node name
- Site name
- Storage grade
- Storage usage: The percentage of the total usable space for object data that has been used for the Storage Node.



The same Storage usage (%) value is also shown in the Storage Used - Object Data chart for each Storage Node (select **NODES > Storage Node > Storage**).

4. View the **ILM usage** tab to determine if the storage pool is currently being used in any ILM rules or erasure-coding profiles.

5. Optionally, go to the **ILM rules page** to learn about and manage any rules that use the storage pool.

See the [instructions for working with ILM rules](#).

Edit storage pool

You can edit a storage pool to change its name or to update sites and storage grades.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- You have reviewed the [guidelines for creating storage pools](#).
- If you plan to edit a storage pool that is used by a rule in the active ILM policy, you have considered how your changes will affect object data placement.

About this task

If you are adding a new site or storage grade to a storage pool that is used in the active ILM policy, be aware that the Storage Nodes in the new site or storage grade will not be used automatically. To force StorageGRID to use a new site or storage grade, you must activate a new ILM policy after saving the edited storage pool.

Steps

1. Select **ILM > Storage pools**.
2. Select the checkbox for the storage pool you want to edit.

You can't edit the All Storage Nodes storage pool (StorageGRID 11.6 and earlier).

3. Select **Edit**.
4. As required, change the storage pool name.
5. As required, select other sites and storage grades.

You are prevented from changing the site or storage grade if the storage pool is used in an erasure-coding profile and the change would cause the erasure-coding scheme to become invalid. For example, if a storage pool used in a erasure-coding profile currently includes a storage grade with only one site, you are prevented from using a storage grade with two sites because the change would make the erasure-coding scheme invalid.



Adding or removing sites from an existing storage pool won't move any existing erasure-encoded data. If you want to move the existing data from the site, you must create a new storage pool and EC profile to re-encode the data.

6. Select **Save**.

After you finish

If you added a new site or storage grade to a storage pool used in the active ILM policy, activate a new ILM policy to force StorageGRID to use the new site or storage grade. For example, clone your existing ILM policy and then activate the clone. See [Work with ILM rules and ILM policies](#).

Remove a storage pool

You can remove a storage pool that is not being used.

Before you begin

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

Steps

1. Select **ILM > Storage pools**.
2. Look at the ILM usage column in the table to determine whether you can remove the storage pool.

You can't remove a storage pool if it is being used in an ILM rule or in an erasure-coding profile. As required, select **storage pool name > ILM usage** to determine where the storage pool is used.

3. If the storage pool you want to remove is not being used, select the checkbox.
4. Select **Remove**.
5. Select **OK**.

Use Cloud Storage Pools

What is a Cloud Storage Pool?

A Cloud Storage Pool lets you use ILM to move object data outside of your StorageGRID system. For example, you might want to move infrequently accessed objects to lower-cost cloud storage, such as Amazon S3 Glacier, S3 Glacier Deep Archive, Google Cloud, or the Archive access tier in Microsoft Azure Blob storage. Or, you might want to maintain a cloud backup of StorageGRID objects to enhance disaster recovery.

From an ILM perspective, a Cloud Storage Pool is similar to a storage pool. To store objects in either location, you select the pool when creating the placement instructions for an ILM rule. However, while storage pools consist of Storage Nodes within the StorageGRID system, a Cloud Storage Pool consists of an external bucket (S3) or container (Azure Blob storage).

The table compares storage pools to Cloud Storage Pools and shows the high-level similarities and differences.

	Storage pool	Cloud Storage Pool
How is it created?	Using the ILM > Storage pools option in Grid Manager.	Using the ILM > Storage pools > Cloud Storage Pools option in Grid Manager. You must set up the external bucket or container before you can create the Cloud Storage Pool.
How many pools can you create?	Unlimited.	Up to 10.

	Storage pool	Cloud Storage Pool
Where are objects stored?	On one or more Storage Nodes within StorageGRID.	In an Amazon S3 bucket, Azure Blob storage container, or Google Cloud that is external to the StorageGRID system. If the Cloud Storage Pool is an Amazon S3 bucket: - You can optionally configure a bucket lifecycle to transition objects to low-cost, long-term storage, such as Amazon S3 Glacier or S3 Glacier Deep Archive. The external storage system must support the Glacier storage class and the S3 RestoreObject API. - You can create Cloud Storage Pools for use with AWS Commercial Cloud Services (C2S), which supports the AWS Secret Region. If the Cloud Storage Pool is an Azure Blob storage container, StorageGRID transitions the object to the Archive tier. Note: In general, don't configure Azure Blob storage lifecycle management for the container used for a Cloud Storage Pool. RestoreObject operations on objects in the Cloud Storage Pool can be affected by the configured lifecycle.
What controls object placement?	An ILM rule in the active ILM policies.	An ILM rule in the active ILM policies.
Which data protection method is used?	Replication or erasure coding.	Replication.
How many copies of each object are allowed?	Multiple.	One copy in the Cloud Storage Pool and, optionally, one or more copies in StorageGRID. Note: You can't store an object in more than one Cloud Storage Pool at any given time.
What are the advantages?	Objects are quickly accessible at any time.	Low-cost storage. Note: FabricPool data can't be tiered to Cloud Storage Pools.

Lifecycle of a Cloud Storage Pool object

Before implementing Cloud Storage Pools, review the lifecycle of objects that are stored in each type of Cloud Storage Pool.

S3: Lifecycle of a Cloud Storage Pool object

The steps describe the lifecycle stages of an object that is stored in an S3 Cloud Storage Pool.



"Glacier" refers to both the Glacier storage class and the Glacier Deep Archive storage class, with one exception: the Glacier Deep Archive storage class does not support the Expedited restore tier. Only Bulk or Standard retrieval is supported.



The Google Cloud Platform (GCP) supports object retrieval from long-term storage without requiring a POST Restore operation.

1. Object stored in StorageGRID

To start the lifecycle, a client application stores an object in StorageGRID.

2. Object moved to S3 Cloud Storage Pool

- When the object is matched by an ILM rule that uses an S3 Cloud Storage Pool as its placement location, StorageGRID moves the object to the external S3 bucket specified by the Cloud Storage Pool.
- When the object has been moved to the S3 Cloud Storage Pool, the client application can retrieve it using an S3 GetObject request from StorageGRID, unless the object has been transitioned to Glacier storage.

3. Object transitioned to Glacier (non-retrievable state)

- Optionally, the object can be transitioned to Glacier storage. For example, the external S3 bucket might use lifecycle configuration to transition an object to Glacier storage immediately or after some number of days.



If you want to transition objects, you must create a lifecycle configuration for the external S3 bucket, and you must use a storage solution that implements the Glacier storage class and supports the S3 RestoreObject API.

- During the transition, the client application can use an S3 HeadObject request to monitor the object's status.

4. Object restored from Glacier storage

If an object has been transitioned to Glacier storage, the client application can issue an S3 RestoreObject request to restore a retrievable copy to the S3 Cloud Storage Pool. The request specifies how many days the copy should be available in the Cloud Storage Pool and the data-access tier to use for the restore operation (Expedited, Standard, or Bulk). When the expiration date of the retrievable copy is reached, the copy is automatically returned to a non-retrievable state.



If one or more copies of the object also exist on Storage Nodes within StorageGRID, there is no need to restore the object from Glacier by issuing a RestoreObject request. Instead, the local copy can be retrieved directly, using a GetObject request.

5. Object retrieved

Once an object has been restored, the client application can issue a GetObject request to retrieve the restored object.

Azure: Lifecycle of a Cloud Storage Pool object

The steps describe the lifecycle stages of an object that is stored in an Azure Cloud Storage Pool.

1. Object stored in StorageGRID

To start the lifecycle, a client application stores an object in StorageGRID.

2. Object moved to Azure Cloud Storage Pool

When the object is matched by an ILM rule that uses an Azure Cloud Storage Pool as its placement location, StorageGRID moves the object to the external Azure Blob storage container specified by the Cloud Storage Pool.

3. Object transitioned to Archive tier (non-retrievable state)

Immediately after moving the object to the Azure Cloud Storage Pool, StorageGRID automatically transitions the object to the Azure Blob storage Archive tier.

4. Object restored from Archive tier

If an object has been transitioned to the Archive tier, the client application can issue an S3 RestoreObject request to restore a retrievable copy to the Azure Cloud Storage Pool.

When StorageGRID receives the RestoreObject, it temporarily transitions the object to the Azure Blob storage Cool tier. As soon as the expiration date in the RestoreObject request is reached, StorageGRID transitions the object back to the Archive tier.



If one or more copies of the object also exist on Storage Nodes within StorageGRID, there is no need to restore the object from the Archive access tier by issuing a RestoreObject request. Instead, the local copy can be retrieved directly, using a GetObject request.

5. Object retrieved

Once an object has been restored to the Azure Cloud Storage Pool, the client application can issue a GetObject request to retrieve the restored object.

Related information

[Use S3 REST API](#)

When to use Cloud Storage Pools

Using Cloud Storage Pools, you can back up or tier data to an external location. Additionally, you can back up or tier data to more than one cloud.

Back up StorageGRID data to external location

You can use a Cloud Storage Pool to back up StorageGRID objects to an external location.

If the copies in StorageGRID are inaccessible, the object data in the Cloud Storage Pool can be used to serve client requests. However, you might need to issue S3 RestoreObject request to access the backup object copy in the Cloud Storage Pool.

The object data in a Cloud Storage Pool can also be used to recover data lost from StorageGRID because of a

storage volume or Storage Node failure. If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID temporarily restores the object and creates a new copy on the recovered Storage Node.

To implement a backup solution:

1. Create a single Cloud Storage Pool.
2. Configure an ILM rule that simultaneously stores object copies on Storage Nodes (as replicated or erasure-coded copies) and a single object copy in the Cloud Storage Pool.
3. Add the rule to your ILM policy. Then, simulate and activate the policy.

Tier data from StorageGRID to external location

You can use a Cloud Storage Pool to store objects outside of the StorageGRID system. For example, suppose you have a large number of objects that you need to retain, but you expect to access those objects rarely, if ever. You can use a Cloud Storage Pool to tier the objects to lower-cost storage and to free up space in StorageGRID.

To implement a tiering solution:

1. Create a single Cloud Storage Pool.
2. Configure an ILM rule that moves rarely used objects from Storage Nodes to the Cloud Storage Pool.
3. Add the rule to your ILM policy. Then, simulate and activate the policy.

Maintain multiple cloud endpoints

You can configure multiple Cloud Storage Pool endpoints if you want to tier or back up object data to more than one cloud. The filters in your ILM rules let you specify which objects are stored in each Cloud Storage Pool. For example, you might want to store objects from some tenants or buckets in Amazon S3 Glacier and objects from other tenants or buckets in Azure Blob storage. Or, you might want to move data between Amazon S3 Glacier and Azure Blob storage.



When using multiple Cloud Storage Pool endpoints, keep in mind that an object can be stored in only one Cloud Storage Pool at a time.

To implement multiple cloud endpoints:

1. Create up to 10 Cloud Storage Pools.
2. Configure ILM rules to store the appropriate object data at the appropriate time in each Cloud Storage Pool. For example, store objects from bucket A in Cloud Storage Pool A, and store objects from bucket B in Cloud Storage Pool B. Or, store objects in Cloud Storage Pool A for some amount of time and then move them to Cloud Storage Pool B.
3. Add the rules to your ILM policy. Then, simulate and activate the policy.

Considerations for Cloud Storage Pools

If you plan to use a Cloud Storage Pool to move objects out of the StorageGRID system, you must review the considerations for configuring and using Cloud Storage Pools.

General considerations

- In general, cloud archival storage, such as Amazon S3 Glacier or Azure Blob storage, is an inexpensive

place to store object data. However, the costs to retrieve data from cloud archival storage are relatively high. To achieve the lowest overall cost, you must consider when and how often you will access the objects in the Cloud Storage Pool. Using a Cloud Storage Pool is recommended only for content that you expect to access infrequently.

- Using Cloud Storage Pools with FabricPool is not supported because of the added latency to retrieve an object from the Cloud Storage Pool target.
- Objects with S3 Object Lock enabled can't be placed in Cloud Storage Pools.
- If the destination S3 bucket for a Cloud Storage Pool has S3 Object Lock enabled, the attempt to configure bucket replication (PutBucketReplication) will fail with an AccessDenied error.
- The following platform, authentication, and protocol combinations with S3 Object lock aren't supported for Cloud Storage Pools:
 - **Platforms:** Google Cloud Platform and Azure
 - **Authentication types:** IAM Roles Anywhere and anonymous access
 - **Protocol:** HTTP

Considerations for the ports used for Cloud Storage Pools

To ensure that the ILM rules can move objects to and from the specified Cloud Storage Pool, you must configure the network or networks that contain your system's Storage Nodes. You must ensure that the following ports can communicate with the Cloud Storage Pool.

By default, Cloud Storage Pools use the following ports:

- **80:** For endpoint URIs that begin with http
- **443:** For endpoint URIs that begin with https

You can specify a different port when you create or edit a Cloud Storage Pool.

If you use a non-transparent proxy server, you must also [configure a storage proxy](#) to allow messages to be sent to external endpoints, such as an endpoint on the internet.

Considerations for costs

Access to storage in the cloud using a Cloud Storage Pool requires network connectivity to the cloud. You must consider the cost of the network infrastructure you will use to access the cloud and provision it appropriately, based on the amount of data you expect to move between StorageGRID and the cloud using the Cloud Storage Pool.

When StorageGRID connects to the external Cloud Storage Pool endpoint, it issues various requests to monitor connectivity and to ensure it can perform the required operations. While some additional costs will be associated with these requests, the cost of monitoring a Cloud Storage Pool should only be a small fraction of the overall cost of storing objects in S3 or Azure.

More significant costs might be incurred if you need to move objects from an external Cloud Storage Pool endpoint back to StorageGRID. Objects might be moved back to StorageGRID in either of these cases:

- The only copy of the object is in a Cloud Storage Pool and you decide to store the object in StorageGRID instead. In this case, you reconfigure your ILM rules and policy. When ILM evaluation occurs, StorageGRID issues multiple requests to retrieve the object from the Cloud Storage Pool. StorageGRID then creates the specified number of replicated or erasure-coded copies locally. After the object is moved back to StorageGRID, the copy in the Cloud Storage Pool is deleted.

- Objects are lost because of Storage Node failure. If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID temporarily restores the object and creates a new copy on the recovered Storage Node.



When objects are moved back to StorageGRID from a Cloud Storage Pool, StorageGRID issues multiple requests to the Cloud Storage Pool endpoint for each object. Before moving large numbers of objects, contact technical support for help in estimating the time frame and associated costs.

S3: Permissions required for the Cloud Storage Pool bucket

The policies for the external S3 bucket used for a Cloud Storage Pool must grant StorageGRID permission to move an object to the bucket, get an object's status, restore an object from Glacier storage when required, and more. Ideally, StorageGRID should have full-control access to the bucket (`s3:*`); however, if this is not possible, the bucket policy must grant the following S3 permissions to StorageGRID:

- `s3:AbortMultipartUpload`
- `s3:DeleteObject`
- `s3:GetObject`
- `s3:ListBucket`
- `s3:ListBucketMultipartUploads`
- `s3:ListMultipartUploadParts`
- `s3:PutObject`
- `s3:RestoreObject`

S3: Considerations for the external bucket's lifecycle

The movement of objects between StorageGRID and the external S3 bucket specified in the Cloud Storage Pool is controlled by ILM rules and the active ILM policies in StorageGRID. In contrast, the transition of objects from the external S3 bucket specified in the Cloud Storage Pool to Amazon S3 Glacier or S3 Glacier Deep Archive (or to a storage solution that implements the Glacier storage class) is controlled by that bucket's lifecycle configuration.

If you want to transition objects from the Cloud Storage Pool, you must create the appropriate lifecycle configuration on the external S3 bucket, and you must use a storage solution that implements the Glacier storage class and supports the S3 `RestoreObject` API.

For example, suppose you want all objects that are moved from StorageGRID to the Cloud Storage Pool to be transitioned to Amazon S3 Glacier storage immediately. You would create a lifecycle configuration on the external S3 bucket that specifies a single action (**Transition**) as follows:


```
<LifecycleConfiguration>
  <Rule>
    <ID>Transition Rule</ID>
    <Filter>
      <Prefix></Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>0</Days>
      <StorageClass>GLACIER</StorageClass>
    </Transition>
  </Rule>
</LifecycleConfiguration>
```

This rule would transition all bucket objects to Amazon S3 Glacier on the day they were created (that is, on the day they were moved from StorageGRID to the Cloud Storage Pool).



When configuring the external bucket's lifecycle, never use **Expiration** actions to define when objects expire. Expiration actions cause the external storage system to delete expired objects. If you later attempt to access an expired object from StorageGRID, the deleted object will not be found.

If you want to transition objects in the Cloud Storage Pool to S3 Glacier Deep Archive (instead of to Amazon S3 Glacier), specify `<StorageClass>DEEP_ARCHIVE</StorageClass>` in the bucket lifecycle. However, be aware that you can't use the Expedited tier to restore objects from S3 Glacier Deep Archive.

Azure: Considerations for Access tier

When you configure an Azure storage account, you can set the default Access tier to Hot or Cool. When creating a storage account for use with a Cloud Storage Pool, you should use the Hot tier as the default tier. Even though StorageGRID immediately sets the tier to Archive when it moves objects to the Cloud Storage Pool, using a default setting of Hot ensures that you will not be charged an early deletion fee for objects removed from the Cool tier before the 30-day minimum.

Azure: Lifecycle management not supported

Don't use Azure Blob storage lifecycle management for the container used with a Cloud Storage Pool. The lifecycle operations might interfere with Cloud Storage Pool operations.

Related information

[Create a Cloud Storage Pool](#)

Compare Cloud Storage Pools and CloudMirror replication

As you begin using Cloud Storage Pools, it might be helpful to understand the similarities and differences between Cloud Storage Pools and the StorageGRID CloudMirror replication service.

	Cloud Storage Pool	CloudMirror replication service
What is the primary purpose?	Acts as an archive target. The object copy in the Cloud Storage Pool can be the only copy of the object, or it can be an additional copy. That is, instead of keeping two copies onsite, you can keep one copy within StorageGRID and send a copy to the Cloud Storage Pool.	Enables a tenant to automatically replicate objects from a bucket in StorageGRID (source) to an external S3 bucket (destination). Creates an independent copy of an object in an independent S3 infrastructure.
How is it set up?	Defined in the same way as storage pools, using the Grid Manager or the Grid Management API. Can be selected as the placement location in an ILM rule. While a storage pool consists of a group of Storage Nodes, a Cloud Storage Pool is defined using a remote S3 or Azure endpoint (IP address, credentials, and so on).	A tenant user configures CloudMirror replication by defining a CloudMirror endpoint (IP address, credentials, and so on) using the Tenant Manager or the S3 API. After the CloudMirror endpoint is set up, any bucket owned by that tenant account can be configured to point to the CloudMirror endpoint.
Who is responsible for setting it up?	Typically, a grid administrator	Typically, a tenant user
What is the destination?	- Any compatible S3 infrastructure (including Amazon S3) - Azure Blob Archive tier - Google Cloud Platform (GCP)	- Any compatible S3 infrastructure (including Amazon S3) - Google Cloud Platform (GCP)
What causes objects to be moved to the destination?	One or more ILM rules in the active ILM policies. The ILM rules define which objects StorageGRID moves to the Cloud Storage Pool and when the objects are moved.	The act of ingesting a new object into a source bucket that has been configured with a CloudMirror endpoint. Objects that existed in the source bucket before the bucket was configured with the CloudMirror endpoint aren't replicated, unless they are modified.
How are objects retrieved?	Applications must make requests to StorageGRID to retrieve objects that have been moved to a Cloud Storage Pool. If the only copy of an object has been transitioned to archival storage, StorageGRID manages the process of restoring the object so it can be retrieved.	Because the mirrored copy in the destination bucket is an independent copy, applications can retrieve the object by making requests either to StorageGRID or to the S3 destination. For example, suppose you use CloudMirror replication to mirror objects to a partner organization. The partner can use its own applications to read or update objects directly from the S3 destination. Using StorageGRID is not required.

	Cloud Storage Pool	CloudMirror replication service
Can you read from the destination directly?	No. Objects moved to a Cloud Storage Pool are managed by StorageGRID. Read requests must be directed to StorageGRID (and StorageGRID will be responsible for retrieval from Cloud Storage Pool).	Yes, because the mirrored copy is an independent copy.
What happens if an object is deleted from the source?	The object is also deleted from the Cloud Storage Pool.	The delete action is not replicated. A deleted object no longer exists in the StorageGRID bucket, but it continues to exist in the destination bucket. Similarly, objects in the destination bucket can be deleted without affecting the source.
How do you access objects after a disaster (StorageGRID system not operational)?	Failed StorageGRID nodes must be recovered. During this process, copies of replicated objects might be restored using the copies in the Cloud Storage Pool.	The object copies in the CloudMirror destination are independent of StorageGRID, so they can be accessed directly before the StorageGRID nodes are recovered.

Create a Cloud Storage Pool

A Cloud Storage Pool specifies a single external Amazon S3 bucket or other S3-compatible provider or an Azure Blob storage container.

When you create a Cloud Storage Pool, you specify the name and location of the external bucket or container that StorageGRID will use to store objects, the cloud provider type (Amazon S3/GCP or Azure Blob storage), and the information StorageGRID needs to access the external bucket or container.

StorageGRID validates the Cloud Storage Pool as soon as you save it, so you must ensure that the bucket or container specified in the Cloud Storage Pool exists and is reachable.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [required access permissions](#).
- You have reviewed the [considerations for Cloud Storage Pools](#).
- The external bucket or container referenced by the Cloud Storage Pool already exists, and you have the [service endpoint information](#).
- To access the bucket or container, you have the [account information for the authentication type](#) you will choose.

Steps

1. Select **ILM > Storage pools > Cloud Storage Pools**.
2. Select **Create**, then enter the following information:

Field	Description
Cloud Storage Pool name	A name that briefly describes the Cloud Storage Pool and its purpose. Use a name that will be easy to identify when you configure ILM rules.
Provider type	Which cloud provider you will use for this Cloud Storage Pool: • Amazon S3/GCP: Select this option for an Amazon S3, Commercial Cloud Services (C2S) S3, Google Cloud Platform (GCP), or other S3-compatible provider. • Azure Blob Storage
Bucket or container	The name of the external S3 bucket or Azure container. You can't change this value after the Cloud Storage Pool is saved.

- Based on your Provider type selection, enter the service endpoint information.

Amazon S3/GCP

- a. For the protocol, select either HTTPS or HTTP.



Don't use HTTP connections for sensitive data.

- b. Enter the hostname. Example:

`s3-aws-region.amazonaws.com`

- c. Select the URL style:

Option	Description
Auto-detect	Attempt to automatically detect which URL style to use, based on the information provided. For example, if you specify an IP address, StorageGRID will use a path-style URL. Select this option only if you don't know which specific style to use.
Virtual-hosted-style	Use a virtual-hosted-style URL to access the bucket. Virtual-hosted-style URLs include the bucket name as part of the domain name. Example: <code>https://bucket-name.s3.company.com/key-name</code>
Path-style	Use a path-style URL to access the bucket. Path-style URLs include the bucket name at the end. Example: <code>https://s3.company.com/bucket-name/key-name</code> Note: The path-style URL option is not recommended and will be deprecated in a future release of StorageGRID.

- d. Optionally, enter the port number, or use the default port: 443 for HTTPS or 80 for HTTP.

Azure Blob Storage

- a. Using one of the following formats, enter the URI for the service endpoint.

- `https://host:port`
- `http://host:port`

Example: `https://myaccount.blob.core.windows.net:443`

If you don't specify a port, by default port 443 is used for HTTPS and port 80 is used for HTTP.

4. Select **Continue**. Then select the authentication type and enter the required information for the Cloud Storage Pool endpoint:

Access key

For Amazon S3/GCP or other S3-compatible provider

- a. **Access key ID:** Enter the access key ID for the account that owns the external bucket.
- b. **Secret access key:** Enter the secret access key.

IAM Roles Anywhere

For AWS IAM Roles Anywhere service

StorageGRID uses the AWS Security Token Service (STS) to dynamically generate a short-lived token to access AWS resources.

- a. **AWS IAM Roles Anywhere region:** Select the region for the Cloud Storage Pool. For example, `us-east-1`.
- b. **Trust anchor URN:** Enter the URN of the trust anchor that validates requests for short-lived STS credentials. Can be a root or intermediate CA.
- c. **Profile URN:** Enter the URN of the IAM Roles Anywhere profile that lists the roles that are assumable for anyone trusted.
- d. **Role URN:** Enter the URN of the IAM role that is assumable for anyone trusted.
- e. **Session duration:** Enter the duration of the temporary security credentials and role session. Enter at least 15 minutes and no more than 12 hours.
- f. **Server CA certificate** (optional): One or more trusted CA certificates, in PEM format, for verifying the IAM Roles Anywhere server. If omitted, the server won't be verified.
- g. **End-entity certificate:** The public key, in PEM format, of the X509 certificate signed by the trust anchor. AWS IAM Roles Anywhere uses this key to issue an STS token.
- h. **End-entity private key:** The private key for the end-entity certificate.

CAP (C2S access portal)

For Commercial Cloud Services (C2S) S3 service

- a. **Temporary credentials URL:** Enter the complete URL that StorageGRID will use to obtain temporary credentials from the CAP server, including all the required and optional API parameters assigned to your C2S account.
- b. **Server CA certificate:** Select **Browse** and upload the CA certificate that StorageGRID will use to verify the CAP server. Certificate must be PEM-encoded and issued by an appropriate Government Certificate Authority (CA).
- c. **Client certificate:** Select **Browse** and upload the certificate that StorageGRID will use to identify itself to the CAP server. The client certificate must be PEM-encoded, issued by an appropriate Government Certificate Authority (CA), and granted access to your C2S account.
- d. **Client private key:** Select **Browse** and upload the PEM-encoded private key for the client certificate.
- e. If the client private key is encrypted, enter the passphrase for decrypting the client private key. Otherwise, leave the **Client private key passphrase** field blank.



If the client certificate will be encrypted, use the traditional format for the encryption. PKCS #8 encrypted format is not supported.

Azure Blob Storage

For Azure Blob Storage, Shared Key only

- a. **Account name:** Enter the name of the storage account that owns the external container
- b. **Account key:** Enter the secret key for the storage account

You can use the Azure portal to find these values.

Anonymous

No additional information is required.

5. Select **Continue**. Then choose the type of server verification you want to use:

Option	Description
Use root CA certificates in Storage Node OS	Use the Grid CA certificates installed on the operating system to secure connections.
Use custom CA certificate	Use a custom CA certificate. Select Browse and upload the PEM-encoded certificate.
Do not verify certificate	Selecting this option means that TLS connections to the Cloud Storage Pool aren't secure.

6. Select **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the bucket or container and the service endpoint exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the bucket or container to identify it as a Cloud Storage Pool. Never remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

If Cloud Storage Pool validation fails, you receive an error message that explains why validation failed. For example, an error might be reported if there is a certificate error or if the bucket or container you specified does not already exist.

7. If an error occurs, see the [instructions for troubleshooting Cloud Storage Pools](#), resolve any issues, and then try saving the Cloud Storage Pool again.

View Cloud Storage Pool details

You can view the details of a Cloud Storage Pool to determine where it's used and to see which nodes and storage grades are included.

Before you begin

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

Steps

1. Select **ILM > Storage pools > Cloud Storage Pools**.

The Cloud Storage Pools table includes the following information for each Cloud Storage Pool that includes Storage Nodes:

- **Name:** The unique display name of the pool.
- **URI:** The Uniform Resource Identifier of the Cloud Storage Pool.
- **Provider type:** Which cloud provider is used for this Cloud Storage Pool.
- **Container:** The name of the bucket used for the Cloud Storage Pool.
- **ILM usage:** How the pool is currently being used. A Cloud Storage Pool might be unused or it might be used in one or more ILM rules, erasure-coding profiles, or both.
- **Last error:** The last error detected during a health check of this Cloud Storage Pool.

2. To view details for a specific Cloud Storage Pool, select its name.

The details page for the pool appears.

3. View the **Authentication** tab to learn about the authentication type for this Cloud Storage Pool and to edit the authentication details.
4. View the **Server verification** tab to learn about verification details, edit verification, download a new certificate, or copy the certificate PEM.
5. View the **ILM usage** tab to determine if the Cloud Storage Pool is currently being used in any ILM rules or erasure-coding profiles.
6. Optionally, go to the **ILM rules page** to [learn about and manage any rules](#) that use the Cloud Storage Pool.

Edit a Cloud Storage Pool

You can edit a Cloud Storage Pool to change its name, service endpoint, or other details; however, you can't change the S3 bucket or Azure container for a Cloud Storage Pool.

Before you begin

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

Steps

1. Select **ILM > Storage pools > Cloud Storage Pools**.

The Cloud Storage Pools table lists the existing Cloud Storage Pools.

2. Select the checkbox for the Cloud Storage Pool you want to edit, then select **Actions > Edit**.

Alternatively, select the name of the Cloud Storage Pool, then select **Edit**.

3. As required, change the Cloud Storage Pool name, service endpoint, authentication credentials, or certificate verification method.



You can't change the provider type or the S3 bucket or Azure container for a Cloud Storage Pool.

If you previously uploaded a server or client certificate, you can expand the **Certificate details** accordion to review the certificate that is currently in use.

4. Select **Save**.

When you save a Cloud Storage Pool, StorageGRID validates that the bucket or container and the service endpoint exist, and that they can be reached using the credentials that you specified.

If Cloud Storage Pool validation fails, an error message is displayed. For example, an error might be reported if there is a certificate error.

See the instructions for [troubleshooting Cloud Storage Pools](#), resolve the issue, and then try saving the Cloud Storage Pool again.

Remove a Cloud Storage Pool

You can remove a Cloud Storage Pool if it not used in an ILM rule and it does not contain object data.

Before you begin

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

If needed, use ILM to move object data

If the Cloud Storage Pool you want to remove contains object data, you must use ILM to move the data to a different location. For example, you can move the data to Storage Nodes on your grid or to a different Cloud Storage Pool.

Steps

1. Select **ILM > Storage pools > Cloud Storage Pools**.
2. Look at the ILM usage column in the table to determine whether you can remove the Cloud Storage Pool.

You can't remove a Cloud Storage Pool if it is being used in an ILM rule or in an erasure-coding profile.

3. If the Cloud Storage Pool is being used, select **cloud storage pool name > ILM usage**.
4. [Clone each ILM rule](#) that currently places objects in the Cloud Storage Pool you want to remove.
5. Determine where you want to move the existing objects managed by each rule you cloned.

You can use one or more storage pools or a different Cloud Storage Pool.

6. Edit each of the rules you cloned.

For Step 2 of the Create ILM rule wizard, select the new location from the **copies at** field.

7. [Create a new ILM policy](#) and replace each of the old rules with a cloned rule.
8. Activate the new policy.
9. Wait for ILM to remove objects from the Cloud Storage Pool and place them in the new location.

Delete Cloud Storage Pool

When the Cloud Storage Pool is empty and not used in any ILM rules, you can delete it.

Before you begin

- You have removed any ILM rules that might have used the pool.
- You have confirmed that the S3 bucket or Azure container does not contain any objects.

An error occurs if you attempt to remove a Cloud Storage Pool if it contains objects. See [Troubleshoot Cloud Storage Pools](#).



When you create a Cloud Storage Pool, StorageGRID writes a marker file to the bucket or container to identify it as a Cloud Storage Pool. Don't remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

Steps

1. Select **ILM > Storage pools > Cloud Storage Pools**.
2. If the ILM usage column indicates that Cloud Storage Pool is not being used, select the checkbox.
3. Select **Actions > Remove**.
4. Select **OK**.

Troubleshoot Cloud Storage Pools

Use these troubleshooting steps to help resolve errors you might encounter when creating, editing, or deleting a Cloud Storage Pool.

Determine if an error has occurred

StorageGRID performs a simple health check on every Cloud Storage Pool by reading the known object `x-ntap-sgws-cloud-pool-uuid` to ensure that the Cloud Storage Pool can be accessed and is functioning correctly. When StorageGRID encounters an error on the endpoint, it performs a health check every minute from each Storage Node. When the error is resolved, the health checks stop. If a health check detects an issue, a message is shown in the Last error column of the Cloud Storage Pools table on the Storage pools page.

The table shows the most recent error detected for each Cloud Storage Pool and indicates how long ago the error occurred.

In addition, a **Cloud Storage Pool connectivity error** alert is triggered if the health check detects that one or more new Cloud Storage Pool errors have occurred within the past 5 minutes. If you receive an email notification for this alert, go to the Storage pools page (select **ILM > Storage pools**), review the error messages in the Last error column, and refer to the troubleshooting guidelines below.

Check if an error has been resolved

After resolving any underlying issues, you can determine if the error has been resolved. From the Cloud Storage Pool page, select the endpoint, and select **Clear error**. A confirmation message indicates that StorageGRID has cleared the error for the Cloud Storage Pool.

If the underlying problem has been resolved, the error message is no longer displayed. However, if the underlying problem has not been fixed (or if a different error is encountered), the error message will be shown

in the Last error column within a few minutes.

Error: Health check failed. Error from endpoint

You might encounter this error when you enable S3 Object Lock with default retention for your Amazon S3 bucket after you start using this bucket for a Cloud Storage Pool. This error occurs when the PUT operation doesn't have an HTTP header with a payload checksum value such as `Content-MD5`. This header value is required by AWS for PUT operations into buckets with S3 Object Lock enabled.

To correct this issue, follow the steps in [Edit a Cloud Storage Pool](#) without making any changes. This action triggers the validation of the Cloud Storage Pool configuration that automatically detects and updates the S3 Object Lock flag on a Cloud Storage Pool endpoint configuration.

Error: This Cloud Storage Pool contains unexpected content

You might encounter this error when you try to create, edit, or delete a Cloud Storage Pool. This error occurs if the bucket or container includes the `x-ntap-sgws-cloud-pool-uuid` marker file, but that file doesn't have the metadata field with the expected UUID.

Typically, you will only see this error if you are creating a new Cloud Storage Pool and another instance of StorageGRID is already using the same Cloud Storage Pool.

Try one of these steps to correct the issue:

- If you are configuring a new Cloud Storage Pool and the bucket contains the `x-ntap-sgws-cloud-pool-uuid` file and additional object keys similar to the following example, create a new bucket and use this new bucket instead.

Example of an additional object key: `my-bucket.3E64CF2C-B74D-4B7D-AFE7-AD28BC18B2F6.1727326606730410`

- If the `x-ntap-sgws-cloud-pool-uuid` file is the only object in the bucket, delete this file.

If these steps don't apply to your scenario, contact support.

Error: Could not create or update Cloud Storage Pool. Error from endpoint

You might encounter this error under the following circumstances:

- When you try to create or edit a Cloud Storage Pool.
- When you select an unsupported platform, authentication, or protocol combination with S3 Object Lock during the configuration of a new Cloud Storage Pool. See [Considerations for Cloud Storage Pools](#).

This error indicates that a connectivity or configuration issue is preventing StorageGRID from writing to the Cloud Storage Pool.

To correct the issue, review the error message from the endpoint.

- If the error message contains `Get url: EOF`, check that the service endpoint used for the Cloud Storage Pool does not use HTTP for a container or bucket that requires HTTPS.
- If the error message contains `Get url: net/http: request canceled while waiting for connection`, verify that the network configuration allows Storage Nodes to access the service endpoint used for the Cloud Storage Pool.

- If the error is due to an unsupported platform, authentication, or protocol, change to a supported configuration with S3 Object Lock and try to save the new Cloud Storage Pool again.
- For all other endpoint error messages, try one or more of the following:
 - Create an external container or bucket with the same name you entered for the Cloud Storage Pool, and try to save the new Cloud Storage Pool again.
 - Correct the container or bucket name you specified for the Cloud Storage Pool, and try to save the new Cloud Storage Pool again.

Error: Failed to parse CA certificate

You might encounter this error when you try to create or edit a Cloud Storage Pool. The error occurs if StorageGRID could not parse the certificate you entered when configuring the Cloud Storage Pool.

To correct the issue, check the CA certificate you provided for issues.

Error: A Cloud Storage Pool with this ID was not found

You might encounter this error when you try to edit or delete a Cloud Storage Pool. This error occurs if the endpoint returns a 404 response, which can mean either of the following:

- The credentials used for the Cloud Storage Pool don't have read permission for the bucket.
- The bucket used for the Cloud Storage Pool does not include the `x-ntap-sgws-cloud-pool-uuid` marker file.

Try one or more of these steps to correct the issue:

- Check that the user associated with the configured Access Key has the requisite permissions.
- Edit the Cloud Storage Pool with credentials that have the requisite permissions.
- If the permissions are correct, contact support.

Error: Could not check the content of the Cloud Storage Pool. Error from endpoint

You might encounter this error when you try to delete a Cloud Storage Pool. This error indicates that some kind of connectivity or configuration issue is preventing StorageGRID from reading the contents of the Cloud Storage Pool bucket.

To correct the issue, review the error message from the endpoint.

Error: Objects have already been placed in this bucket

You might encounter this error when you try to delete a Cloud Storage Pool. You can't delete a Cloud Storage Pool if it contains data that was moved there by ILM, data that was in the bucket before you configured the Cloud Storage Pool, or data that was put in the bucket by some other source after the Cloud Storage Pool was created.

Try one or more of these steps to correct the issue:

- Follow the instructions for moving objects back to StorageGRID in "Lifecycle of a Cloud Storage Pool object."
- If you are certain the remaining objects were not placed in the Cloud Storage Pool by ILM, manually delete the objects from the bucket.



Never manually delete objects from a Cloud Storage Pool that might have been placed there by ILM. If you later attempt to access a manually deleted object from StorageGRID, the deleted object will not be found.

Error: Proxy encountered an external error while trying to reach the Cloud Storage Pool

You might encounter this error if you have configured a non-transparent storage proxy between Storage Nodes and the external S3 endpoint used for the Cloud Storage Pool. This error occurs if the external proxy server can't reach the Cloud Storage Pool endpoint. For example, the DNS server might not be able to resolve the hostname or there might be an external networking issue.

Try one or more of these steps to correct the issue:

- Check the settings for the Cloud Storage Pool (**ILM > Storage pools**).
- Check the networking configuration of the storage proxy server.

Error: X.509 certificate is out of validity period

You might encounter this error when you try to delete a Cloud Storage Pool. This error occurs when the authentication requires an X.509 certificate to ensure the correct external Cloud Storage Pool is validated and the external pool is empty before the Cloud Storage Pool configuration is deleted.

Try these steps to correct the issue:

- Update the certificate configured for authentication to the Cloud Storage Pool.
- Make sure any certificate expiration alert on this Cloud Storage Pool is resolved.

Related information

[Lifecycle of a Cloud Storage Pool object](#)

Manage erasure-coding profiles

You can view the details for an erasure-coding profile and rename a profile if needed. You can deactivate an erasure-coding profile if it is not currently used in any ILM rules.

Before you begin

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

View erasure-coding profile details

You can view the details of an erasure-coding profile to determine its status, the erasure-coding scheme used, and other information.

Steps

1. Select **CONFIGURATION > System > Erasure coding**.
2. Select the profile. The detail page for the profile appears.
3. Optionally, view the ILM rules tab for a list of ILM rules that use the profile, and the ILM policies that use those rules.

4. Optionally, view the Storage Nodes tab for details about each Storage Node in the profile's storage pool, such as the site where it's located and the storage usage.

Rename an erasure-coding profile

You might want to rename an erasure-coding profile to make it more obvious what the profile does.

Steps

1. Select **CONFIGURATION > System > Erasure coding**.
2. Select the profile you want to rename.
3. Select **Rename**.
4. Enter a unique name for the erasure-coding profile.

The erasure-coding profile name is appended to the storage pool name in the placement instruction for an ILM rule.



Erasure-coding profile names must be unique. A validation error occurs if you use the name of an existing profile, even if that profile has been deactivated.

5. Select **Save**.

Deactivate an erasure-coding profile

You can deactivate an erasure-coding profile if you no longer plan to use it and if the profile is not currently used in any ILM rules.



Confirm that no erasure-coded data repair operations or decommission procedures are in process. An error message is returned if you attempt to deactivate an erasure-coding profile while either of these operations are in progress.

About this task

StorageGRID prevents you from deactivating an erasure-coding profile if either of the following is true:

- The erasure-coding profile is currently used in an ILM rule.
- The erasure-coding profile is no longer used in any ILM rules, but object data and parity fragments for the profile still exist.

Steps

1. Select **CONFIGURATION > System > Erasure coding**.
2. On the Active tab, review the **Status** column to confirm that the erasure-coding profile you want to deactivate is not used in any ILM rules.

You can't deactivate an erasure-coding profile if it is used in any ILM rule. In the example, the 2+1 Data Center 1 profile is used in at least one ILM rule.

☐	Profile name ? ▾	Status ? ▾	Storage pool ? ▾	Erasure-coding scheme ? ▾
☐	2+1 Data Center 1	Used in <u>5 rules</u>	Data Center 1	2+1
☐	New profile	Deactivated	Data Center 1	2+1

3. If the profile is used in an ILM rule, follow these steps:

- a. Select **ILM > Rules**.
- b. Select each rule and review the retention diagram to determine if the rule uses the erasure-coding profile you want to deactivate.
- c. If the ILM rule uses the erasure-coding profile you want to deactivate, determine if the rule is used in any ILM policy.
- d. Complete the additional steps in the table, based on where the erasure-coding profile is used.

Where has the profile been used?	Additional steps to perform before deactivating the profile	Refer to these additional instructions
Never used in any ILM rule	No additional steps required. Continue with this procedure.	<i>None</i>
In an ILM rule that has never been used in any ILM policy	- Edit or delete all affected ILM rules. If you edit the rule, remove all placements that use the erasure-coding profile. - Continue with this procedure.	Work with ILM rules and ILM policies

Where has the profile been used?	Additional steps to perform before deactivating the profile	Refer to these additional instructions
In an ILM rule that is currently in an active ILM policy	1. Clone the policy. 2. Remove the ILM rule that uses the erasure-coding profile. 3. Add one or more new ILM rules to ensure objects are protected. 4. Save, simulate, and activate the new policy. 5. Wait for the new policy to be applied and for existing objects to be moved to new locations based on the new rules you added. Note: Depending on the number of objects and the size of your StorageGRID system, it might take weeks or even months for ILM operations to move the objects to new locations, based on the new ILM rules. While you can safely attempt to deactivate an erasure-coding profile while it is still associated with data, the deactivation operation will fail. An error message will inform you if the profile is not yet ready to be deactivated. 6. Edit or delete the rule you removed from the policy. If you edit the rule, remove all placements that use the erasure-coding profile. 7. Continue with this procedure.	Create an ILM policy Work with ILM rules and ILM policies
In an ILM rule that is currently in an ILM policy	1. Edit the policy. 2. Remove the ILM rule that uses the erasure-coding profile. 3. Add one or more new ILM rules to ensure all objects are protected. 4. Save the policy. 5. Edit or delete the rule you removed from the policy. If you edit the rule, remove all placements that use the erasure-coding profile. 6. Continue with this procedure.	Create an ILM policy Work with ILM rules and ILM policies

e. Refresh the Erasure-Coding Profiles page to ensure that the profile is not used in an ILM rule.

4. If the profile is not used in an ILM rule, select the radio button and select **Deactivate**. The Deactivate erasure-coding profile dialog box appears.



You can select multiple profiles to deactivate at the same time, as long as each profile is not used in any rule.

5. If you are sure you want to deactivate the profile, select **Deactivate**.

Results

- If StorageGRID is able to deactivate the erasure-coding profile, its status is Deactivated. You can no longer select this profile for any ILM rule. You can't reactivate a deactivated profile.
- If StorageGRID is not able to deactivate the profile, an error message appears. For example, an error message appears if object data is still associated with this profile. You might need to wait several weeks before trying the deactivation process again.

Configure regions (optional and S3 only)

ILM rules can filter objects based on the regions where S3 buckets are created, allowing you to store objects from different regions in different storage locations.

If you want to use an S3 bucket region as a filter in a rule, you must first create the regions that can be used by the buckets in your system.



You can't change the region for a bucket after the bucket has been created.

Before you begin

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

About this task

When creating an S3 bucket, you can specify that the bucket be created in a specific region. Specifying a region allows the bucket to be geographically close to its users, which can help optimize latency, minimize costs, and address regulatory requirements.

When you create an ILM rule, you might want to use the region associated with an S3 bucket as an advanced filter. For example, you can design a rule that applies only to objects in S3 buckets created in the `us-west-2` region. You can then specify that copies of those objects be placed on Storage Nodes at a data center site within that region to optimize latency.

When configuring regions, follow these guidelines:

- By default, all buckets are considered to belong to the `us-east-1` region.
- You must create the regions using the Grid Manager before you can specify a non-default region when creating buckets using the Tenant Manager or Tenant Management API or with the LocationConstraint request element for S3 PUT Bucket API requests. An error occurs if a PUT Bucket request uses a region that has not been defined in StorageGRID.
- You must use the exact region name when you create the S3 bucket. Region names are case sensitive. Valid characters are numbers, letters, and hyphens.



EU is not considered to be an alias for eu-west-1. If you want to use the EU or eu-west-1 region, you must use the exact name.

- You can't delete or modify a region if it's used in a rule that is assigned to any policy (active or inactive).
- If you use an invalid region as the advanced filter in an ILM rule, you can't add that rule to a policy.

An invalid region can result if you use a region as an advanced filter in an ILM rule but you later delete that region, or if you use the Grid Management API to create a rule and specify a region that you have not defined.

- If you delete a region after using it to create an S3 bucket, you will need to re-add the region if you ever want to use the Location Constraint advanced filter to find objects in that bucket.

Steps

1. Select **ILM > Regions**.

The Regions page appears, with the currently defined regions listed. **Region 1** shows the default region, `us-east-1`, which can't be modified or removed.

2. To add a region:

- a. Select **Add another region**.
- b. Enter the name of a region that you want to use when creating S3 buckets.

You must use this exact region name as the LocationConstraint request element when you create the corresponding S3 bucket.

3. To remove an unused region, select the delete icon .

An error message appears if you attempt to remove a region that is currently used in any policy (active or inactive).

4. When you are done making changes, select **Save**.

You can now select these regions from the Advanced filters section in step 1 of the Create ILM rule wizard. See [Use advanced filters in ILM rules](#).

Create ILM rule

Use ILM rules to manage objects

To manage objects, you create a set of information lifecycle management (ILM) rules and organize them into an ILM policy.

Every object ingested into the system is evaluated against the active policy. When a rule in the policy matches an object's metadata, the instructions in the rule determine what actions StorageGRID takes to copy and store that object.



Object metadata is not managed by ILM rules. Instead, object metadata is stored in a Cassandra database in what is known as a metadata store. Three copies of object metadata are automatically maintained at each site to protect the data from loss.

Elements of an ILM rule

An ILM rule has three elements:

- **Filtering criteria:** A rule's basic and advanced filters define which objects the rule applies to. If an object matches all filters, StorageGRID applies the rule and creates the object copies specified in the rule's placement instructions.
- **Placement instructions:** A rule's placement instructions define the number, type, and location of object copies. Each rule can include a sequence of placement instructions to change the number, type, and location of object copies over time. When the time period for one placement expires, the instructions in the next placement are automatically applied by the next ILM evaluation.
- **Ingest behavior:** A rule's ingest behavior allows you to choose how the objects filtered by the rule are protected as they are ingested (when an S3 client saves an object to the grid).

ILM rule filtering

When you create an ILM rule, you specify filters to identify which objects the rule applies to.

In the simplest case, a rule might not use any filters. Any rule that does not use filters applies to all objects, so it must be the last (default) rule in an ILM policy. The default rule provides storage instructions for objects that don't match the filters in another rule.

- Basic filters allow you to apply different rules to large, distinct groups of objects. These filters allow you to apply a rule to specific tenant accounts, specific S3 buckets, or both.

Basic filters give you a simple way to apply different rules to large numbers of objects. For example, your company's financial records might need to be stored to meet regulatory requirements, while data from the marketing department might need to be stored to facilitate daily operations. After creating separate tenant accounts for each department or after segregating data from the different departments into separate S3 buckets, you can easily create one rule that applies to all financial records and a second rule that applies to all marketing data.

- Advanced filters give you granular control. You can create filters to select objects based on the following object properties:
 - Ingest time
 - Last access time
 - All or part of the object name (Key)
 - Location constraint (S3 only)
 - Object size
 - User metadata
 - Object tag (S3 only)

You can filter objects on very specific criteria. For example, objects stored by a hospital's imaging department might be used frequently when they are less than 30 days old and infrequently afterwards, while objects that contain patient visit information might need to be copied to the billing department at the health network's headquarters. You can create filters that identify each type of object based on object name, size, S3 object tags, or any other relevant criteria, and then create separate rules to store each set of objects appropriately.

You can combine filters as needed in a single rule. For example, the marketing department might want to store large image files differently than their vendor records, while the Human Resources department might need to store personnel records in a specific geography and policy information centrally. In this case you can create rules that filter by tenant account to segregate the records from each department, while using filters in each rule to identify the specific type of objects that the rule applies to.

ILM rule placement instructions

Placement instructions determine where, when, and how object data is stored. An ILM rule can include one or more placement instructions. Each placement instruction applies to a single period of time.

When you create placement instructions:

- You start by specifying the reference time, which determines when the placement instructions start. The reference time might be when an object is ingested, when an object is accessed, when a versioned object becomes noncurrent, or a user-defined time.
- Next, you specify when the placement will apply, relative to the reference time. For example, a placement might start on day 0 and continue for 365 days, relative to when the object was ingested.
- Finally, you specify the type of copies (replication or erasure coding) and the location where the copies are stored. For example, you might want to store two replicated copies at two different sites.

Each rule can define multiple placements for a single time period and different placements for different time periods.

- To place objects in multiple locations during a single time period, select **Add other type or location** to add more than one line for that time period.
- To place objects in different locations in different time periods, select **Add another time period** to add the next time period. Then, specify one or more lines within the time period.

The example shows two placement instructions on the Define placements page of the Create ILM rule wizard.

Time period and placements Sort by start date

If you want a rule to apply only to specific objects, select **Previous** and add advanced filters. When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the filter.

Time period 1

From Day store for days

Store objects by copies at ,

and store objects by using

[Add other type or location](#)

Time period 2

From Day store forever

Store objects by copies at

[Add other type or location](#)

The first placement instruction 1 has two lines for the first year:

- The first line creates two replicated object copies at two data center sites.
- The second line creates a 6+3 erasure-coded copy using all data center sites.

633

The second placement instruction  creates two copies after one year and keeps those copies forever.

When you define the set of placement instructions for a rule, you must ensure that at least one placement instruction begins at day 0, that there are no gaps between the time periods you have defined, and that the final placement instruction continues either forever or until you no longer require any object copies.

As each time period in the rule expires, the content placement instructions for the next time period are applied. New object copies are created and any unneeded copies are deleted.

ILM rule ingest behavior

Ingest behavior controls whether object copies are immediately placed according to the instructions in the rule, or if interim copies are made and the placement instructions are applied later. The following ingest behaviors are available for ILM rules:

- **Balanced:** StorageGRID attempts to make all copies specified in the ILM rule at ingest; if this is not possible, interim copies are made and success is returned to the client. The copies specified in the ILM rule are made when possible.
- **Strict:** All copies specified in the ILM rule must be made before success is returned to the client.
- **Dual commit:** StorageGRID immediately makes interim copies of the object and returns success to the client. Copies specified in the ILM rule are made when possible.

Related information

- [Ingest options](#)
- [Advantages, disadvantages, and limitations of the ingest options](#)
- [How consistency and ILM rules interact to affect data protection](#)

Example ILM rule

As an example, an ILM rule could specify the following:

- Apply only to the objects belonging to Tenant A.
- Make two replicated copies of those objects and store each copy at a different site.
- Retain the two copies "forever," which means that StorageGRID will not automatically delete them. Instead, StorageGRID will retain these objects until they are deleted by a client delete request or by the expiration of a bucket lifecycle.
- Use the Balanced option for ingest behavior: the two-site placement instruction is applied as soon as Tenant A saves an object to StorageGRID, unless it is not possible to immediately make both required copies.

For example, if Site 2 is unreachable when Tenant A saves an object, StorageGRID will make two interim copies on Storage Nodes at Site 1. As soon as Site 2 becomes available, StorageGRID will make the required copy at that site.

Related information

- [What is a storage pool](#)
- [What is a Cloud Storage Pool](#)

Access the Create an ILM rule wizard


ILM rules allow you to manage the placement of object data over time. To create an ILM rule, you use the Create an ILM rule wizard.



If you want to create the default ILM rule for a policy, follow the [instructions for creating a default ILM rule](#) instead.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have [specific access permissions](#).
- If you want to specify which tenant accounts this rule applies to, you have the [Tenant accounts permission](#) or you know the account ID for each account.
- If you want the rule to filter objects on last access time metadata, Last access time updates must be enabled by S3 bucket.
- You have configured any Cloud Storage Pools you plan to use. See [Create Cloud Storage Pool](#).
- You are familiar with the [ingest options](#).
- If you need to create a compliant rule for use with S3 Object Lock, you are familiar with the [requirements for S3 Object Lock](#).
- Optionally, you have watched the video: [Video: ILM rules overview](#).



About this task

When creating ILM rules:

- Consider the StorageGRID system's topology and storage configurations.
- Consider what types of object copies you want to make (replicated or erasure-coded) and the number of copies of each object that are required.
- Determine what types of object metadata are used in the applications that connect to the StorageGRID system. ILM rules filter objects based on their metadata.
- Consider where you want object copies to be placed over time.
- Decide which ingest option to use (Balanced, Strict, or Dual commit).

Steps

1. Select **ILM > Rules**.
2. Select **Create**. [Step 1 \(Enter details\)](#) of the Create an ILM rule wizard appears.

Step 1 of 3: Enter details

The **Enter details** step of the Create an ILM rule wizard allows you to enter a name and description for the rule and to define filters for the rule.

Entering a description and defining filters for the rule are optional.

About this task

When evaluating an object against an [ILM rule](#), StorageGRID compares the object metadata to the rule's filters. If the object metadata matches all filters, StorageGRID uses the rule to place the object. You can design a rule to apply to all objects, or you can specify basic filters, such as one or more tenant accounts or bucket names, or advanced filters, such as the object's size or user metadata.

Steps

1. Enter a unique name for the rule in the **Name** field.
2. Optionally, enter a short description for the rule in the **Description** field.

You should describe the rule's purpose or function so you can recognize the rule later.

3. Optionally, select one or more S3 tenant accounts to which this rule applies. If this rule applies to all tenants, leave this field blank.

If you don't have either the Root access permission or the Tenant accounts permission, you can't select tenants from the list. Instead, enter the tenant ID or enter multiple IDs as a comma-delimited string.

4. Optionally, specify the S3 buckets to which this rule applies.

If **applies to all buckets** is selected (default), the rule applies to all S3 buckets.

5. For S3 tenants, optionally select **Yes** to apply the rule only to older object versions in S3 buckets that have versioning enabled.

If you select **Yes**, "Noncurrent time" will be automatically selected for Reference time in [Step 2 of the Create an ILM rule wizard](#).



Noncurrent time applies only to S3 objects in versioning-enabled buckets. See [Operations on buckets](#), [PutBucketVersioning](#) and [Manage objects with S3 Object Lock](#).

You can use this option to reduce the storage impact of versioned objects by filtering for noncurrent object versions. See [Example 4: ILM rules and policy for S3 versioned objects](#).

6. Optionally, select **Add an advanced filter** to specify additional filters.

If you don't configure advanced filtering, the rule applies to all objects that match the basic filters. For more information about advanced filtering, see [Use advanced filters in ILM rules](#) and [Specify multiple metadata types and values](#).

7. Select **Continue**. [Step 2 \(Define placements\)](#) of the Create an ILM rule wizard appears.

Use advanced filters in ILM rules

Advanced filtering allows you to create ILM rules that apply only to specific objects based on their metadata. When you set up advanced filtering for a rule, you select the type of metadata you want to match, select an operator, and specify a metadata value. When objects are evaluated, the ILM rule is applied only to those

objects that have metadata matching the advanced filter.

The table shows the types of metadata you can specify in advanced filters, the operators you can use for each type of metadata, and the metadata values expected.

Metadata type	Supported operators	Metadata value
Ingest time	• is • is not • is before • is on or before • is after • is on or after	Time and date the object was ingested. Note: To avoid resource issues when activating a new ILM policy, you can use the Ingest time advanced filter in any rule that might change the location of large numbers of existing objects. Set Ingest time to be greater than or equal to the approximate time when the new policy will go into effect to ensure that existing objects aren't moved unnecessarily.
Key	• equals • does not equal • contains • does not contain • starts with • does not start with • ends with • does not end with	All or part of a unique S3 object key. For example, you might want to match objects that end with <code>.txt</code> or start with <code>test-object/</code>.
Last access time	• is • is not • is before • is on or before • is after • is on or after	Time and date the object was last retrieved (read or viewed). Note: If you plan to use last access time as an advanced filter, Last access time updates must be enabled for the S3 bucket.
Location constraint (S3 only)	• equals • does not equal	The region where an S3 bucket was created. Use ILM > Regions to define the regions that are shown. Note: A value of <code>us-east-1</code> will match objects in buckets created in the <code>us-east-1</code> region as well as objects in buckets that have no region specified. See Configure regions (optional and S3 only).

Metadata type	Supported operators	Metadata value
Object size	• equals • does not equal • less than • less than or equal to • greater than • greater than or equal to	The object's size. Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.
User metadata	• contains • ends with • equals • exists • starts with • does not contain • does not end with • does not equal • does not exist • does not start with	Key-value pair, where User metadata name is the key and Metadata value is the value. For example, to filter on objects that have user metadata of <code>color=blue</code>, specify <code>color</code> for User metadata name, <code>equals</code> for the operator, and <code>blue</code> for Metadata value. Note: User-metadata names aren't case sensitive; user-metadata values are case sensitive.
Object tag (S3 only)	• contains • ends with • equals • exists • starts with • does not contain • does not end with • does not equal • does not exist • does not start with	Key-value pair, where Object tag name is the key and Object tag value is the value. For example, to filter on objects that have an object tag of <code>Image=True</code>, specify <code>Image</code> for Object tag name, <code>equals</code> for the operator, and <code>True</code> for Object tag value. Note: Object tag names and object tag values are case sensitive. You must enter these items exactly as they were defined for the object.

Specify multiple metadata types and values

When you define advanced filtering, you can specify multiple types of metadata and multiple metadata values. For example, if you want a rule to match objects between 10 MB and 100 MB in size, you would select the **Object size** metadata type and specify two metadata values.

- The first metadata value specifies objects greater than or equal to 10 MB.
- The second metadata value specifies objects less than or equal to 100 MB.

Filter group 1
Objects with all of following metadata will be evaluated by this rule:

Object size
greater than or equal to
10
MB

and
Object size
less than or equal to
100
MB

Using multiple entries allows you to have precise control over which objects are matched. In the following example, the rule applies to objects that have Brand A or Brand B as the value of the camera_type user metadata. However, the rule only applies to those Brand B objects that are smaller than 10 MB.

Filter group 1
Objects with all of following metadata will be evaluated by this rule:

User metadata
camera_type
equals
Brand A

Add another advanced filter

or
Filter group 2
Objects with all of following metadata will be evaluated by this rule:

User metadata
camera_type
equals
Brand B

and
Object size
less than or equal to
10
MB

Add another advanced filter

Step 2 of 3: Define placements

The **Define placements** step of the Create ILM Rule wizard allows you to define the placement instructions that determine how long objects are stored, the type of copies (replicated or erasure-coded), the storage location, and the number of copies.



The screenshots shown are examples. Your results might vary depending on your StorageGRID version.

About this task

An ILM rule can include one or more placement instructions. Each placement instruction applies to a single period of time. When you use more than one instruction, the time periods must be contiguous, and at least one instruction must start on day 0. The instructions can continue either forever, or until you no longer require any object copies.

Each placement instruction can have multiple lines if you want to create different types of copies or use different locations during that time period.

In this example, the ILM rule stores one replicated copy in Site 1 and one replicated copy in Site 2 for the first year. After one year, a 2+1 erasure-coded copy is made and saved at only one site.

Time period 1
From Day
0
store
for
365
days

Store objects by
replicating
1
copies at
Site 1

and store objects by
replicating
1
copies at
Site 2

Add other type or location

Time period 2
From Day
365
store
forever

Store objects by
erasure coding
using
2+1 EC scheme at Site 3

Add other type or location

Steps

- For **Reference time**, select the type of time to use when calculating the start time for a placement instruction.

Option	Description
Ingest time	The time when the object was ingested.
Last access time	The time when the object was last retrieved (read or viewed). To use this option, updates to Last access time must be enabled for the S3 bucket. Refer to Use Last access time in ILM rules.
User defined creation time	A time specified in user-defined metadata.
Noncurrent time	"Noncurrent time" is automatically selected if you selected Yes for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?" in Step 1 of the Create an ILM rule wizard .

If you want to create a *compliant* rule, you must select **Ingest time**. Refer to [Manage objects with S3 Object Lock](#).

- In the **Time period and placements** section, enter a starting time and a duration for the first time period.

For example, you might want to specify where to store objects for the first year (*From day 0 store for 365 days*). At least one instruction must start at day 0.

- If you want to create replicated copies:
 - From the **Store objects by** drop-down list, select **replicating**.
 - Select the number of copies you want to make.

A warning appears if you change the number of copies to 1. An ILM rule that creates only one

replicated copy for any time period puts data at risk of permanent loss. Refer to [Why you should not use single-copy replication](#).

To avoid the risk, do one or more of the following:

- Increase the number of copies for the time period.
- Add copies to other storage pools or to a Cloud Storage Pool.
- Select **erasure coding** instead of **replicating**.

You can safely ignore this warning if this rule already creates multiple copies for all time periods.

c. In the **copies at** field, select the storage pools you want to add.

If you specify only one storage pool, be aware that StorageGRID can store only one replicated copy of an object on any given Storage Node. If your grid includes three Storage Nodes and you select 4 as the number of copies, only three copies will be made—one copy for each Storage Node.

The **ILM placement unachievable** alert is triggered to indicate that the ILM rule could not be completely applied.

If you specify more than one storage pool, keep these rules in mind:

- The number of copies can't be greater than the number of storage pools.
- If the number of copies equals the number of storage pools, one copy of the object is stored in each storage pool.
- If the number of copies is less than the number of storage pools, one copy is stored at the ingest site, and then the system distributes the remaining copies to keep disk usage among the pools balanced, while ensuring that no site gets more than one copy of an object.
- If the storage pools overlap (contain the same Storage Nodes), all copies of the object might be saved at only one site. For this reason, don't specify the All Storage Nodes storage pool (StorageGRID 11.6 and earlier) and another storage pool.

4. If you want to create an erasure-coded copy:

a. From the **Store objects by** drop-down list, select **erasure coding**.



Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

b. If you didn't add an Object size filter for a value greater than 200 KB, select **Previous** to return to Step 1. Then, select **Add an advanced filter** and set an **Object size** filter to any value greater than 200 KB.

c. Select the storage pool you want to add and the erasure-coding scheme you want to use.

The storage location for an erasure-coded copy includes the name of the erasure-coding scheme, followed by the name of the storage pool.

Available erasure-coding schemes are limited by the number of Storage Nodes in the storage pool you select. A **Recommended** badge appears next to the schemes that provide either the [best protection](#) or [the lowest storage overhead](#).

5. Optionally:

- a. Select **Add other type or location** to create additional copies at different locations.
- b. Select **Add another time period** to add different time periods.



Object deletions occur based on the following settings:

- Objects are automatically deleted at the end of the final time period unless another time period ends with **forever**.
- Depending on [bucket and tenant retention period settings](#), objects might not be deleted even if the ILM retention period ends.

6. If you want to store objects in a Cloud Storage Pool:
 - a. In the **Store objects by** drop-down list, select **replicating**.
 - b. Select the **copies at** field, then select a Cloud Storage Pool.

When using Cloud Storage Pools, keep these rules in mind:

- You can't select more than one Cloud Storage Pool in a single placement instruction. Similarly, you can't select a Cloud Storage Pool and a storage pool in the same placement instruction.
- You can store only one copy of an object in any given Cloud Storage Pool. An error message appears if you set **Copies** to 2 or more.
- You can't store more than one object copy in any Cloud Storage Pool at the same time. An error message appears if multiple placements that use a Cloud Storage Pool have overlapping dates or if multiple lines in the same placement use a Cloud Storage Pool.
- You can store an object in a Cloud Storage Pool at the same time that object is being stored as replicated or erasure-coded copies in StorageGRID. However, you must include more than one line in the placement instruction for the time period, so you can specify the number and types of copies for each location.


7. In the Retention diagram, confirm your placement instructions.

In this example, the ILM rule stores one replicated copy in Site 1 and one replicated copy in Site 2 for the first year. After one year and for an additional 10 years, a 6+3 erasure-coded copy will be saved at three sites. After 11 years total, the objects will be deleted from StorageGRID.

The Rule analysis section of the Retention diagram states:

- StorageGRID site-loss protection will apply for the duration of this rule.
- Objects processed by this rule will be deleted after Day 4015.

Refer to [Enable site-loss protection](#).



8. Select **Continue**. [Step 3 \(Select ingest behavior\)](#) of the Create an ILM rule wizard appears.

Use Last access time in ILM rules

You can use Last access time as the reference time in an ILM rule. For example, you might want to leave objects that have been viewed in the last three months on local Storage Nodes, while moving objects that have not been viewed as recently to an off-site location. You can also use Last access time as an advanced filter if you want an ILM rule to apply only to objects that were last accessed on a specific date.

About this task

Before using Last access time in an ILM rule, review the following considerations:

- When using Last access time as a reference time, be aware that changing the Last access time for an object does not trigger an immediate ILM evaluation. Instead, the object's placements are assessed and the object is moved as required when background ILM evaluates the object. This could take two weeks or more after the object is accessed.

Take this latency into account when creating ILM rules based on Last access time and avoid placements that use short time periods (less than one month).

- When using Last access time as an advanced filter or as a reference time, you must enable last access time updates for S3 buckets. You can use the [Tenant Manager](#) or the [Tenant Management API](#).



Last access time updates are disabled by default for S3 buckets.



Be aware that enabling last access time updates can reduce performance, especially in systems with small objects. The performance impact occurs because StorageGRID must update the objects with new timestamps every time the objects are retrieved.

The following table summarizes whether the Last access time is updated for all objects in the bucket for different types of requests.

Type of request	Whether Last access time is updated when last access time updates are disabled	Whether Last access time is updated when last access time updates are enabled
Request to retrieve an object, its access control list, or its metadata	No	Yes
Request to update an object's metadata	Yes	Yes
Request to copy an object from one bucket to another	• No, for the source copy • Yes, for the destination copy	• Yes, for the source copy • Yes, for the destination copy
Request to complete a multipart upload	Yes, for the assembled object	Yes, for the assembled object

Step 3 of 3: Select ingest behavior

The **Select ingest behavior** step of the Create ILM Rule wizard allows you to choose how the objects filtered by this rule are protected as they are ingested.

About this task

StorageGRID can make interim copies and queue the objects for ILM evaluation later, or it can make copies to meet the rule's placement instructions immediately.

Steps

1. Select the [ingest behavior](#) to use.

For more information, see [Advantages, disadvantages, and limitations of the ingest options](#).



You can't use the Balanced or Strict option if the rule uses one of these placements:

- A Cloud Storage Pool at day 0
- A Cloud Storage Pool when the rule uses a User defined creation time as a Reference time

See [Example 5: ILM rules and policy for Strict ingest behavior](#).

2. Select **Create**.

The ILM rule is created. The rule does not become active until it is added to an [ILM policy](#) and that policy is activated.

To view the details of the rule, select the rule's name on the ILM rules page.

Create a default ILM rule

Before creating an ILM policy, you must create a default rule to place any objects not matched by another rule in the policy. The default rule can't use any filters. It must apply

to all tenants, all buckets, and all object versions.

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 default rule is the last rule to be evaluated in an ILM policy, so it can't use any filters. The placement instructions for the default rule are applied to any objects that aren't matched by another rule in the policy.

In this example policy, the first rule applies only to objects belonging to test-tenant-1. The default rule, which is last, applies to objects belonging to all other tenant accounts.

Proposed policy name

Example ILM policy

Reason for change

Example

Manage rules

1. Select the rules you want to add to the policy.

2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

Select rules

Rule order	Rule name	Filters
1	EC for test-tenant-1	Tenant is test-tenant-1
Default	Default rule	—

When you create the default rule, keep these requirements in mind:

- The default rule will automatically be placed as the last rule when you add it to a policy.
- The default rule can't use any basic or advanced filters.
- The default rule must apply to all object versions.
- The default rule should create replicated copies.



Don't use a rule that creates erasure-coded copies as the default rule for a policy. Erasure-coding rules should use an advanced filter to prevent smaller objects from being erasure-coded.

- In general, the default rule should retain objects forever.
- If you are using (or you plan to enable) the global S3 Object Lock setting, the default rule must be compliant.

Steps

1. Select **ILM > Rules**.
2. Select **Create**.

Step 1 (Enter details) of the Create ILM rule wizard appears.

3. Enter a unique name for the rule in the **Rule name** field.
4. Optionally, enter a short description for the rule in the **Description** field.
5. Leave the **Tenant accounts** field blank.

The default rule must apply to all tenant accounts.

6. Leave the Bucket name drop-down selection as **applies to all buckets**.

The default rule must apply to all S3 buckets.

7. Keep the default answer, **No**, for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?"
8. Don't add advanced filters.

The default rule can't specify any filters.

9. Select **Next**.

Step 2 (Define placements) appears.

10. For Reference time, select any option.

If you kept the default answer, **No**, for the question, "Apply this rule to older object versions only?" Noncurrent time will not be included in the pull-down list. The default rule must apply all object versions.

11. Specify the placement instructions for the default rule.
 - The default rule should retain objects forever. A warning appears when you activate a new policy if the default rule does not retain objects forever. You must confirm this is the behavior you expect.
 - The default rule should create replicated copies.



Don't use a rule that creates erasure-coded copies as the default rule for a policy. Erasure-coding rules should include the **Object size (MB) greater than 200 KB** advanced filter to prevent smaller objects from being erasure-coded.

- If you are using (or you plan to enable) the global S3 Object Lock setting, the default rule must be compliant:
 - It must create at least two replicated object copies or one erasure-coded copy.
 - These copies must exist on Storage Nodes for the entire duration of each line in the placement instructions.
 - Object copies can't be saved in a Cloud Storage Pool.
 - At least one line of the placement instructions must start at day 0, using Ingest time as the reference time.
 - At least one line of the placement instructions must be "forever."

12. Look at the Retention diagram to confirm your placement instructions.

13. Select **Continue**.

Step 3 (Select ingest behavior) appears.

14. Select the ingest option to use, and select **Create**.

Manage ILM policies

Use ILM policies

An information lifecycle management (ILM) policy is an ordered set of ILM rules that determines how the StorageGRID system manages object data over time.



An ILM policy that has been incorrectly configured can result in unrecoverable data loss. Before activating an ILM policy, carefully review the ILM policy and its ILM rules, and then simulate the ILM policy. Always confirm that the ILM policy will work as intended.

Default ILM policy

When you install StorageGRID and add sites, a default ILM policy is automatically created, as follows:

- If your grid contains one site, the default policy contains a default rule that replicates two copies of each object at that site.
- If your grid contains more than one site, the default rule replicates one copy of each object at each site.


If the default policy does not meet your storage requirements, you can create your own rules and policy. See [Create an ILM rule](#) and [Create an ILM policy](#).

One or many active ILM policies?

You can have one or more active ILM policies at a time.

One policy

If your grid will use a simple data protection scheme with few tenant-specific and bucket-specific rules, use a single active ILM policy. The ILM rules can contain filters to manage different buckets or tenants.



When you have only one policy and a tenant's requirements change, you must create a new ILM policy or clone the existing policy to apply changes, simulate, and then activate the new ILM policy. Changes to the ILM policy could result in object moves that could take many days and cause system latency.

Multiple policies

To provide different quality-of-service options to tenants, you can have more than one active policy at a time. Each policy can manage specific tenants, S3 buckets, and objects. When you apply or change one policy for a specific set of tenants or objects, the policies applied to other tenants and objects are not affected.


ILM policy tags

If you want to allow tenants to easily switch between multiple data protection policies on a per-bucket basis, use multiple ILM policies with *ILM policy tags*. You assign each ILM policy to a tag, then tenants tag a bucket to apply the policy to that bucket. You can set ILM policy tags on S3 buckets only.

For example, you might have three tags named Gold, Silver, and Bronze. You can assign an ILM policy to each tag, based on how long and where that policy stores objects. Tenants can choose which policy to use by tagging their buckets. A bucket tagged Gold is managed by the Gold policy and receives the Gold level of data protection and performance.

Default ILM policy tag

A default ILM policy tag is automatically created when you install StorageGRID. Every grid must have one active policy that is assigned to the Default tag. The default policy applies to any untagged S3 buckets.



How does an ILM policy evaluate objects?

An active ILM policy controls the placement, duration, and data protection of objects.

When clients save objects to StorageGRID, the objects are evaluated against the ordered set of ILM rules in the policy, as follows:


1. If the filters for the first rule in the policy match an object, the object is ingested according to that rule's ingest behavior and stored according to that rule's placement instructions.
2. If the filters for the first rule don't match the object, the object is evaluated against each subsequent rule in the policy until a match is made.
3. If no rules match an object, the ingest behavior and placement instructions for the default rule in the policy are applied. The default rule is the last rule in a policy. The default rule must apply to all tenants, all S3 buckets, and all object versions, and can't use any advanced filters.

Example ILM policy

As an example, an ILM policy could contain three ILM rules that specify the following:

- **Rule 1: Replicated copies for Tenant A**
 - Match all objects belonging to Tenant A.
 - Store these objects as three replicated copies at three sites.
 - Objects belonging to other tenants aren't matched by Rule 1, so they are evaluated against Rule 2.
- **Rule 2: Erasure coding for objects greater than 1 MB**
 - Match all objects from other tenants, but only if they are greater than 1 MB. These larger objects are stored using 6+3 erasure coding at three sites.
 - Does not match objects 1 MB or smaller, so these objects are evaluated against Rule 3.
- **Rule 3: 2 copies 2 data centers (default)**
 - Is the last and default rule in the policy. Does not use filters.

- Make two replicated copies of all objects not matched by Rule 1 or Rule 2 (objects not belonging to Tenant A that are 1 MB or smaller).



What are active and inactive policies?

Every StorageGRID system must have at least one active ILM policy. If you want to have more than one active ILM policy, you create ILM policy tags and assign a policy to each tag. Tenants then apply tags to S3 buckets. The default policy is applied to all objects in buckets that do not have a policy tag assigned.

When you first create an ILM policy, you select one or more ILM rules and arrange them in a specific order. After you have simulated the policy to confirm its behavior, you activate it.

When you activate one ILM policy, StorageGRID uses that policy to manage all objects, including existing objects and newly ingested objects. Existing objects might be moved to new locations when the ILM rules in the new policy are implemented.

If you activate more than one ILM policy at a time, and tenants apply policy tags to S3 buckets, the objects in each bucket are managed according to the policy assigned to the tag.

A StorageGRID system tracks the history of policies that have been activated or deactivated.

Considerations for creating an ILM policy

- Only use the system-provided policy, Baseline 2 copies policy, in test systems. For StorageGRID 11.6 and earlier, the Make 2 Copies rule in this policy uses the All Storage Nodes storage pool, which contains all sites. If your StorageGRID system has more than one site, two copies of an object might be placed on the same site.



The All Storage Nodes storage pool is automatically created during the installation of StorageGRID 11.6 and earlier. If you upgrade to a later version of StorageGRID, the All Storage Nodes pool will still exist. If you install StorageGRID 11.7 or later as a new installation, the All Storage Nodes pool is not created.

- When designing a new policy, consider all of the different types of objects that might be ingested into your grid. Make sure the policy includes rules to match and place these objects as required.
- Keep the ILM policy as simple as possible. This avoids potentially dangerous situations where object data is not protected as intended when changes are made to the StorageGRID system over time.
- Make sure that the rules in the policy are in the correct order. When the policy is activated, new and existing objects are evaluated by the rules in the order listed, starting at the top. For example, if the first rule in a policy matches an object, that object will not be evaluated by any other rule.
- The last rule in every ILM policy is the default ILM rule, which can't use any filters. If an object has not been matched by another rule, the default rule controls where that object is placed and for how long it is retained.
- Before activating a new policy, review any changes that the policy is making to the placement of existing objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

Create ILM policies

Create one or more ILM policies to meet your quality-of-service requirements.

Having one active ILM policy allows you to apply the same ILM rules to all tenants and buckets.

Having multiple active ILM policies allows you to apply the appropriate ILM rules to specific tenants and buckets to fulfill multiple quality-of-service requirements.

Create an ILM policy

About this task

Before creating your own policy, verify that the [default ILM policy](#) does not meet your storage requirements.



Only use the system-provided policies, 2 copies Policy (for one-site grids) or 1 Copy per Site (for multi-site grids), in test systems. For StorageGRID 11.6 and earlier, the default rule in this policy uses the All Storage Nodes storage pool, which contains all sites. If your StorageGRID system has more than one site, two copies of an object might be placed on the same site.



If the [global S3 Object Lock setting has been enabled](#), you must ensure that the ILM policy is compliant with the requirements of buckets that have S3 Object Lock enabled. In this section, follow the instructions that mention having S3 Object Lock enabled.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [required access permissions](#).
- You have [created ILM rules](#) based on whether S3 Object Lock is enabled.


S3 Object Lock not enabled

- You have [created the ILM rules](#) you want to add to the policy. As required, you can save a policy, create additional rules, and then edit the policy to add the new rules.
- You have [created a default ILM rule](#) that does not contain any filters.

S3 Object Lock enabled

- The [global S3 Object Lock setting is already enabled](#) for the StorageGRID system.
- You have [created the compliant and non-compliant ILM rules](#) you want to add to the policy. As required, you can save a policy, create additional rules, and then edit the policy to add the new rules.
- You have [created a default ILM rule](#) for the policy that is compliant.

- Optionally, you have watched the video: [Video: ILM policies overview](#)



See also [Use ILM policies](#).

Steps

1. Select **ILM > Policies**.

If the global S3 Object Lock setting is enabled, the ILM policies page indicates which ILM rules are compliant.

2. Determine how you want to create the ILM policy.

Create new policy

- a. Select **Create policy**.

Clone existing policy

- a. Select the checkbox for the policy you want to start with, then select **Clone**.

Edit existing policy

- a. If a policy is inactive, you can edit it. Select the checkbox for the inactive policy you want to start with, then select **Edit**.

3. In the **Policy name** field, enter a unique name for the policy.
4. Optionally, in the **Reason for change** field, enter the reason you are creating a new policy.

5. To add rules to the policy, select **Select rules**. Select a rule name to view the settings for that rule.

If you are cloning a policy:

- The rules used by the policy you are cloning are selected.
- If the policy you are cloning used any rules with no filters that were not the default rule, you are prompted to remove all but one of those rules.
- If the default rule used a filter, you are prompted to select a new default rule.
- If the default rule was not the last rule, you can move the rule to the end of the new policy.

S3 Object Lock not enabled

- a. Select one default rule for the policy. To create a new default rule, select **ILM rules page**.

The default rule applies to any objects that don't match another rule in the policy. The default rule can't use any filters and is always evaluated last.



Don't use the Make 2 Copies rule as the default rule for a policy. The Make 2 Copies rule uses a single storage pool, All Storage Nodes, which contains all sites. If your StorageGRID system has more than one site, two copies of an object might be placed on the same site.

S3 Object Lock enabled

- a. Select one default rule for the policy. To create a new default rule, select **ILM rules page**.

The list of rules contains only the rules that are compliant and don't use any filters.



Don't use the Make 2 Copies rule as the default rule for a policy. The Make 2 Copies rule uses a single storage pool, All Storage Nodes, which contains all sites. If you use this rule, multiple copies of an object might be placed on the same site.

- b. If you need a different "default" rule for objects in non-compliant S3 buckets, select **Include a rule without filters for non-compliant S3 buckets**, and select one non-compliant rule that does not use a filter.

For example, you might want to use a Cloud Storage Pool to store objects in buckets that don't have S3 Object Lock enabled.



You can only select one non-compliant rule that does not use a filter.

See also [Example 7: Compliant ILM policy for S3 Object Lock](#).

6. When you are done selecting the default rule, select **Continue**.
7. For the Other rules step, select any other rules you want to add to the policy. These rules use at least one filter (tenant account, bucket name, advanced filter, or the Noncurrent reference time). Then select **Select**.

The Create a policy window now lists the rules you selected. The default rule is at the end, with the other rules above it.

If S3 Object Lock is enabled and you also selected a non-compliant "default" rule, that rule is added as the second-to-last rule in the policy.



A warning appears if any rule does not retain objects forever. When you activate this policy, you must confirm that you want StorageGRID to delete objects when the placement instructions for the default rule elapse (unless a bucket lifecycle keeps the objects for a longer time period).

8. Drag the rows for the non-default rules to determine the order in which these rules will be evaluated.

You can't move the default rule. If S3 Object Lock is enabled, you also can't move the non-compliant "default" rule if one was selected.



You must confirm that the ILM rules are in the correct order. When the policy is activated, new and existing objects are evaluated by the rules in the order listed, starting at the top.

9. As required, select **Select rules** to add or remove rules.

10. When you are done, select **Save**.

11. Repeat these steps to create additional ILM policies.

12. [Simulate an ILM policy](#). You should always simulate a policy before activating it to ensure it works as expected.

Simulate a policy

Simulate a policy on test objects before activating the policy and applying it to your production data.

Before you begin

- You know the S3 bucket/object-key for each object you want to test.

Steps

1. Using an S3 client or the [S3 Console](#), ingest the objects required to test each rule.
2. On the ILM policies page, select the checkbox for the policy, then select **Simulate**.
3. In the **Object** field, enter the S3 bucket/object-key for a test object. For example, bucket-01/filename.png.
4. If S3 versioning is enabled, optionally enter a version ID for the object in the **Version ID** field.
5. Select **Simulate**.
6. In the Simulation results section, confirm that each object was matched by the correct rule.
7. To determine which storage pool or erasure-coding profile is in effect, select the name of the matched rule to go to the rule details page.



Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

Results

Any edits to the policy's rules will be reflected in the Simulation results and show the new match and previous match. The Simulate policy window retains the objects you tested until you select either **Clear all** or the remove icon  for each object in the Simulation results list.

Related information

[Example ILM policy simulations](#)

Activate a policy

When you activate a single new ILM policy, existing objects and newly ingested objects are managed by that policy. When you activate multiple policies, ILM policy tags assigned to buckets determine the objects to be managed.

Before you activate a new policy:

1. Simulate the policy to confirm that it behaves as you expect.
2. Review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.



Errors in an ILM policy can cause unrecoverable data loss.

About this task

When you activate an ILM policy, the system distributes the new policy to all nodes. However, the new active policy might not actually take effect until all grid nodes are available to receive the new policy. In some cases, the system waits to implement a new active policy to ensure that grid objects aren't accidentally removed. Specifically:

- If you make policy changes that **increase data redundancy or durability**, those changes are implemented immediately. For example, if you activate a new policy that includes a three-copies rule instead of a two-copies rule, that policy will be implemented right away because it increases data redundancy.
- If you make policy changes that **could decrease data redundancy or durability**, those changes will not be implemented until all grid nodes are available. For example, if you activate a new policy that uses a two-copies rule instead of a three-copies rule, the new policy will appear in the Active policy tab but it will not take effect until all nodes are online and available.

Steps

Follow the steps for activating one policy or multiple policies:

Activate one policy

Follow these steps if you will have only one active policy. If you already have one or more active policies and you are activating additional policies, follow the steps for activating multiple policies.

1. When you are ready to activate a policy, select **ILM > Policies**.

Alternatively, you can activate a single policy from the **ILM > Policy tags** page.

2. On the Policies tab, select the checkbox for the policy you want to activate, then select **Activate**.
3. Follow the appropriate step:
 - If a warning message prompts you to confirm that you want to activate the policy, select **OK**.
 - If a warning message containing details about the policy appears:
 - a. Review the details to ensure the policy would manage data as expected.
 - b. If the default rule stores objects for a limited number of days, review the retention diagram and then type in that number of days into the text box.
 - c. If the default rule stores objects forever, but one or more other rules has limited retention, type **yes** in the text box.
 - d. Select **Activate policy**.

Activate multiple policies

To activate multiple policies, you must create tags and assign a policy to each tag.



When multiple tags are in use, if tenants frequently reassign policy tags to buckets, grid performance might be impacted. If you have untrusted tenants, consider using only the Default tag.

1. Select **ILM > Policy tags**.
2. Select **Create**.
3. In the Create policy tag dialog box, type a tag name and, optionally, a description for the tag.



Tag names and descriptions are visible to tenants. Choose values that will help tenants make an informed decision when selecting policy tags to assign to their buckets. For example, if the assigned policy will delete objects after a period of time, you could communicate that in the description. Do not include sensitive information in these fields.

4. Select **Create tag**.
5. In the ILM policy tags table, use the pull-down to select a policy to assign to the tag.
6. If warnings appear in the Policy limitations column, select **View policy details** to review the policy.
7. Ensure each policy would manage data as expected.
8. Select **Activate assigned policies**. Or, select **Clear changes** to remove the policy assignment.
9. In the Activate policies with new tags dialog box, review the descriptions of how each tag, policy, and rule will manage objects. Make changes as needed to ensure the policies will manage objects as expected.
10. When you are sure you want to activate the policies, type **yes** in the text box, then select **Activate**

policies.

Related information

[Example 6: Changing an ILM policy](#)

Example ILM policy simulations

The examples of ILM policy simulations provide guidelines for structuring and modifying simulations for your environment.

Example 1: Verify rules when simulating an ILM policy

This example describes how to verify rules when simulating a policy.

In this example, the **Example ILM policy** is being simulated against the ingested objects in two buckets. The policy includes three rules, as follows:

- The first rule, **Two copies, two years for bucket-a**, applies only to objects in bucket-a.
- The second rule, **EC objects > 1 MB**, applies to all buckets but filters on objects greater than 1 MB.
- The third rule, **Two copies, two data centers**, is the default rule. It does not include any filters and does not use the Noncurrent reference time.

After simulating the policy, confirm that each object was matched by the correct rule.

Simulation results				
Use this table to confirm the results of applying this policy to the selected objects.				
Clear all ?				
Object	Version ID	Rule matched	Previous match	Actions
bucket-a/bucket-a object.pdf	—	Two copies, two years for bucket-a	—	✕
bucket-b/test object greater than 1 MB.pdf	—	EC objects > 1 MB	—	✕
bucket-b/test object less than 1 MB.pdf	—	Two copies, two data centers	—	✕

In this example:

- bucket-a/bucket-a object.pdf correctly matched the first rule, which filters on objects in bucket-a.
- bucket-b/test object greater than 1 MB.pdf is in bucket-b, so it did not match the first rule. Instead, it was correctly matched by the second rule, which filters on objects greater than 1 MB.
- bucket-b/test object less than 1 MB.pdf did not match the filters in the first two rules, so it will be placed by the default rule, which includes no filters.

Example 2: Reorder rules when simulating an ILM policy

This example shows how you can reorder rules to change the results when simulating a policy.

In this example, the **Demo** policy is being simulated. This policy, which is intended to find objects that have series=x-men user metadata, includes three rules, as follows:

- The first rule, **PNGs**, filters for key names that end in .png.
- The second rule, **X-men**, applies only to objects for Tenant A and filters for series=x-men user metadata.
- The last rule, **Two copies two data centers**, is the default rule, which matches any objects that don't match the first two rules.

Steps

1. After adding the rules and saving the policy, select **Simulate**.
2. In the **Object** field, enter the S3 bucket/object-key for a test object and select **Simulate**.

The Simulation results appear, showing that the `Havok.png` object was matched by the **PNGs** rule.

Simulation results

Use this table to confirm the results of applying this policy to the selected objects.

Clear all

Object	Version ID	Rule matched	Previous match	Actions
photos/Havok.png	—	PNGs	—	X

However, `Havok.png` was meant to test the **X-men** rule.

3. To resolve the issue, reorder the rules.
 - a. Select **Finish** to close the Simulate ILM Policy window.
 - b. Select **Edit** to edit the policy.
 - c. Drag the **X-men** rule to the top of the list.
 - d. Select **Save**.
4. Select **Simulate**.

The objects you previously tested are re-evaluated against the updated policy, and the new simulation results are shown. In the example, the Rule matched column shows that the `Havok.png` object now matches the X-men metadata rule, as expected. The Previous match column shows that the PNGs rule matched the object in the previous simulation.

Simulation results

Use this table to confirm the results of applying this policy to the selected objects.

Clear all

Object	Version ID	Rule matched	Previous match	Actions
photos/Havok.png	—	X-men	PNGs	X

Example 3: Correct a rule when simulating an ILM policy

This example shows how to simulate a policy, correct a rule in the policy, and continue the simulation.

In this example, the **Demo** policy is being simulated. This policy is intended to find objects that have `series=x-men` user metadata. However, unexpected results occurred when simulating this policy against the `Beast.jpg` object. Instead of matching the X-men metadata rule, the object matched the default rule, Two copies two data centers.

Simulation results				
Use this table to confirm the results of applying this policy to the selected objects.				
<button>Clear all</button> ?				
Object	Version ID	Rule matched	Previous match	Actions
photos/Beast.jpg	—	Two copies two data centers	—	X

When a test object is not matched by the expected rule in the policy, you must examine each rule in the policy and correct any errors.

Steps

1. Select **Finish** to close the Simulate policy dialog. On the details page for the policy, select **Retention diagram**. Then select **Expand all** or **View details** for each rule as needed.
2. Review the rule's tenant account, reference time, and filtering criteria.

As an example, suppose the metadata for the X-men rule was entered as "x-men01" instead of "x-men."

3. To resolve the error, correct the rule as follows:
 - If the rule is part of the policy, you can either clone the rule or remove the rule from the policy and then edit it.
 - If the rule is part of the active policy, you must clone the rule. You can't edit or remove a rule from the active policy.
4. Perform the simulation again.

In this example, the corrected X-men rule now matches the `Beast.jpg` object based on the `series=x-men` user metadata, as expected.

Simulation results				
Use this table to confirm the results of applying this policy to the selected objects.				
<button>Clear all</button> ?				
Object	Version ID	Rule matched	Previous match	Actions
photos/Beast.jpg	—	X-men	—	X

Manage ILM policy tags

You can view ILM policy tag details, edit a tag, or remove a tag.

Before you begin

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

View ILM policy tag details

To view the details for a tag:

1. Select **ILM > Policy tags**.
2. Select the name of the policy from the table. The details page for the tag appears.
3. On the details page, view the previous history of assigned policies.
4. View a policy by selecting it.

Edit ILM policy tag



Tag names and descriptions are visible to tenants. Choose values that will help tenants make an informed decision when selecting policy tags to assign to their buckets. For example, if the assigned policy will delete objects after a period of time, you could communicate that in the description. Do not include sensitive information in these fields.

To edit the description for an existing tag:

1. Select **ILM > Policy tags**.
2. Select the checkbox for the tag, then select **Edit**.

Alternatively, select the name of the tag. The details page for the tag appears, and you can select **Edit** on that page.

3. Change the tag description as needed
4. Select **Save**.

Remove ILM policy tag

When you remove a policy tag, any buckets that are assigned that tag will have the Default policy applied.

To remove a tag:

1. Select **ILM > Policy tags**.
2. Select the checkbox for the tag, then select **Remove**. A confirmation dialog box appears.

Alternatively, select the name of the tag. The details page for the tag appears, and you can select **Remove** on that page.

3. Select **Yes** to delete the tag.

Verify an ILM policy with object metadata lookup

After you have activated an ILM policy, ingest representative test objects into the StorageGRID system, then perform an object metadata lookup to confirm that copies are being made as intended and placed in the correct locations.

Before you begin

You have an object identifier, which can be one of:

* **UUID**: The object's Universally Unique Identifier.

* **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. If the S3 bucket is versioned and you want to look up a specific version of an S3 object using the bucket and object key, you have the **version ID**.

Steps

1. Ingest the object.
2. Select **ILM > Object metadata lookup**.
3. Type the object's identifier in the **Identifier** field. You can enter a UUID, CBID, or S3 bucket/object-key.
4. Optionally, enter a version ID for the object (S3 only).
5. Select **Look Up**.

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

- System metadata, such as object ID (UUID), result type (object, delete marker, S3 bucket), and logical size of the object. Refer to the example screenshot below for more details.
 - 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.
6. Confirm that the object is stored in the correct location or locations and that it's the correct type of copy.

If the Audit option is enabled, you can also monitor the audit log for the ORLM Object Rules Met message. The ORLM audit message can provide you with more information about the status of the ILM evaluation process, but it can't give you information about the correctness of the object data's placement or the completeness of the ILM policy. You must evaluate this yourself. For details, see [Review audit logs](#).

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



The following screenshot is an example. Your results will vary depending on your StorageGRID version.

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": {
      "PAIIS": "2",

```

Related information

[Use S3 REST API](#)

Work with ILM policies and ILM rules

As your storage requirements change, you might need to put additional policies in place or modify the ILM rules associated with a policy. You can view ILM metrics to determine system performance.

Before you begin

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

View ILM policies

To view active and inactive ILM policies and policy activation history:

1. Select **ILM > Policies**.
2. Select **Policies** to view a list of active and inactive policies. The table lists the name of each policy, the tags the policy is assigned to, and whether the policy is active or inactive.
3. Select **Activation history** to view a list of activation start and end dates for policies.
4. Select a policy name to view the details for the policy.



If you view the details for a policy whose status is Edited or Deleted, a message appears explaining that you are viewing the version of the policy that was active for the specified time span and has since been edited or deleted.

Edit an ILM policy

You can only edit an inactive policy. If you want to edit an active policy, deactivate it or create a clone and edit the clone.

To edit a policy:

1. Select **ILM > Policies**.
2. Select the checkbox for the policy you want to edit, then select **Edit**.
3. Edit the policy by following the instructions in [Create ILM policies](#).
4. Simulate the policy before you re-activate it.



An ILM policy that has been incorrectly configured can result in unrecoverable data loss. Before activating an ILM policy, carefully review the ILM policy and its ILM rules, and then simulate the ILM policy. Always confirm that the ILM policy will work as intended.

Clone an ILM policy

To clone an ILM policy:

1. Select **ILM > Policies**.
2. Select the checkbox for the policy you want to clone, then select **Clone**.
3. Create a new policy starting with the policy you've cloned by following the instructions in [Create ILM policies](#).



An ILM policy that has been incorrectly configured can result in unrecoverable data loss. Before activating an ILM policy, carefully review the ILM policy and its ILM rules, and then simulate the ILM policy. Always confirm that the ILM policy will work as intended.

Remove an ILM policy

You can only remove an ILM policy if it is inactive. To remove a policy:

1. Select **ILM > Policies**.
2. Select the checkbox for the inactive policy you want to remove.
3. Select **Remove**.

View ILM rule details

To view the details for an ILM rule, including the retention diagram and placement instructions for the rule:

1. Select **ILM > Rules**.
2. Select the name of the rule whose details you want to view. Example:

2 copies 2 data centers

Compliant: No
Ingest behavior: Strict
Reference time: Noncurrent time

[Clone](#) [Edit](#) [Remove](#)

[Rule detail](#) [Used in policies](#)

Time period and placements

[Retention diagram](#) [Placement instructions](#)

Sort placements by [Time period](#) [Storage pool](#) ● Replicated copy ● Erasure-coded (EC) copy

Rule analysis:

- Objects processed by this rule will not be deleted by ILM.

Reference time: **Noncurrent time** Ingest behavior: **Strict**

Day 0

Day 0 - forever

2 replicated copies - Data Center 1

EC 2+1 - Data Center 1

Duration Forever

Additionally, you can use the details page to clone, edit, or remove a rule. You can't edit or remove a rule if it's used in any policy.

Clone an ILM rule

You can clone an existing rule if you want to create a new rule that uses some of the settings of the existing rule. If you need to edit a rule that's used in any policy, you clone the rule instead and make changes to the clone. After you make changes to the clone, you can remove the original rule from the policy and replace it with the modified version as required.



You can't clone an ILM rule if it was created using StorageGRID version 10.2 or earlier.

Steps

1. Select **ILM > Rules**.
2. Select the checkbox for the rule you want to clone, then select **Clone**. Alternatively, select the rule name, then select **Clone** from the rule details page.
3. Update the cloned rule by following the steps for [editing an ILM rule](#) and [using advanced filters in ILM rules](#).

When cloning an ILM rule, you must enter a new name.

Edit an ILM rule

You might need to edit an ILM rule to change a filter or placement instruction.

You can't edit a rule if it is used in any ILM policy. Instead, you can [clone the rule](#) and make any required changes to the cloned copy.



An ILM policy that has been incorrectly configured can result in unrecoverable data loss. Before activating an ILM policy, carefully review the ILM policy and its ILM rules, and then simulate the ILM policy. Always confirm that the ILM policy will work as intended.

Steps

1. Select **ILM > Rules**.
2. Confirm that the rule you want to edit is not used in any ILM policy.
3. If the rule you want to edit is not in use, select the checkbox for the rule and select **Actions > Edit**. Alternatively, select the name of the rule, then select **Edit** on the rule details page.
4. Complete the steps of the Edit ILM rule wizard. As necessary, follow the steps for [creating an ILM rule](#) and [using advanced filters in ILM rules](#).

When editing an ILM rule, you can't change its name.

Remove an ILM rule

To keep the list of current ILM rules manageable, remove any ILM rules that you aren't likely to use.

Steps

To remove an ILM rule that is currently used in an active policy:

1. Clone the policy.
2. Remove the ILM rule from the policy clone.
3. Save, simulate, and activate the new policy to make sure objects are protected as expected.
4. Go to the steps for removing an ILM rule that is currently used in an inactive policy.

To remove an ILM rule that is currently used in an inactive policy:

1. Select the inactive policy.
2. Remove the ILM rule from the policy or [remove the policy](#).
3. Go to the steps for removing an ILM rule that is not currently used.

To remove an ILM rule that is not currently used:

1. Select **ILM > Rules**.
2. Confirm that the rule you want to remove is not used in any policy.
3. If the rule you want to remove is not in use, select the rule and select **Actions > Remove**. You can select multiple rules and remove all of them at the same time.
4. Select **Yes** to confirm that you want to remove the ILM rule.

View ILM metrics

You can view metrics for ILM, such as the number of objects in the queue and the evaluation rate. You can monitor these metrics to determine system performance. A large queue or evaluation rate might indicate that the system is not able to keep up with the ingest rate, the load from the client applications is excessive, or that some abnormal condition exists.

Steps


1. Select **Dashboard > ILM**.



Because the dashboard can be customized, the ILM tab might not be available.

2. Monitor the metrics on the ILM tab.

You can select the question mark  to see a description of the items on the ILM tab.



Use S3 Object Lock

Manage objects with S3 Object Lock

As a grid administrator, you can enable S3 Object Lock for your StorageGRID system and implement a compliant ILM policy to help ensure that objects in specific S3 buckets aren't deleted or overwritten for a specified amount of time.

What is S3 Object Lock?

The StorageGRID S3 Object Lock feature is an object-protection solution that is equivalent to S3 Object Lock in Amazon Simple Storage Service (Amazon S3).

When the global S3 Object Lock setting is enabled for a StorageGRID system, an S3 tenant account can create buckets with or without S3 Object Lock enabled. If a bucket has S3 Object Lock enabled, bucket

versioning is required and is enabled automatically.

A bucket without S3 Object Lock can only have objects without retention settings specified. No ingested objects will have retention settings.

A bucket with S3 Object Lock can have objects with and without retention settings specified by S3 client applications. Some objects ingested will have retention settings.

A bucket with S3 Object Lock and default retention configured can have uploaded objects with retention settings specified and new objects without retention settings. The new objects use the default setting, because the retention setting hasn't been configured at the object-level.

Effectively, all newly ingested objects have retention settings when default retention is configured. Existing objects without object retention settings remain unaffected.

Retention modes

The StorageGRID S3 Object Lock feature supports two retention modes to apply different levels of protection to objects. These modes are equivalent to the Amazon S3 retention modes.

- In compliance mode:
 - The object can't be deleted until its retain-until-date is reached.
 - The object's retain-until-date can be increased, but it can't be decreased.
 - The object's retain-until-date can't be removed until that date is reached.
- In governance mode:
 - Users with special permission can use a bypass header in requests to modify certain retention settings.
 - These users can delete an object version before its retain-until-date is reached.
 - These users can increase, decrease, or remove an object's retain-until-date.

Retention settings for object versions

If a bucket is created with S3 Object Lock enabled, users can use the S3 client application to optionally specify the following retention settings for each object that is added to the bucket:

- **Retention mode:** Either compliance or governance.
- **Retain-until-date:** If an object version's retain-until-date is in the future, the object can be retrieved, but it can't be deleted.
- **Legal hold:** Applying a legal hold to an object version immediately locks that object. For example, you might need to put a legal hold on an object that is related to an investigation or legal dispute. A legal hold has no expiration date, but remains in place until it is explicitly removed. Legal holds are independent of the retain-until-date.



If an object is under a legal hold, no one can delete the object, regardless of its retention mode.

For details on the object settings, see [Use S3 REST API to configure S3 Object Lock](#).

Default retention setting for buckets

If a bucket is created with S3 Object Lock enabled, users can optionally specify the following default settings for the bucket:

- **Default retention mode:** Either compliance or governance.
- **Default retention period:** How long new object versions added to this bucket should be retained, starting from the day they are added.

The default bucket settings apply only to new objects that don't have their own retention settings. Existing bucket objects aren't affected when you add or change these default settings.

See [Create an S3 bucket](#) and [Update S3 Object Lock default retention](#).

Comparing S3 Object Lock to legacy Compliance

The S3 Object Lock replaces the Compliance feature that was available in earlier StorageGRID versions. Because the S3 Object Lock feature conforms to Amazon S3 requirements, it deprecates the proprietary StorageGRID Compliance feature, which is now referred to as "legacy Compliance."



The global Compliance setting is deprecated. If you enabled this setting using a previous version of StorageGRID, the S3 Object Lock setting is enabled automatically. You can continue to use StorageGRID to manage the settings of existing compliant buckets; however, you can't create new compliant buckets. For details, see [NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#).

If you used the legacy Compliance feature in a previous version of StorageGRID, refer to the following table to learn how it compares to the S3 Object Lock feature in StorageGRID.

	S3 Object Lock	Compliance (legacy)
How is the feature enabled globally?	From the Grid Manager, select CONFIGURATION > System > S3 Object Lock .	No longer supported.
How is the feature enabled for a bucket?	Users must enable S3 Object Lock when creating a new bucket using the Tenant Manager, the Tenant Management API, or the S3 REST API.	No longer supported.
Is bucket versioning supported?	Yes. Bucket versioning is required and is enabled automatically when S3 Object Lock is enabled for the bucket.	No.
How is object retention set?	Users can set a retain-until-date for each object version, or they can set a default retention period for each bucket.	Users must set a retention period for the entire bucket. The retention period applies to all objects in the bucket.

	S3 Object Lock	Compliance (legacy)
Can the retention period be changed?	• In compliance mode, the retain-until-date for an object version can be increased but never decreased. • In governance mode, users with special permissions can decrease or even remove an object's retention settings.	A bucket's retention period can be increased but never decreased.
Where is legal hold controlled?	Users can place a legal hold or lift a legal hold for any object version in the bucket.	A legal hold is placed on the bucket and affects all objects in the bucket.
When can objects be deleted?	• In compliance mode, an object version can be deleted after the retain-until-date is reached, assuming the object is not under legal hold. • In governance mode, users with special permissions can delete an object before its retain-until-date is reached, assuming the object is not under legal hold.	An object can be deleted after the retention period expires, assuming the bucket is not under legal hold. Objects can be deleted automatically or manually.
Is bucket lifecycle configuration supported?	Yes	No

S3 Object Lock tasks

As a grid administrator, you must coordinate closely with tenant users to ensure that the objects are protected in a manner that satisfies their retention requirements.



Applying tenant settings across the grid could take 15 minutes or longer based on network connectivity, node status, and Cassandra operations.

The following lists for grid administrators and tenant users contain the high-level tasks for using the S3 Object Lock feature.

Grid administrator

- Enable global S3 Object Lock setting for entire StorageGRID system.
- Ensure that information lifecycle management (ILM) policies are *compliant*; that is, they meet the [requirements of buckets with S3 Object Lock enabled](#).
- As needed, allow a tenant to use Compliance as the retention mode. Otherwise, only Governance mode is allowed.
- As needed, set a maximum retention period for a tenant.

Tenant user

- Review considerations for buckets and objects with S3 Object Lock.
- As needed, contact grid administrator to enable global S3 Object Lock setting and set permissions.
- Create buckets with S3 Object Lock enabled.
- Optionally, configure default retention settings for a bucket:
 - Default retention mode: Governance or Compliance, if allowed by the grid administrator.
 - Default retention period: Must be less than or equal to maximum retention period set by grid administrator.
- Use the S3 client application to add objects and optionally set object-specific retention:
 - Retention mode. Governance or Compliance, if allowed by the grid administrator.
 - Retain Until Date: Must be less than or equal to what is allowed by the maximum retention period set by grid administrator.

Requirements for S3 Object Lock

You must review the requirements for enabling the global S3 Object Lock setting, the requirements for creating compliant ILM rules and ILM policies, and the restrictions StorageGRID places on buckets and objects that use S3 Object Lock.

Requirements for using the global S3 Object Lock setting

- You must enable the global S3 Object Lock setting using the Grid Manager or the Grid Management API before any S3 tenant can create a bucket with S3 Object Lock enabled.
- Enabling the global S3 Object Lock setting allows all S3 tenant accounts to create buckets with S3 Object Lock enabled.
- After you enable the global S3 Object Lock setting, you can't disable the setting.
- You can't enable the global S3 Object Lock unless the default rule in all active ILM policies is *compliant* (that is, the default rule must comply with the requirements of buckets with S3 Object Lock enabled).
- When the global S3 Object Lock setting is enabled, you can't create a new ILM policy or activate an existing ILM policy unless the default rule in the policy is compliant. After the global S3 Object Lock setting has been enabled, the ILM rules and ILM policies pages indicate which ILM rules are compliant.

Requirements for compliant ILM rules

If you want to enable the global S3 Object Lock setting, you must ensure that the default rule in all active ILM policies is compliant. A compliant rule satisfies the requirements of both buckets with S3 Object Lock enabled and any existing buckets that have legacy Compliance enabled:

- It must create at least two replicated object copies or one erasure-coded copy.
- These copies must exist on Storage Nodes for the entire duration of each line in the placement instructions.
- Object copies can't be saved in a Cloud Storage Pool.
- At least one line of the placement instructions must start at day 0, using **Ingest time** as the reference time.
- At least one line of the placement instructions must be "forever."

Requirements for ILM policies

When the global S3 Object Lock setting is enabled, active and inactive ILM policies can include both compliant and non-compliant rules.

- The default rule in an active or inactive ILM policy must be compliant.
- Non-compliant rules only apply to objects in buckets that don't have S3 Object Lock enabled or that don't have the legacy Compliance feature enabled.
- Compliant rules can apply to objects in any bucket; S3 Object Lock or legacy Compliance does not need to be enabled for the bucket.

Example of a compliant ILM policy for S3 Object Lock

Requirements for buckets with S3 Object Lock enabled

- If the global S3 Object Lock setting is enabled for the StorageGRID system, you can use the Tenant Manager, the Tenant Management API, or the S3 REST API to create buckets with S3 Object Lock enabled.
- If you plan to use S3 Object Lock, you must enable S3 Object Lock when you create the bucket. You can't enable S3 Object Lock for an existing bucket.
- When S3 Object Lock is enabled for a bucket, StorageGRID automatically enables versioning for that bucket. You can't disable S3 Object Lock or suspend versioning for the bucket.
- Optionally, you can specify a default retention mode and retention period for each bucket using the Tenant Manager, the Tenant Management API, or the S3 REST API. The bucket's default retention settings apply only to new objects added to the bucket that don't have their own retention settings. You can override these default settings by specifying a retention mode and retain-until-date for each object version when it is uploaded.
- Bucket lifecycle configuration is supported for buckets with S3 Object Lock enabled.
- CloudMirror replication is not supported for buckets with S3 Object Lock enabled.

Requirements for objects in buckets with S3 Object Lock enabled

- To protect an object version, you can specify default retention settings for the bucket, or you can specify retention settings for each object version. Object-level retention settings can be specified using the S3 client application or the S3 REST API.
- Retention settings apply to individual object versions. An object version can have both a retain-until-date and a legal hold setting, one but not the other, or neither. Specifying a retain-until-date or a legal hold setting for an object protects only the version specified in the request. You can create new versions of the object, while the previous version of the object remains locked.

Lifecycle of objects in buckets with S3 Object Lock enabled

Each object that is saved in a bucket with S3 Object Lock enabled goes through these stages:

1. Object ingest

When an object version is added to bucket that has S3 Object Lock enabled, retention settings are applied as follows:

- If retention settings are specified for the object, the object-level settings are applied. Any default bucket settings are ignored.

- If no retention settings are specified for the object, the default bucket settings are applied, if they exist.
- If no retention settings are specified for the object or the bucket, the object is not protected by S3 Object Lock.

If retention settings are applied, both the object and any S3 user-defined metadata are protected.

2. Object retention and deletion

Multiple copies of each protected object are stored by StorageGRID for the specified retention period. The exact number and type of object copies and the storage locations are determined by the compliant rules in the active ILM policies. Whether a protected object can be deleted before its retain-until-date is reached depends on its retention mode.

- If an object is under a legal hold, no one can delete the object, regardless of its retention mode.

Related information

- [Create an S3 bucket](#)
- [Update S3 Object Lock default retention](#)
- [Use S3 REST API to configure S3 Object Lock](#)
- [Example 7: Compliant ILM policy for S3 Object Lock](#)

Enable S3 Object Lock globally

If an S3 tenant account needs to comply with regulatory requirements when saving object data, you must enable S3 Object Lock for your entire StorageGRID system. Enabling the global S3 Object Lock setting allows any S3 tenant user to create and manage buckets and objects with S3 Object Lock.

Before you begin

- You have the [Root access permission](#).
- You are signed in to the Grid Manager using a [supported web browser](#).
- You have reviewed the S3 Object Lock workflow, and you understand the considerations.
- You have confirmed that the default rule in the active ILM policy is compliant. See [Create a default ILM rule](#) for details.

About this task

A grid administrator must enable the global S3 Object Lock setting to allow tenant users to create new buckets that have S3 Object Lock enabled. After this setting is enabled, it can't be disabled.

Review the compliance settings of existing tenants after you enable the global S3 Object Lock setting. When you enable this setting, the S3 Object Lock per-tenant settings depend on the StorageGRID release at the time the tenant was created.



The global Compliance setting is deprecated. If you enabled this setting using a previous version of StorageGRID, the S3 Object Lock setting is enabled automatically. You can continue to use StorageGRID to manage the settings of existing compliant buckets; however, you can't create new compliant buckets. For details, see [NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#).

Steps

1. Select **CONFIGURATION > System > S3 Object Lock**.

The S3 Object Lock Settings page appears.

2. Select **Enable S3 Object Lock**.

3. Select **Apply**.

A confirmation dialog box appears and reminds you that you can't disable S3 Object Lock after it is enabled.

4. If you are sure you want to permanently enable S3 Object Lock for your entire system, select **OK**.

When you select **OK**:

- If the default rule in the active ILM policy is compliant, S3 Object Lock is now enabled for the entire grid and can't be disabled.
- If the default rule is not compliant, an error appears. You must create and activate a new ILM policy that includes a compliant rule as its default rule. Select **OK**. Then, create a new policy, simulate it, and activate it. See [Create ILM policy](#) for instructions.

Resolve consistency errors when updating the S3 Object Lock or legacy Compliance configuration

If a data center site or multiple Storage Nodes at a site become unavailable, you might need to help S3 tenant users apply changes to the S3 Object Lock or legacy Compliance configuration.

Tenant users who have buckets with S3 Object Lock (or legacy Compliance) enabled can change certain settings. For example, a tenant user using S3 Object Lock might need to put an object version under legal hold.

When a tenant user updates the settings for an S3 bucket or an object version, StorageGRID attempts to immediately update the bucket or object metadata across the grid. If the system is unable to update the metadata because a data center site or multiple Storage Nodes are unavailable, it returns an error:

```
503: Service Unavailable
Unable to update compliance settings because the settings can't be
consistently applied on enough storage services. Contact your grid
administrator for assistance.
```

To resolve this error, follow these steps:

1. Attempt to make all Storage Nodes or sites available again as soon as possible.
2. If you are unable to make enough of the Storage Nodes at each site available, contact technical support, who can help you recover nodes and ensure that changes are consistently applied across the grid.
3. Once the underlying issue has been resolved, remind the tenant user to retry their configuration changes.

Related information

- [Use a tenant account](#)
- [Use S3 REST API](#)

- [Recover and maintain](#)

Example ILM rules and policies

Example 1: ILM rules and policy for object storage

You can use the following example rules and policy as a starting point when defining an ILM policy to meet your object protection and retention requirements.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

ILM rule 1 for example 1: Copy object data to two sites

This example ILM rule copies object data to storage pools in two sites.

Rule definition	Example value
One-site storage pools	Two storage pools, each containing different sites, named Site 1 and Site 2.
Rule name	Two Copies Two Sites
Reference time	Ingest time
Placements	On Day 0 to forever, keep one replicated copy at Site 1 and one replicated copy at Site 2.

The Rule analysis section of the Retention diagram states:

- StorageGRID site-loss protection will apply for the duration of this rule.
- Objects processed by this rule will not be deleted by ILM.

Reference time ⓘ

Ingest time

Time period and placements

Sort by start date

If you want a rule to apply only to specific objects, select **Previous** and add advanced filters. When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the filter.

Time period 1
From Day 0
store forever

Store objects by replicating
1
copies at Site 1

and store objects by replicating
1
copies at Site 2

Add other type or location

Add another time period

Retention diagram

● Replicated copy

Rule analysis:

- StorageGRID site-loss protection will apply for the duration of this rule.
- Objects processed by this rule will not be deleted by ILM.

Reference time: Ingest time

Day 0

Day 0 - forever

1 replicated copy - Site 1
1 replicated copy - Site 2

Duration Forever

ILM rule 2 for example 1: Erasure-coding profile with bucket matching

This example ILM rule uses an erasure-coding profile and an S3 bucket to determine where and how long the object is stored.

Rule definition	Example value
Storage pool with multiple sites	• One storage pool across three sites (Sites 1, 2, 3) • Use 6+3 erasure-coding scheme
Rule name	S3 Bucket finance-records
Reference time	Ingest time
Placements	For objects in the S3 bucket named finance-records, create one erasure-coded copy in the pool specified by the erasure-coding profile. Keep this copy forever.

Time period and placements

Sort by start date

If you want a rule to apply only to specific objects, select **Previous** and add advanced filters. When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the filter.

Time period 1

From Day 0

store forever

×

Store objects by

erasure coding

using 6+3 EC scheme at Sites 1, 2, 3

✎

×

Add other type or location

Add another time period

Retention diagram

Erasure-coded (EC) copy

Rule analysis:

- StorageGRID site-loss protection will apply for the duration of this rule.
- Objects processed by this rule will not be deleted by ILM.

Reference time: Ingest time

Day 0

Day 0 - forever

EC 6+3 - Sites 1, 2, 3

Forever

ILM policy for example 1

In practice, most ILM policies are simple, even though the StorageGRID system allows you to design sophisticated and complex ILM policies.

A typical ILM policy for a multi-site grid might include ILM rules such as the following:

- At ingest, store all objects belonging to the S3 bucket named `finance-records` in a storage pool that contains three sites. Use 6+3 erasure coding.
- If an object does not match the first ILM rule, use the policy's default ILM rule, Two Copies Two Data Centers, to store one copy of that object in Site 1, and one copy in Site 2.

Proposed policy name

Object Storage Policy

Reason for change

example 1

Manage rules

1. Select the rules you want to add to the policy.

2. Determine the order in which the rules will be evaluated by dragging and dropping the rows. The default rule will be automatically placed at the end of the policy and cannot be moved.

Select rules

Rule order	Rule name	Filters
1	S3 Bucket finance-records	Tenant is Finance Bucket name is finance-records
Default	Two Copies Two Data Centers	—

Related information

- [Use ILM policies](#)
- [Create ILM policies](#)

Example 2: ILM rules and policy for EC object size filtering

You can use the following example rules and policy as starting points to define an ILM policy that filters by object size to meet recommended EC requirements.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

ILM rule 1 for example 2: Use EC for objects greater than 1 MB

This example ILM rule erasure codes objects that are greater than 1 MB.



Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

Rule definition	Example value
Rule name	EC Only Objects > 1 MB
Reference time	Ingest time
Advanced filter for Object size	Object size greater than 1 MB
Placements	Create a 2+1 erasure-coded copy using three sites

Filter group 1 Objects with all of following metadata will be evaluated by this rule: ×

Object size ▼

greater than ▼

1 ⬆ ⬇ ⬆

MB ▼ ×

ILM rule 2 for example 2: Two replicated copies

This example ILM rule creates two replicated copies and does not filter by object size. This rule is the default rule for the policy. Because the first rule filters out all objects greater than 1 MB, this rule only applies to objects that are 1 MB or smaller.

Rule definition	Example value
Rule name	Two Replicated Copies
Reference time	Ingest time
Advanced filter for Object size	None

Rule definition	Example value
Placements	On Day 0 to forever, keep one replicated copy at Site 1 and one replicated copy at Site 2.

ILM policy for example 2: Use EC for objects greater than 1 MB

This example ILM policy includes two ILM rules:

- The first rule erasure codes all objects that are greater than 1 MB.
- The second (default) ILM rule creates two replicated copies. Because objects greater than 1 MB have been filtered out by rule 1, rule 2 only applies to objects that are 1 MB or smaller.

Example 3: ILM rules and policy for better protection for image files

You can use the following example rules and policy to ensure that images greater than 1 MB are erasure-coded and that two copies are made of smaller images.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

ILM rule 1 for example 3: Use EC for image files greater than 1 MB

This example ILM rule uses advanced filtering to erasure code all image files greater than 1 MB.



Erasure coding is best suited for objects greater than 1 MB. Don't use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

Rule definition	Example value
Rule name	EC Image Files > 1 MB
Reference time	Ingest time
Advanced filter for Object size	Object size greater than 1 MB
Advanced filters for Key	• Ends with .jpg • Ends with .png
Placements	Create a 2+1 erasure-coded copy using three sites

Filter group 1
Objects with all of following metadata will be evaluated by this rule:

Object size

greater than

1

MB

X

and

Key

ends with

.jpg

X

or

Filter group 2
Objects with all of following metadata will be evaluated by this rule:

Object size

greater than

1

MB

X

and

Key

ends with

.png

X

Because this rule is configured as the first rule in the policy, the erasure-coding placement instruction only applies to .jpg and .png files that are greater than 1 MB.

ILM rule 2 for example 3: Create 2 replicated copies for all remaining image files

This example ILM rule uses advanced filtering to specify that smaller image files be replicated. Because the first rule in the policy has already matched image files greater than 1 MB, this rule applies to image files that are 1 MB or smaller.

Rule definition	Example value
Rule name	2 Copies for Image Files
Reference time	Ingest time
Advanced filters for Key	• Ends with .jpg • Ends with .png
Placements	Create 2 replicated copies in two storage pools

ILM policy for example 3: Better protection for image files

This example ILM policy includes three rules:

- The first rule erasure codes all image files greater than 1 MB.
- The second rule creates two copies of any remaining image files (that is, images that are 1 MB or smaller).
- The default rule applies to all remaining objects (that is, any non-image files).

Rule order	Rule name	Filters
1	↕ EC image files > 1 MB	Object size is greater than 1 MB
2	↕ 2 copies for small images	Object size is less than or equal to 200 KB
Default	Default rule	—

Example 4: ILM rules and policy for S3 versioned objects

If you have an S3 bucket with versioning enabled, you can manage the noncurrent object versions by including rules in your ILM policy that use "Noncurrent time" as the reference time.



If you specify a limited retention time for objects, those objects will be deleted permanently after the time period is reached. Make sure you understand how long the objects will be retained.

As this example shows, you can control the amount of storage used by versioned objects by using different placement instructions for noncurrent object versions.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.



To perform ILM policy simulation on a noncurrent version of an object, you must know the object version's UUID or CBID. To find the UUID and CBID, use [object metadata lookup](#) while the object is still current.

Related information

[How objects are deleted](#)

ILM rule 1 for example 4: Save three copies for 10 years

This example ILM rule stores a copy of each object at three sites for 10 years.

This rule applies to all objects, whether or not they are versioned.

Rule definition	Example value
Storage pools	Three storage pools, each consisting of different data centers, named Site 1, Site 2, and Site 3.
Rule name	Three Copies Ten Years

Rule definition	Example value
Reference time	Ingest time
Placements	On Day 0, keep three replicated copies for 10 years (3,652 days), one in Site 1, one in Site 2, and one in Site 3. At the end of 10 years, delete all copies of the object.

ILM rule 2 for example 4: Save two copies of noncurrent versions for 2 years

This example ILM rule stores two copies of the noncurrent versions of an S3 versioned object for 2 years.

Because ILM rule 1 applies to all versions of the object, you must create another rule to filter out any noncurrent versions.

To create a rule that uses "Noncurrent time" as the reference time, select **Yes** for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?" in Step 1 (Enter details) of the Create an ILM rule wizard. When you select **Yes**, *Noncurrent time* is automatically selected for the reference time, and you can't select a different reference time.

1 Enter details

2 Define placements

3 Select ingest behavior

Rule name

Older Object Versions: Two Copies Two Years

Description (optional)

Older versions only

Basic filters (optional)

Specify which tenant accounts and buckets this rule applies to.

Tenant accounts ?

Select tenant accounts

Bucket name ?

matches all

Apply this rule to older object versions only (in S3 buckets with versioning enabled)? ?

☐ No
☒ Yes

In this example, only two copies of the noncurrent versions are stored, and those copies will be stored for two years.

Rule definition	Example value
Storage Pools	Two storage pools, each at different data centers, Site 1 and Site 2.
Rule name	Noncurrent Versions: Two Copies Two Years
Reference time	Noncurrent time Automatically selected when you select Yes for the question, "Apply this rule to older object versions only (in S3 buckets with versioning enabled)?" in the Create an ILM rule wizard.
Placements	On Day 0 relative to noncurrent time (that is, starting from the day the object version becomes the noncurrent version), keep two replicated copies of the noncurrent object versions for 2 years (730 days), one in Site 1 and one in Site 2. At the end of 2 years, delete the noncurrent versions.

ILM policy for example 4: S3 versioned objects

If you want to manage older versions of an object differently than the current version, rules that use "Noncurrent time" as the reference time must appear in the ILM policy before rules that apply to the current object version.

An ILM policy for S3 versioned objects might include ILM rules such as the following:

- Keep any older (noncurrent) versions of each object for 2 years, starting from the day the version became noncurrent.



The "Noncurrent time" rules must appear in the policy before the rules that apply to the current object version. Otherwise, the noncurrent object versions will never be matched by the "Noncurrent time" rule.

- At ingest, create three replicated copies and store one copy at each of three sites. Keep copies of the current object version for 10 years.

When you simulate the example policy, you would expect test objects to be evaluated as follows:

- Any noncurrent object versions would be matched by the first rule. If a noncurrent object version is older than 2 years, it is permanently deleted by ILM (all copies of the noncurrent version removed from the grid).
- The current object version would be matched by the second rule. When the current object version has been stored for 10 years, the ILM process adds a delete marker as the current version of the object, and it makes the previous object version "noncurrent". The next time ILM evaluation occurs, this noncurrent version is matched by the first rule. As a result, the copy at Site 3 is purged and the two copies at Site 1 and Site 2 are stored for 2 more years.

Example 5: ILM rules and policy for Strict ingest behavior

You can use a location filter and the Strict ingest behavior in a rule to prevent objects from being saved at a particular data center location.

In this example, a Paris-based tenant does not want to store some objects outside of the EU because of regulatory concerns. Other objects, including all objects from other tenant accounts, can be stored at either the Paris data center or the US data center.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

Related information

- [Ingest options](#)
- [Create ILM rule: Select ingest behavior](#)

ILM rule 1 for example 5: Strict ingest to guarantee Paris data center

This example ILM rule uses the Strict ingest behavior to guarantee that objects saved by a Paris-based tenant to S3 buckets with the region set to eu-west-3 region (Paris) are never stored at the US data center.

This rule applies to objects that belong to the Paris tenant and that have the S3 bucket region set to eu-west-3 (Paris).

Rule definition	Example value
Tenant account	Paris tenant
Advanced filter	Location constraint equals eu-west-3
Storage pools	Site 1 (Paris)
Rule name	Strict ingest to guarantee Paris data center
Reference time	Ingest time
Placements	On Day 0, keep two replicated copies forever in Site 1 (Paris)
Ingest behavior	Strict. Always use this rule’s placements on ingest. Ingest fails if it is not possible to store two copies of the object at the Paris data center.

Strict ingest to guarantee Paris data center

Compliant: Yes
 Used in active policy: No
 Used in proposed policy: No

Ingest behavior: Strict
 Reference time: Ingest time

[Clone](#) [Edit](#) [Remove](#)

Filters

This rule applies if:


- Tenant is Paris tenant

And it only applies if objects have this metadata:

- Location constraint is eu-west-3

Time period and placements

Retention diagram [Placement instructions](#)



ILM rule 2 for example 5: Balanced ingest for other objects

This example ILM rule uses the Balanced ingest behavior to provide optimum ILM efficiency for any objects not matched by the first rule. Two copies of all objects matched by this rule will be stored—one at the US data center and one at the Paris data center. If the rule can't be satisfied immediately, interim copies are stored at any available location.

This rule applies to objects that belong to any tenant and any region.

Rule definition	Example value
Tenant account	Ignore
Advanced filter	<i>Not specified</i>
Storage pools	Site 1 (Paris) and Site 2 (US)
Rule name	2 Copies 2 Data Centers
Reference time	Ingest time
Placements	On Day 0, keep two replicated copies forever at two data centers

Rule definition	Example value
Ingest behavior	Balanced. Objects that match this rule are placed according to the rule's placement instructions if possible. Otherwise, interim copies are made at any available location.

ILM policy for example 5: Combining ingest behaviors

The example ILM policy includes two rules that have different ingest behaviors.

An ILM policy that uses two different ingest behaviors might include ILM rules such as the following:

- Store objects that belong to the Paris tenant and that have the S3 bucket region set to eu-west-3 (Paris) only in the Paris data center. Fail ingest if the Paris data center is not available.
- Store all other objects (including those that belong to the Paris tenant but that have a different bucket region) in both the US data center and the Paris data center. Make interim copies in any available location if the placement instruction can't be satisfied.

When you simulate the example policy, you expect test objects to be evaluated as follows:

- Any objects that belong to the Paris tenant and that have the S3 bucket region set to eu-west-3 are matched by the first rule and are stored at the Paris data center. Because the first rule uses Strict ingest, these objects are never stored at the US data center. If the Storage Nodes at the Paris data center aren't available, ingest fails.
- All other objects are matched by the second rule, including objects that belong to the Paris tenant and that don't have the S3 bucket region set to eu-west-3. One copy of each object is saved at each data center. However, because the second rule uses Balanced ingest, if one data center is unavailable, two interim copies are saved at any available location.

Example 6: Change an ILM policy

If your data protection needs to be changed or you add new sites, you can create and activate a new ILM policy.

Before changing a policy, you must understand how changes in ILM placements can temporarily affect the overall performance of a StorageGRID system.

In this example, a new StorageGRID site has been added in an expansion, and a new active ILM policy needs to be implemented to store data at the new site. To implement a new active policy, first [create a policy](#). Afterward, you must [simulate](#) and then [activate](#) the new policy.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

How changing an ILM policy affects performance

When you activate a new ILM policy, the performance of your StorageGRID system might be temporarily affected, especially if the placement instructions in the new policy require many existing objects to be moved to new locations.

When you activate a new ILM policy, StorageGRID uses it to manage all objects, including existing objects and

newly ingested objects. Before activating a new ILM policy, review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object’s location might result in temporary resource issues when the new placements are evaluated and implemented.

To ensure a new ILM policy does not affect the placement of existing replicated and erasure-coded objects, you can [create an ILM rule with an ingest time filter](#). For example, **Ingest time is on or after <date and time>**, so that the new rule applies only to objects ingested on or after the date and time specified.

The types of ILM policy changes that can temporarily affect StorageGRID performance include the following:

- Applying a different erasure-coding profile to existing erasure-coded objects.



StorageGRID considers each erasure-coding profile to be unique and does not reuse erasure-coding fragments when a new profile is used.

- Changing the type of copies required for existing objects; for example, converting a large percentage of replicated objects to erasure-coded objects.
- Moving copies of existing objects to a completely different location; for example, moving a large number of objects to or from a Cloud Storage Pool or to or from a remote site.

Active ILM policy for example 6: Data protection at two sites

In this example, the active ILM policy was initially designed for a two-site StorageGRID system and uses two ILM rules.

Active policyPolicy history

Policy name:Data Protection for Two Sites (2 rules)

Reason for change :Data protection for two sites (using 2 rules)

Start date:2022-10-11 10:37:11 MDT

Simulate

Policy rulesRetention diagram

Rule order ?	Rule name	Filters ?
1	One-Site Erasure Coding for Tenant A	Tenant is Tenant A
Default	Two-Site Replication for Other Tenants	—

In this ILM policy, objects belonging to Tenant A are protected by 2+1 erasure coding at a single site, while objects belonging to all other tenants are protected across two sites using 2-copy replication.

Rule 1: One-site erasure coding for Tenant A

Rule definition	Example value
Rule name	One-Site Erasure Coding for Tenant A

Rule definition	Example value
Tenant Account	Tenant A
Storage Pool	Site 1
Placements	2+1 erasure coding in Site 1 from day 0 to forever

Rule 2: Two-site replication for other tenants

Rule definition	Example value
Rule name	Two-Site Replication for Other Tenants
Tenant Account	Ignore
Storage Pools	Site 1 and Site 2
Placements	Two replicated copies from day 0 to forever: one copy at Site 1 and one copy at Site 2.

ILM policy for example 6: Data protection at three sites

In this example, the ILM policy is being replaced with a new policy for a three-site StorageGRID system.

After performing an expansion to add the new site, the grid administrator created two new storage pools: a storage pool for Site 3 and a storage pool containing all three sites (not the same as the All Storage Nodes default storage pool). Then, the administrator created two new ILM rules and a new ILM policy, which is designed to protect data at all three sites.

When this new ILM policy is activated, objects belonging to Tenant A will be protected by 2+1 erasure coding at three sites, while objects belonging to other tenants (and smaller objects belonging to Tenant A) will be protected across three sites using 3-copy replication.

Rule 1: Three-site erasure coding for Tenant A

Rule definition	Example value
Rule name	Three-Site Erasure Coding for Tenant A
Tenant Account	Tenant A
Storage Pool	All 3 Sites (includes Site 1, Site 2, and Site 3)
Placements	2+1 erasure coding in All 3 Sites from day 0 to forever

Rule 2: Three-site replication for other tenants

Rule definition	Example value
Rule name	Three-Site Replication for Other Tenants
Tenant Account	Ignore
Storage Pools	Site 1, Site 2, and Site 3
Placements	Three replicated copies from day 0 to forever: one copy at Site 1, one copy at Site 2, and one copy at Site 3.

Activating the ILM policy for example 6

When you activate a new ILM policy, existing objects might be moved to new locations or new object copies might be created for existing objects, based on the placement instructions in any new or updated rules.



Errors in an ILM policy can cause unrecoverable data loss. Carefully review and simulate the policy before activating it to confirm that it will work as intended.



When you activate a new ILM policy, StorageGRID uses it to manage all objects, including existing objects and newly ingested objects. Before activating a new ILM policy, review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

What happens when erasure-coding instructions change

In the currently active ILM policy for this example, objects belonging to Tenant A are protected using 2+1 erasure coding at Site 1. In the new ILM policy, objects belonging to Tenant A will be protected using 2+1 erasure coding at Sites 1, 2, and 3.

When the new ILM policy is activated, the following ILM operations occur:

- New objects ingested by Tenant A are split into two data fragments and one parity fragment is added. Then, each of the three fragments is stored at a different site.
- The existing objects belonging to Tenant A are re-evaluated during the ongoing ILM scanning process. Because the ILM placement instructions use a new erasure-coding profile, entirely new erasure-coded fragments are created and distributed to the three sites.



The existing 2+1 fragments at Site 1 aren't reused. StorageGRID considers each erasure-coding profile to be unique and does not reuse erasure-coding fragments when a new profile is used.

What happens when replication instructions change

In the currently active ILM policy for this example, objects belonging to other tenants are protected using two replicated copies in storage pools at Sites 1 and 2. In the new ILM policy, objects belonging to other tenants will be protected using three replicated copies in storage pools at Sites 1, 2, and 3.

When the new ILM policy is activated, the following ILM operations occur:

- When any tenant other than Tenant A ingests a new object, StorageGRID creates three copies and saves one copy at each site.
- Existing objects belonging to these other tenants are re-evaluated during the ongoing ILM scanning process. Because the existing object copies at Site 1 and Site 2 continue to satisfy the replication requirements of the new ILM rule, StorageGRID only needs to create one new copy of the object for Site 3.

Performance impact of activating this policy

When the ILM policy in this example is activated, the overall performance of this StorageGRID system will be temporarily affected. Higher than normal levels of grid resources will be required to create new erasure-coded fragments for Tenant A's existing objects and new replicated copies at Site 3 for other tenants' existing objects.

As a result of the ILM policy change, client read and write requests might temporarily experience higher than normal latencies. Latencies will return to normal levels after the placement instructions are fully implemented across the grid.

To avoid resource issues when activating a new ILM policy, you can use the Ingest time advanced filter in any rule that might change the location of large numbers of existing objects. Set Ingest time to be greater than or equal to the approximate time when the new policy will go into effect to ensure that existing objects aren't moved unnecessarily.



Contact technical support if you need to slow or increase the rate at which objects are processed after an ILM policy change.

Example 7: Compliant ILM policy for S3 Object Lock

You can use the S3 bucket, ILM rules, and ILM policy in this example as a starting point when defining an ILM policy to meet the object protection and retention requirements for objects in buckets with S3 Object Lock enabled.



If you used the legacy Compliance feature in previous StorageGRID releases, you can also use this example to help manage any existing buckets that have the legacy Compliance feature enabled.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss.

Related information

- [Manage objects with S3 Object Lock](#)
- [Create an ILM policy](#)

Bucket and objects for S3 Object Lock example

In this example, an S3 tenant account named Bank of ABC has used the Tenant Manager to create a bucket with S3 Object Lock enabled to store critical bank records.

Bucket definition	Example value
Tenant Account Name	Bank of ABC
Bucket Name	bank-records
Bucket Region	us-east-1 (default)

Each object and object version that is added to the bank-records bucket will use the following values for `retain-until-date` and `legal hold` settings.

Setting for each object	Example value
<code>retain-until-date</code>	"2030-12-30T23:59:59Z" (December 30, 2030) Each object version has its own <code>retain-until-date</code> setting. This setting can be increased, but not decreased.
<code>legal hold</code>	"OFF" (Not in effect) A legal hold can be placed or lifted on any object version at any time during the retention period. If an object is under a legal hold, the object can't be deleted even if the <code>retain-until-date</code> has been reached.

ILM rule 1 for S3 Object Lock example: Erasure-coding profile with bucket matching

This example ILM rule applies only to the S3 tenant account named Bank of ABC. It matches any object in the `bank-records` bucket and then uses erasure coding to store the object on Storage Nodes at three data center sites using a 6+3 erasure-coding profile. This rule satisfies the requirements of buckets with S3 Object Lock enabled: a copy is kept on Storage Nodes from day 0 to forever, using Ingest time as the reference time.

Rule definition	Example value
Rule name	Compliant Rule: EC Objects in bank-records Bucket - Bank of ABC
Tenant Account	Bank of ABC
Bucket Name	<code>bank-records</code>
Advanced filter	Object Size (MB) greater than 1 Note: This filter ensures that erasure coding is not used for objects 1 MB or smaller.

Rule definition	Example value
Reference time	Ingest time

Rule definition	Example value
Placements	From day 0 store forever
Erasure-coding profile	• Create an erasure-coded copy on Storage Nodes at three data center sites • Uses 6+3 erasure-coding scheme

ILM rule 2 for S3 Object Lock example: Non-compliant rule

This example ILM rule initially stores two replicated object copies on Storage Nodes. After one year, it stores one copy on a Cloud Storage Pool forever. Because this rule uses a Cloud Storage Pool, it is not compliant and will not apply to the objects in buckets with S3 Object Lock enabled.

Rule definition	Example value
Rule name	Non-compliant rule: Use Cloud Storage Pool
Tenant accounts	Not specified
Bucket name	Not specified, but will only apply to buckets that don't have S3 Object Lock (or the legacy Compliance feature) enabled.
Advanced filter	Not specified

Rule definition	Example value
Reference time	Ingest time
Placements	• On Day 0, keep two replicated copies on Storage Nodes in Data Center 1 and Data Center 2 for 365 days • After 1 year, keep one replicated copy in a Cloud Storage Pool forever

ILM rule 3 for S3 Object Lock example: Default rule

This example ILM rule copies object data to storage pools in two data centers. This compliant rule is designed to be the default rule in the ILM policy. It does not include any filters, does not use the Noncurrent reference time, and satisfies the requirements of buckets with S3 Object Lock enabled: two object copies are kept on Storage Nodes from day 0 to forever, using Ingest as the reference time.

Rule definition	Example value
Rule name	Default compliant rule: Two Copies Two Data Centers
Tenant account	Not specified

Rule definition	Example value
Bucket name	Not specified
Advanced filter	Not specified

Rule definition	Example value
Reference time	Ingest time
Placements	From Day 0 to forever, keep two replicated copies—one on Storage Nodes in Data Center 1 and one on Storage Nodes in Data Center 2.

Compliant ILM policy for S3 Object Lock example

To create an ILM policy that will effectively protect all objects in your system, including those in buckets with S3 Object Lock enabled, you must select ILM rules that satisfy the storage requirements for all objects. Then, you must simulate and activate the policy.

Add rules to the policy

In this example, the ILM policy includes three ILM rules, in the following order:

1. A compliant rule that uses erasure coding to protect objects greater than 1 MB in a specific bucket with S3 Object Lock enabled. The objects are stored on Storage Nodes from day 0 to forever.
2. A non-compliant rule that creates two replicated object copies on Storage Nodes for a year and then moves one object copy to a Cloud Storage Pool forever. This rule does not apply to buckets with S3 Object Lock enabled because it uses a Cloud Storage Pool.
3. The default compliant rule that creates two replicated object copies on Storage Nodes from day 0 to forever.

Simulate the policy

After you have added rules to your policy, chosen a default compliant rule, and arranged the other rules, you should simulate the policy by testing objects from the bucket with S3 Object Lock enabled and from other buckets. For example, when you simulate the example policy, you would expect test objects to be evaluated as follows:

- The first rule will only match test objects that are greater than 1 MB in the bucket bank-records for the Bank of ABC tenant.
- The second rule will match all objects in all non-compliant buckets for all other tenant accounts.
- The default rule will match these objects:
 - Objects 1 MB or smaller in the bucket bank-records for the Bank of ABC tenant.
 - Objects in any other bucket that has S3 Object Lock enabled for all other tenant accounts.

Activate the policy

When you are completely satisfied that the new policy protects object data as expected, you can activate it.

Example 8: Priorities for S3 bucket lifecycle and ILM policy

Depending on your lifecycle configuration, objects follow the retention settings of either the S3 bucket lifecycle or an ILM policy.

Example of bucket lifecycle taking priority over ILM policy

ILM policy

- Rule based on noncurrent-time reference: On Day 0, keep X copies for 20 days
- Rule based on ingest-time reference (default): On Day 0, keep X copies for 50 days

Bucket Lifecycle

```
"Filter": {"Prefix": "docs/"}, "Expiration": {"Days": 100},  
"NoncurrentVersionExpiration": {"NoncurrentDays": 5}
```

Result

- An object named "docs/text" is ingested. It matches the bucket lifecycle filter of "docs/" prefix.
 - After 100 days a delete-marker is created and "docs/text" becomes noncurrent.
 - After 5 days, a total of 105 days since ingest, "docs/text" is deleted.
 - After 95 days, a total of 200 days since the ingest and 100 days since the delete-marker was created, the expired delete-marker is deleted.
- An object named "video/movie" is ingested. It does not match the filter and uses the ILM retention policy.
 - After 50 days a delete-marker is created and "video/movie" becomes noncurrent.
 - After 20 days, a total of 70 days since the ingest, "video/movie" is deleted.
 - After 30 days, a total of 100 days since the ingest and 50 days since the delete-marker was created, the expired delete-marker is deleted.

Example of bucket lifecycle implicitly keeping-forever

ILM policy

- Rule based on noncurrent-time reference: On Day 0, keep X copies for 20 days
- Rule based on ingest-time reference (default): On Day 0, keep X copies for 50 days

Bucket Lifecycle

```
"Filter": {"Prefix": "docs/"}, "Expiration": {"ExpiredObjectDeleteMarker":  
true}
```

Result

- An object named "docs/text" is ingested. It matches the bucket lifecycle filter of "docs/" prefix.

The `Expiration` action applies only to expired delete markers, which implies keeping everything else forever (starting with "docs/").

Delete markers that start with "docs/" are removed when they become expired.

- An object named "video/movie" is ingested. It does not match the filter and uses the ILM retention policy.

- After 50 days a delete-marker is created and "video/movie" becomes noncurrent.
- After 20 days, a total of 70 days since the ingest, "video/movie" is deleted.
- After 30 days, a total of 100 days since the ingest and 50 days since the delete-marker was created, the expired delete-marker is deleted.

Example of using bucket lifecycle to duplicate ILM and clean up expired delete markers

ILM policy

- Rule based on noncurrent-time reference: On Day 0, keep X copies for 20 days
- Rule based on ingest-time reference (default): On Day 0, keep X copies forever

Bucket Lifecycle

```
"Filter": {}, "Expiration": {"ExpiredObjectDeleteMarker": true},
"NoncurrentVersionExpiration": {"NoncurrentDays": 20}
```

Result

- The ILM policy is duplicated in the bucket lifecycle.
 - The ILM policy's forever rule is designed for removing objects manually and cleaning up noncurrent versions after 20 days. Consequently, the ingest-time rule will keep expired delete markers forever.
 - The bucket lifecycle duplicates the ILM policy's behavior while adding `"ExpiredObjectDeleteMarker": true`, which removes delete markers once they are expired
- An object is ingested. No filter means that the bucket lifecycle applies to all objects and overrides the ILM retention settings.
 - When a tenant issues an object delete request, a delete-marker is created and the object becomes noncurrent.
 - After 20 days, the noncurrent object is deleted and the delete-marker becomes expired.
 - Shortly after, the expired delete-marker is deleted.

System hardening

General considerations for system hardening

System hardening is the process of eliminating as many security risks as possible from a StorageGRID system.

As you install and configure StorageGRID, use these guidelines to help you meet any prescribed security objectives for confidentiality, integrity, and availability.

You should already be using industry-standard best practices for system hardening. For example, you use strong passwords for StorageGRID, use HTTPS instead of HTTP, and enable certificate-based authentication where available.

StorageGRID follows the [NetApp Vulnerability Handling Policy](#). Reported vulnerabilities are verified and addressed according to the product security incident response process.

When hardening a StorageGRID system, consider the following:

- **Which of the three StorageGRID networks** you have implemented. All StorageGRID systems must use

the Grid Network, but you might also be using the Admin Network, the Client Network, or both. Each network has different security considerations.

- **The type of platforms** you use for the individual nodes in your StorageGRID system. StorageGRID nodes can be deployed on VMware virtual machines, within a container engine on Linux hosts, or as dedicated hardware appliances. Each type of platform has its own set of hardening best practices.
- **How trusted the tenant accounts are.** If you are a service provider with untrusted tenant accounts, you will have different security concerns than if you only use trusted, in-house tenants.
- **Which security requirements and conventions** your organization follows. You might need to comply with specific regulatory or corporate requirements.

Hardening guidelines for software upgrades

You must keep your StorageGRID system and related services up to date to defend against attacks.

Upgrades to StorageGRID software

Whenever possible, you should upgrade StorageGRID software to the most recent major release or to the previous major release. Keeping StorageGRID up to date helps reduce the amount of time that known vulnerabilities are active and reduces the overall attack surface area. In addition, the most recent releases of StorageGRID often contain security hardening features that aren't included in earlier releases.

Consult the [NetApp Interoperability Matrix Tool](#) (IMT) to determine which version of StorageGRID software you should be using. When a hotfix is required, NetApp prioritizes creating updates for the most recent releases. Some patches might not be compatible with earlier releases.

- To download the most recent StorageGRID releases and hotfixes, go to [NetApp Downloads: StorageGRID](#).
- To upgrade StorageGRID software, see the [upgrade instructions](#).
- To apply a hotfix, see the [StorageGRID hotfix procedure](#).

Upgrades to external services

External services can have vulnerabilities that affect StorageGRID indirectly. You should ensure that the services that StorageGRID depends on are kept up to date. These services include LDAP, KMS (or KMIP server), DNS, and NTP.

For a list of supported versions, see the [NetApp Interoperability Matrix Tool](#).

Upgrades to hypervisors

If your StorageGRID nodes are running on VMware or another hypervisor, you must ensure that the hypervisor software and firmware are up to date.

For a list of supported versions, see the [NetApp Interoperability Matrix Tool](#).

Upgrades to Linux nodes

If your StorageGRID nodes are using Linux host platforms, you must ensure that security updates and kernel updates are applied to the host OS. Additionally, you must apply firmware updates to vulnerable hardware when these updates become available.

For a list of supported versions, see the [NetApp Interoperability Matrix Tool](#).

Hardening guidelines for StorageGRID networks

The StorageGRID system supports up to three network interfaces per grid node, allowing you to configure the networking for each individual grid node to match your security and access requirements.

For detailed information about StorageGRID networks, see the [StorageGRID network types](#).

Guidelines for Grid Network

You must configure a Grid Network for all internal StorageGRID traffic. All grid nodes are on the Grid Network, and they must be able to talk to all other nodes.

When configuring the Grid Network, follow these guidelines:

- Ensure that the network is secured from untrusted clients, such as those on the open internet.
- When possible, use the Grid Network exclusively for internal traffic. Both the Admin Network and the Client Network have additional firewall restrictions that block external traffic to internal services. Using the Grid Network for external client traffic is supported, but this use offers fewer layers of protection.
- If the StorageGRID deployment spans multiple data centers, use a virtual private network (VPN) or equivalent on the Grid Network to provide additional protection for internal traffic.
- Some maintenance procedures require secure shell (SSH) access on port 22 between the primary Admin Node and all other grid nodes. Use an external firewall to restrict SSH access to trusted clients.

Guidelines for Admin Network

The Admin Network is typically used for administrative tasks (trusted employees using the Grid Manager or SSH) and for communicating with other trusted services such as LDAP, DNS, NTP, or KMS (or KMIP server). However, StorageGRID does not enforce this usage internally.

If you are using the Admin Network, follow these guidelines:

- Block all internal traffic ports on the Admin Network. See the [list of internal ports](#).
- If untrusted clients can access the Admin Network, block access to StorageGRID on the Admin Network with an external firewall.

Guidelines for Client Network

The Client Network is typically used for tenants and for communicating with external services, such as the CloudMirror replication service or another platform service. However, StorageGRID does not enforce this usage internally.

If you are using the Client Network, follow these guidelines:

- Block all internal traffic ports on the Client Network. See the [list of internal ports](#).
- Accept inbound client traffic only on explicitly configured endpoints. See the information about [managing firewall controls](#).

Hardening guidelines for StorageGRID nodes

StorageGRID nodes can be deployed on VMware virtual machines, within a container

engine on Linux hosts, or as dedicated hardware appliances. Each type of platform and each type of node has its own set of hardening best practices.

Control remote IPMI access to BMC

You can enable or disable remote IPMI access for all appliances containing a BMC. The remote IPMI interface allows low-level hardware access to your StorageGRID appliances by anyone with a BMC account and password. If you do not need remote IPMI access to the BMC, disable this option.

- To control remote IPMI access to the BMC in Grid Manager, go to **CONFIGURATION > Security > Security settings > Appliances**:
 - Clear the **Enable remote IPMI access** checkbox to disable IPMI access to the BMC.
 - Select the **Enable remote IPMI access** checkbox to enable IPMI access to the BMC.

Firewall configuration

As part of the system hardening process, you must review external firewall configurations and modify them so that traffic is accepted only from the IP addresses and on the ports from which it is strictly needed.

StorageGRID includes an internal firewall on each node that enhances the security of your grid by enabling you to control network access to the node. You should [manage internal firewall controls](#) to prevent network access on all ports except those necessary for your specific grid deployment. The configuration changes you make on the Firewall control page are deployed to each node.

Specifically, you can manage these areas:

- **Privileged addresses**: You can allow selected IP addresses or subnets to access ports that are closed by settings on the Manage external access tab.
- **Manage external access**: You can close ports that are open by default, or reopen ports previously closed.
- **Untrusted Client Network**: You can specify whether a node trusts inbound traffic from the Client Network as well as the additional ports you want open when untrusted Client Network is configured.

While this internal firewall provides an additional layer of protection against some common threats, it does not remove the need for an external firewall.

For a list of all internal and external ports used by StorageGRID, see [Network port reference](#).

Disable unused services

For all StorageGRID nodes, you should disable or block access to unused services. For example, if you aren't planning to use DHCP, use the Grid Manager to close port 68. Select **CONFIGURATION > Firewall control > Manage external access**. Then change the Status toggle for port 68 from **Open** to **Closed**.

Virtualization, containers, and shared hardware

For all StorageGRID nodes, avoid running StorageGRID on the same physical hardware as untrusted software. Don't assume that hypervisor protections will prevent malware from accessing StorageGRID-protected data if both StorageGRID and the malware exist on the same the physical hardware. For example, the Meltdown and Spectre attacks exploit critical vulnerabilities in modern processors and allow programs to steal data in memory on the same computer.

Protect nodes during installation

Don't allow untrusted users to access StorageGRID nodes over the network when the nodes are being installed. Nodes aren't fully secure until they have joined the grid.

Guidelines for Admin Nodes

Admin Nodes provide management services such as system configuration, monitoring, and logging. When you sign in to the Grid Manager or the Tenant Manager, you are connecting to an Admin Node.

Follow these guidelines to secure the Admin Nodes in your StorageGRID system:

- Secure all Admin Nodes from untrusted clients, such as those on the open internet. Ensure that no untrusted client can access any Admin Node on the Grid Network, the Admin Network, or the Client Network.
- StorageGRID Groups control access to Grid Manager and Tenant Manager features. Grant each Group of users the minimum required permissions for their role, and use the read-only access mode to prevent users from changing configuration.
- When using StorageGRID load balancer endpoints, use Gateway Nodes instead of Admin Nodes for untrusted client traffic.
- If you have untrusted tenants, don't allow them to have direct access to the Tenant Manager or the Tenant Management API. Instead, have any untrusted tenants use a tenant portal or an external tenant management system, which interacts with the Tenant Management API.
- Optionally, use an admin proxy for more control over AutoSupport communication from Admin Nodes to NetApp Support. See the steps for [creating an admin proxy](#).
- Optionally, use the restricted 8443 and 9443 ports to separate Grid Manager and Tenant Manager communications. Block the shared port 443 and limit tenant requests to port 9443 for additional protection.
- Optionally, use separate Admin Nodes for grid administrators and tenant users.

For more information, see the instructions for [administering StorageGRID](#).

Guidelines for Storage Nodes

Storage Nodes manage and store object data and metadata. Follow these guidelines to secure the Storage Nodes in your StorageGRID system.

- Don't allow untrusted clients to connect directly to Storage Nodes. Use a load balancer endpoint served by a Gateway Node or a third party load balancer.
- Don't enable outbound services for untrusted tenants. For example, when creating the account for an untrusted tenant, don't allow the tenant to use its own identity source and don't allow the use of platform services. See the steps for [creating a tenant account](#).
- Use a third-party load balancer for untrusted client traffic. Third-party load balancing offers more control and additional layers of protection against attack.
- Optionally, use a storage proxy for more control over Cloud Storage Pools and platform services communication from Storage Nodes to external services. See the steps for [creating a storage proxy](#).
- Optionally, connect to external services using the Client Network. Then, select **CONFIGURATION > Security > Firewall control > Untrusted Client Networks** and indicate that the Client Network on the Storage Node is untrusted. The Storage Node no longer accepts any incoming traffic on the Client Network, but it continues to allow outbound requests for Platform Services.

Guidelines for Gateway Nodes

Gateway Nodes provide an optional load-balancing interface that client applications can use to connect to StorageGRID. Follow these guidelines to secure any Gateway Nodes in your StorageGRID system:

- Configure and use load balancer endpoints. See [Considerations for load balancing](#).
- Use a third-party load balancer between the client and the Gateway Node or Storage Nodes for untrusted client traffic. Third-party load balancing offers more control and additional layers of protection against attack. If you do use a third-party load balancer, network traffic can still optionally be configured to go through an internal load balancer endpoint or be sent directly to Storage Nodes.
- If you are using load balancer endpoints, optionally have clients connect over the Client Network. Then, select **CONFIGURATION > Security > Firewall control > Untrusted Client Networks** and indicate that the Client Network on the Gateway Node is untrusted. The Gateway Node only accepts inbound traffic on the ports explicitly configured as load balancer endpoints.

Guidelines for hardware appliance nodes

StorageGRID hardware appliances are specially designed for use in a StorageGRID system. Some appliances can be used as Storage Nodes. Other appliances can be used as Admin Nodes or Gateway Nodes. You can combine appliance nodes with software-based nodes or deploy fully engineered, all-appliance grids.

Follow these guidelines to secure any hardware appliance nodes in your StorageGRID system:

- If the appliance uses SANtricity System Manager for storage controller management, prevent untrusted clients from accessing SANtricity System Manager over the network.
- If the appliance has a baseboard management controller (BMC), be aware that the BMC management port allows low-level hardware access. Connect the BMC management port only to a secure, trusted, internal management network. If no such network is available, leave the BMC management port unconnected or blocked, unless a BMC connection is requested by technical support.
- If the appliance supports remote management of the controller hardware over Ethernet using the Intelligent Platform Management Interface (IPMI) standard, block untrusted traffic on port 623.



You can enable or disable remote IPMI access for all appliances containing a BMC. The remote IPMI interface allows low-level hardware access to your StorageGRID appliances by anyone with a BMC account and password. If you do not need remote IPMI access to the BMC, disable this option using one of the following methods:

In Grid Manager, go to **CONFIGURATION > Security > Security settings > Appliances** and clear the **Enable remote IPMI access** checkbox.

In the Grid management API, use the private endpoint: `PUT /private/bmc`.

- For appliance models containing SED, FDE, or FIPS NL-SAS drives that you manage with SANtricity System Manager, [enable and configure SANtricity Drive Security](#).
- For appliance models containing SED or FIPS NVMe SSDs that you manage using the StorageGRID Appliance Installer and Grid Manager, [enable and configure StorageGRID drive encryption](#).
- For appliances without SED, FDE, or FIPS drives, enable and configure StorageGRID software node encryption [using a Key Management Server \(KMS\)](#).

Hardening guidelines for TLS and SSH

You should replace the default certificates created during installation and select the appropriate security policy for TLS and SSH connections.

Hardening guidelines for certificates

You should replace the default certificates created during installation with your own custom certificates.

For many organizations, the self-signed digital certificate for StorageGRID web access is not compliant with their information security policies. On production systems, you should install a CA-signed digital certificate for use in authenticating StorageGRID.

Specifically, you should use custom server certificates instead of these default certificates:

- **Management interface certificate:** Used to secure access to the Grid Manager, the Tenant Manager, the Grid Management API, and the Tenant Management API.
- **S3 API certificate:** Used to secure access to Storage Nodes and Gateway Nodes, which S3 client applications use to upload and download object data.

See [Manage security certificates](#) for details and instructions.



StorageGRID manages the certificates used for load balancer endpoints separately. To configure load balancer certificates, see [Configure load balancer endpoints](#).

When using custom server certificates, follow these guidelines:

- Certificates should have a `subjectAltName` that matches DNS entries for StorageGRID. For details, see section 4.2.1.6, "Subject Alternative Name," in [RFC 5280: PKIX Certificate and CRL Profile](#).
- When possible, avoid the use of wildcard certificates. An exception to this guideline is the certificate for an S3 virtual hosted style endpoint, which requires the use of a wildcard if bucket names aren't known in advance.
- When you must use wildcards in certificates, you should take additional steps to reduce the risks. Use a wildcard pattern such as `*.s3.example.com`, and don't use the `s3.example.com` suffix for other applications. This pattern also works with path-style S3 access, such as `dc1-s1.s3.example.com/mybucket`.
- Set the certificate expiration times to be short (for example, 2 months), and use the Grid Management API to automate certificate rotation. This is especially important for wildcard certificates.

In addition, clients should use strict hostname checking when communicating with StorageGRID.

Hardening guidelines for TLS and SSH policy

You can select a security policy to determine which protocols and ciphers are used to establish secure TLS connections with client applications and secure SSH connections to internal StorageGRID services.

The security policy controls how TLS and SSH encrypt data in motion. As a best practice, you should disable encryption options that aren't required for application compatibility. Use the default Modern policy, unless your system needs to be Common Criteria-compliant or you need to use other ciphers.

See [Manage the TLS and SSH policy](#) for details and instructions.

Other hardening guidelines

In addition to following the hardening guidelines for StorageGRID networks and nodes, you should follow the hardening guidelines for other areas of the StorageGRID system.

Temporary installation password

To secure the StorageGRID system during installation, set a password on the temporary installer password page in the StorageGRID installation UI or in the Installation API. When set, this password applies to all methods for installing StorageGRID, including the user interface, Installation API, and `configure-storagegrid.py` script.

For more information, refer to:

- [Install StorageGRID on Red Hat Enterprise Linux](#)
- [Install StorageGRID on Ubuntu or Debian](#)
- [Install StorageGRID on VMware](#)
- [Install StorageGRID appliance](#)

Logs and audit messages

Always protect StorageGRID logs and audit message output in a secure manner. StorageGRID logs and audit messages provide invaluable information from a support and system availability standpoint. In addition, the information and details contained in StorageGRID logs and audit message output are generally of a sensitive nature.

Configure StorageGRID to send security events to an external syslog server. If using syslog export, select TLS and RELP/TLS for the transport protocols.

See the [Log files reference](#) for more information about StorageGRID logs. See [Audit messages](#) for more information about StorageGRID audit messages.

NetApp AutoSupport

The AutoSupport feature of StorageGRID allows you to proactively monitor the health of your system and automatically send packages to the NetApp Support Site, your organization's internal support team, or a support partner. By default, sending AutoSupport packages to NetApp is enabled when StorageGRID is configured for the first time.

The AutoSupport feature can be disabled. However, NetApp recommends enabling it because AutoSupport helps speed problem identification and resolution should an issue arise on your StorageGRID system.

AutoSupport supports HTTPS, HTTP, and SMTP for transport protocols. Because of the sensitive nature of AutoSupport packages, NetApp strongly recommends using HTTPS as the default transport protocol for sending AutoSupport packages to NetApp.

Cross-origin resource sharing (CORS)

You can configure cross-origin resource sharing (CORS) for an S3 bucket if you want that bucket and objects in that bucket to be accessible to web applications in other domains. In general, don't enable CORS unless it is required. If CORS is required, restrict it to trusted origins.

See the steps for [configuring cross-origin resource sharing \(CORS\)](#).

External security devices

A complete hardening solution must address security mechanisms outside of StorageGRID. Using additional infrastructure devices for filtering and limiting access to StorageGRID is an effective way to establish and maintain a stringent security posture. These external security devices include firewalls, intrusion prevention

systems (IPSs), and other security devices.

A third-party load balancer is recommended for untrusted client traffic. Third-party load balancing offers more control and additional layers of protection against attack.

Ransomware mitigation

Help protect your object data from ransomware attacks by following the recommendations in [Ransomware defense with StorageGRID](#).

Configure StorageGRID for FabricPool

Configure StorageGRID for FabricPool

If you use NetApp ONTAP software, you can use NetApp FabricPool to tier inactive data to a NetApp StorageGRID object storage system.

Use these instructions to:

- Learn the considerations and best practices for configuring StorageGRID for a FabricPool workload.
- Learn how to configure a StorageGRID object storage system for use with FabricPool.
- Learn how to provide the required values to ONTAP when attaching StorageGRID as a FabricPool cloud tier.

Quick start for configuring StorageGRID for FabricPool

1

Plan your configuration

- Decide which FabricPool volume tiering policy you will use to tier inactive ONTAP data to StorageGRID.
- Plan and install a StorageGRID system to meet your storage capacity and performance needs.
- Become familiar with StorageGRID system software, including the [Grid Manager](#) and the [Tenant Manager](#).
- Review the FabricPool best practices for [HA groups](#), [load balancing](#), [ILM](#), and [more](#).
- Review these additional resources, which provide details about using and configuring ONTAP and FabricPool:

[TR-4598: FabricPool Best Practices in ONTAP](#)

[ONTAP documentation for FabricPool](#)

2

Perform prerequisite tasks

Obtain the [information needed to attach StorageGRID as a cloud tier](#), including:

- IP addresses
- Domain names
- SSL certificate

Optionally, configure [identity federation](#) and [single sign-on](#).

3

Configure StorageGRID settings

Use StorageGRID to obtain the values ONTAP needs to connect to the grid.

Using the [FabricPool setup wizard](#) is the recommended and the fastest way to configure all items, but you can also configure each entity manually, if required.

4

Configure ONTAP and DNS

Use ONTAP to [add a cloud tier](#) that uses the StorageGRID values. Then, [configure DNS entries](#) to associate IP addresses to any domain names you plan to use.

5

Monitor and manage

When your system is up and running, perform ongoing tasks in ONTAP and StorageGRID to manage and monitor FabricPool data tiering over time.

What is FabricPool?

FabricPool is an ONTAP hybrid storage solution that uses a high-performance flash aggregate as the performance tier and an object store as the cloud tier. Using FabricPool-enabled aggregates helps you reduce storage cost without compromising performance, efficiency, or protection.

FabricPool associates a cloud tier (an external object store, such as StorageGRID) with a local tier (an ONTAP storage aggregate) to create a composite collection of discs. Volumes inside the FabricPool can then take advantage of the tiering by keeping active (hot) data on high-performance storage (the local tier) and tiering inactivate (cold) data to the external object store (the cloud tier).

No architectural changes are required, and you can continue managing your data and application environment from the central ONTAP storage system.

What is StorageGRID?

NetApp StorageGRID is a storage architecture that manages data as objects, as opposed to other storage architectures such as file or block storage. Objects are kept inside a single container (such as a bucket) and aren't nested as files inside a directory inside other directories. Although object storage generally provides lower performance than file or block storage, it is significantly more scalable. StorageGRID buckets can hold petabytes of data and billions of objects.

Why use StorageGRID as a FabricPool cloud tier?

FabricPool can tier ONTAP data to a number of object storage providers, including StorageGRID. Unlike public clouds that might set a maximum number of supported input/output operations per second (IOPS) at the bucket or container level, StorageGRID performance scales with the number of nodes in a system. Using StorageGRID as a FabricPool cloud tier allows you to keep your cold data in your own private cloud for highest performance and complete control over your data.

In addition, a FabricPool license is not required when you use StorageGRID as the cloud tier.

Information needed to attach StorageGRID as a cloud tier

Before you can attach StorageGRID as a cloud tier for FabricPool, you must perform configuration steps in StorageGRID and obtain certain values for use in ONTAP.

What values do I need?

The following table shows the values you must configure in StorageGRID and how those values are used by ONTAP and the DNS server.

Value	Where value is configured	Where value is used
Virtual IP (VIP) addresses	StorageGRID > HA group	DNS entry
Port	StorageGRID > Load balancer endpoint	ONTAP System Manager > Add Cloud Tier
SSL certificate	StorageGRID > Load balancer endpoint	ONTAP System Manager > Add Cloud Tier
Server name (FQDN)	StorageGRID > Load balancer endpoint	DNS entry
Access key ID and secret access key	StorageGRID > Tenant and bucket	ONTAP System Manager > Add Cloud Tier
Bucket/Container name	StorageGRID > Tenant and bucket	ONTAP System Manager > Add Cloud Tier

How do I get these values?

Depending on your requirements, you can do either of the following to obtain the information you need:

- Use the [FabricPool setup wizard](#). The FabricPool setup wizard helps you to quickly configure the required values in StorageGRID and outputs a file that you can use to configure ONTAP System Manager. The wizard guides you through the required steps and helps to make sure your settings conform to StorageGRID and FabricPool best practices.
- Configure each item manually. Then, enter the values into ONTAP System Manager or the ONTAP CLI. Follow these steps:
 1. [Configure a high availability \(HA\) group for FabricPool](#).
 2. [Create a load balancer endpoint for FabricPool](#).
 3. [Create a tenant account for FabricPool](#).
 4. Sign in to the tenant account, and [create the bucket and access keys for the root user](#).
 5. Create an ILM rule for FabricPool data and add it to your active ILM policies. See [Configure ILM for FabricPool data](#).
 6. Optionally, [create a traffic classification policy for FabricPool](#).

Use FabricPool setup wizard

Use FabricPool setup wizard: Considerations and requirements

You can use the FabricPool setup wizard to configure StorageGRID as the object storage system for a FabricPool cloud tier. After you complete the setup wizard, you can enter the required details into ONTAP System Manager.

When to use the FabricPool setup wizard

The FabricPool setup wizard guides you through each step of configuring StorageGRID for use with FabricPool and automatically configures certain entities for you, such as the ILM and traffic classification policies. As part of completing the wizard, you download a file that you can use to enter values into ONTAP System Manager. Use the wizard to configure your system more quickly and to make sure your settings conform to StorageGRID and FabricPool best practices.

Assuming you have Root access permission, you can complete the FabricPool setup wizard when you start using the StorageGRID Grid Manager, or you can access and complete the wizard at any later time. Depending on your requirements, you can also configure some or all of the required items manually and then use the wizard to assemble the values that ONTAP needs into a single file.



Use the FabricPool setup wizard unless you know you have special requirements or your implementation will require significant customization.

Before using the wizard

Confirm you have completed these prerequisite steps.

Review best practices

- You have a general understanding of the [information needed to attach StorageGRID as a cloud tier](#).
- You have reviewed the FabricPool best practices for:
 - [High availability \(HA\) groups](#)
 - [Load balancing](#)
 - [ILM rules and policy](#)

Obtain IP addresses and set up VLAN interfaces

If you will configure an HA group, you know which nodes ONTAP will connect to and which StorageGRID network will be used. You also know which values to enter for the subnet CIDR, gateway IP address, and virtual IP (VIP) addresses.

If you plan to use a virtual LAN to segregate FabricPool traffic, you have already configured the VLAN interface. See [Configure VLAN interfaces](#).

Configure identity federation and SSO

If you plan to use identity federation or single sign-on (SSO) for your StorageGRID system, you have enabled these features. You also know which federated group should have root access for the tenant account that ONTAP will use. See [Use identity federation](#) and [Configure single sign-on](#).

Obtain and configure domain names

- You know which fully qualified domain name (FQDN) to use for StorageGRID. Domain name server (DNS) entries will map this FQDN to the virtual IP (VIP) addresses of the HA group that you create using the wizard. See [Configure DNS server](#).
- If you plan to use S3 virtual hosted-style requests, you have [configured S3 endpoint domain names](#). ONTAP uses path-style URLs by default, but using virtual hosted-style requests is recommended.

Review load balancer and security certificate requirements

If you plan to use the StorageGRID load balancer, you have reviewed the general [considerations for load balancing](#). You have the certificates you will upload or the values you need to generate a certificate.

If you plan to use an external (third-party) load balancer endpoint, you have the fully qualified domain name (FQDN), port, and certificate for that load balancer.

Confirm ILM storage pool configuration

if you initially installed StorageGRID 11.6 or earlier, you have configured the storage pool you will use. In general, you should create a storage pool for each StorageGRID site you will use to store ONTAP data.



This prerequisite does not apply if you initially installed StorageGRID 11.7 or 11.8. When you initially install either of these versions, storage pools are automatically created for each site.

Relationship between ONTAP and the StorageGRID cloud tier

The FabricPool wizard guides you through the process of creating a single StorageGRID cloud tier that includes one StorageGRID tenant, one set of access keys, and one StorageGRID bucket. You can attach this StorageGRID cloud tier to one or more ONTAP local tiers.

Attaching a single cloud tier to multiple local tiers in a cluster is the general best practice. However, depending on your requirements, you might want to use more than one bucket or even more than one StorageGRID tenant for the local tiers in a single cluster. Using different buckets and tenants allows you to isolate data and data access between ONTAP local tiers, but is somewhat more complex to configure and manage.

NetApp does not recommend attaching a single cloud tier to local tiers in multiple clusters.



For the best practices for using StorageGRID with NetApp MetroCluster™ and FabricPool Mirror, see [TR-4598: FabricPool Best Practices in ONTAP](#).

Optional: Use a different bucket for each local tier

To use more than one bucket for the local tiers in an ONTAP cluster, add more than one StorageGRID cloud tier in ONTAP. Each cloud tier shares the same HA group, load balancer endpoint, tenant, and access keys, but uses a different container (StorageGRID bucket). Follow these general steps:

1. From StorageGRID Grid Manager, complete the FabricPool setup wizard for the first cloud tier.
2. From ONTAP System Manager, add a cloud tier and use the file you downloaded from StorageGRID to provide the required values.
3. From StorageGRID Tenant Manager, sign in to the tenant that was created by the wizard, and create a second bucket.
4. Complete the FabricPool wizard again. Select the existing HA group, load balancer endpoint, and tenant.

Then, select the new bucket you created manually. Create a new ILM rule for the new bucket and activate an ILM policy to include that rule.

5. From ONTAP, add a second cloud tier but provide the new bucket name.

Optional: Use a different tenant and bucket for each local tier

To use more than one tenant and different sets of access keys for the local tiers in an ONTAP cluster, add more than one StorageGRID cloud tier in ONTAP. Each cloud tier shares the same HA group, load balancer endpoint, but uses a different tenant, access keys, and container (StorageGRID bucket). Follow these general steps:

1. From StorageGRID Grid Manager, complete the FabricPool setup wizard for the first cloud tier.
2. From ONTAP System Manager, add a cloud tier and use the file you downloaded from StorageGRID to provide the required values.
3. Complete the FabricPool wizard again. Select the existing HA group and load balancer endpoint. Create a new tenant and bucket. Create a new ILM rule for the new bucket and activate an ILM policy to include that rule.
4. From ONTAP, add a second cloud tier but provide the new access key, secret key, and bucket name.

Access and complete the FabricPool setup wizard

You can use the FabricPool setup wizard to configure StorageGRID as the object storage system for a FabricPool cloud tier.

Before you begin

- You have reviewed the [considerations and requirements](#) for using the FabricPool setup wizard.



If you want to configure StorageGRID for use with any other S3 client application, go to [Use S3 setup wizard](#).


- You have the [Root access permission](#).

Access the wizard

You can complete the FabricPool setup wizard when you start using the StorageGRID Grid Manager, or you can access and complete the wizard at any later time.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. If the **FabricPool and S3 setup wizard** banner appears on the dashboard, select the link in the banner. If the banner no longer appears, select the help icon from the header bar in the Grid Manager and select **FabricPool and S3 setup wizard**.



3. In the FabricPool section of the FabricPool and S3 setup wizard page, select **Configure now**.

Step 1 of 9: Configure HA group appears.

Step 1 of 9: Configure HA group

A high availability (HA) group is a collection of nodes that each contain the StorageGRID Load Balancer service. An HA group can contain Gateway Nodes, Admin Nodes, or both.

You can use an HA group to help keep FabricPool data connections available. An HA group uses virtual IP addresses (VIPs) to provide highly available access to the Load Balancer service. If the active interface in the HA group fails, a backup interface can manage the workload with little impact to FabricPool operations.

For details about this task, see [Manage high availability groups](#) and [Best practices for high availability groups](#).

Steps

1. If you plan to use an external load balancer, you don't need to create an HA group. Select **Skip this step** and go to [Step 2 of 9: Configure load balancer endpoint](#).
2. To use the StorageGRID load balancer, create a new HA group or use an existing HA group.

Create HA group

- a. To create a new HA group, select **Create HA group**.
- b. For the **Enter details** step, complete the following fields.

Field	Description
HA group name	A unique display name for this HA group.
Description (optional)	The description of this HA group.

- c. For the **Add interfaces** step, select the node interfaces you want to use in this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

You can select one or more nodes, but you can select only one interface for each node.

- d. For the **Prioritize interfaces** step, determine the Primary interface and any backup interfaces for this HA group.

Drag rows to change the values in the **Priority order** column.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

If the HA group includes more than one interface and the active interface fails, the virtual IP (VIP) addresses move to the first backup interface in the priority order. If that interface fails, the VIP addresses move to the next backup interface, and so on. When failures are resolved, the VIP addresses move back to highest priority interface available.

- e. For the **Enter IP addresses** step, complete the following fields.

Field	Description
Subnet CIDR	The address of the VIP subnet in CIDR notation—an IPv4 address followed by a slash and the subnet length (0-32). The network address must not have any host bits set. For example, 192.16.0.0/22.
Gateway IP address (optional)	Optional. If the ONTAP IP addresses used to access StorageGRID aren't on the same subnet as the StorageGRID VIP addresses, enter the StorageGRID VIP local gateway IP address. The local gateway IP address must be within the VIP subnet.
Virtual IP address	Enter at least one and no more than ten VIP addresses for the active interface in the HA group. All VIP addresses must be within the VIP subnet and all will be active at the same time on the active interface. At least one address must be IPv4. Optionally, you can specify additional IPv4 and IPv6 addresses.

- f. Select **Create HA group** and then select **Finish** to return to the FabricPool setup wizard.
- g. Select **Continue** to go to the load balancer step.

Use existing HA group

- a. To use an existing HA group, select the HA group name from the **Select an HA group** drop-down list.
- b. Select **Continue** to go to the load balancer step.

Step 2 of 9: Configure load balancer endpoint

StorageGRID uses a load balancer to manage the workload from client applications, such as FabricPool. Load balancing maximizes speed and connection capacity across multiple Storage Nodes.

You can use the StorageGRID Load Balancer service, which exists on all Gateway and Admin Nodes, or you can connect to an external (third-party) load balancer. Using the StorageGRID load balancer is recommended.

For details about this task, see the general [considerations for load balancing](#) and the [best practices for load balancing for FabricPool](#).

Steps

1. Select or create a StorageGRID load balancer endpoint or use an external load balancer.

Create endpoint

- a. Select **Create endpoint**.
- b. For the **Enter endpoint details** step, complete the following fields.

Field	Description
Name	A descriptive name for the endpoint.
Port	The StorageGRID port you want to use for load balancing. This field defaults to 10433 for the first endpoint you create, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes, because these ports are reserved on Admin Nodes. Note: Ports used by other grid services aren't permitted. See the Network port reference.
Client type	Must be S3 .
Network protocol	Select HTTPS. Note: Communicating with StorageGRID without TLS encryption is supported but not recommended.

- c. For the **Select binding mode** step, specify the binding mode. The binding mode controls how the endpoint is accessed using any IP address or using specific IP addresses and network interfaces.

Mode	Description
Global (default)	Clients can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN. Use the Global setting (default) unless you need to restrict the accessibility of this endpoint.
Virtual IPs of HA groups	Clients must use a virtual IP address (or corresponding FQDN) of an HA group to access this endpoint. Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints don't overlap.
Node interfaces	Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces to access this endpoint.
Node type	Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node to access this endpoint.

d. For the **Tenant access** step, select one of the following:

Field	Description
Allow all tenants (default)	All tenant accounts can use this endpoint to access their buckets. Allow all tenants is almost always the appropriate option for the load balancer endpoint used for FabricPool. You must select this option if you are using the FabricPool setup wizard for a new StorageGRID system and you have not yet created any tenant accounts.
Allow selected tenants	Only the selected tenant accounts can use this endpoint to access their buckets.
Block selected tenants	The selected tenant accounts can't use this endpoint to access their buckets. All other tenants can use this endpoint.

e. For the **Attach certificate** step, select one of the following:

Field	Description
Upload certificate (recommended)	Use this option to upload a CA-signed server certificate, certificate private key, and optional CA bundle.
Generate certificate	Use this option to generate a self-signed certificate. See Configure load balancer endpoints for details of what to enter.
Use StorageGRID S3 certificate	This option is available only if you have already uploaded or generated a custom version of the StorageGRID global certificate. See Configure S3 API certificates for details.

f. Select **Finish** to return to the FabricPool setup wizard.

g. Select **Continue** to go to the tenant and bucket step.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

Use existing load balancer endpoint

- Select the name of an existing endpoint from the **Select a load balancer endpoint** drop-down list.
- Select **Continue** to go to the tenant and bucket step.

Use external load balancer

- Complete the following fields for the external load balancer.

Field	Description
FQDN	The fully qualified domain name (FQDN) of the external load balancer.
Port	The port number that FabricPool will use to connect to the external load balancer.
Certificate	Copy the server certificate for the external load balancer and paste it into this field.

- b. Select **Continue** to go to the tenant and bucket step.

Step 3 of 9: Tenant and bucket

A tenant is an entity that can use S3 applications to store and retrieve objects in StorageGRID. Each tenant has its own users, access keys, buckets, objects, and a specific set of capabilities. You must create a StorageGRID tenant before you can create the bucket that FabricPool will use.

A bucket is a container used to store a tenant's objects and object metadata. Although some tenants might have many buckets, the wizard lets you create or select only one tenant and one bucket at a time. You can use the Tenant Manager later to add any additional buckets you need.

You can create a new tenant and bucket for FabricPool use, or you can select an existing tenant and bucket. If you create a new tenant, the system automatically creates the access key ID and secret access key for the tenant's root user.

For details about this task, see [Create a tenant account for FabricPool](#) and [Create an S3 bucket and obtain an access key](#).

Steps

Create a new tenant and bucket or select an existing tenant.

New tenant and bucket

1. To create a new tenant and bucket, enter a **Tenant name**. For example, `FabricPool tenant`.
2. Define root access for the tenant account, based on whether your StorageGRID system uses [identity federation](#), [single sign-on \(SSO\)](#), or both.

Option	Do this
If identity federation is not enabled	Specify the password to use when signing into the tenant as the local root user.
If identity federation is enabled	1. Select an existing federated group to have Root access permission for the tenant. 2. Optionally, specify the password to use when signing in to the tenant as the local root user.
If both identity federation and single sign-on (SSO) are enabled	Select an existing federated group to have Root access permission for the tenant. No local users can sign in.

3. For **Bucket name**, enter the name of the bucket FabricPool will use to store ONTAP data. For example, `fabricpool-bucket`.



You can't change the bucket name after creating the bucket.

4. Select the **Region** for this bucket.

Use the default region (`us-east-1`) unless you expect to use ILM in the future to filter objects based on the bucket's region.

5. Select **Create and Continue** to create the tenant and bucket and to go to the download data step

Select tenant and bucket

The existing tenant account must have at least one bucket that does not have versioning enabled. You can't select an existing tenant account if no bucket exists for that tenant.

1. Select the existing tenant from the **Tenant name** drop-down list.
2. Select the existing bucket from the **Bucket name** drop-down list.

FabricPool does not support object versioning, so buckets that have versioning enabled aren't shown.



Don't select a bucket that has S3 Object Lock enabled for use with FabricPool.

3. Select **Continue** to go to the download data step.

Step 4 of 9: Download ONTAP settings

During this step, you download a file that you can use to enter values into ONTAP System Manager.

Steps

1. Optionally, select the copy icon () to copy both the access key ID and secret access key to the clipboard.

These values are included in the download file, but you might want to save them separately.

2. Select **Download ONTAP settings** to download a text file that contains the values you've entered so far.

The `ONTAP_FabricPool_settings_bucketname.txt` file includes the information you need to configure StorageGRID as the object storage system for a FabricPool cloud tier, including:

- Load balancer connection details, including the server name (FQDN), port, and certificate
- Bucket name
- Access key ID and secret access key for the root user of the tenant account

3. Save the copied keys and downloaded file to a secure location.



Don't close this page until you have copied both access keys, downloaded the ONTAP settings, or both. The keys will not be available after you close this page. Make sure to save this information in a secure location because it can be used to obtain data from your StorageGRID system.

4. Select the checkbox to confirm you have downloaded or copied the access key ID and secret access key.
5. Select **Continue** to go to the ILM storage pool step.

Step 5 of 9: Select a storage pool

A storage pool is a group of Storage Nodes. When you select a storage pool, you determine which nodes StorageGRID will use to store the data tiered from ONTAP.

For details about this step, see [Create a storage pool](#).

Steps

1. From the **Site** drop-down list, select the StorageGRID site you want to use for the data tiered from ONTAP.
2. From the **Storage pool** drop-down list, select the storage pool for that site.

The storage pool for a site includes all Storage Nodes at that site.

3. Select **Continue** to go to the ILM rule step.

Step 6 of 9: Review ILM rule for FabricPool

Information lifecycle management (ILM) rules control the placement, duration, and ingest behavior for all objects in your StorageGRID system.

The FabricPool setup wizard automatically creates the recommended ILM rule for FabricPool use. This rule applies only to the bucket you specified. It uses 2+1 erasure coding at a single site to store the data that is tiered from ONTAP.

For details about this step, see [Create ILM rule](#) and [Best practices for using ILM with FabricPool data](#).

Steps

1. Review the rule details.

Field	Description
Rule name	Automatically generated and can't be changed
Description	Automatically generated and can't be changed
Filter	The bucket name This rule only applies to objects that are saved in the bucket you specified.
Reference time	Ingest time The placement instruction starts when objects are initially saved to the bucket.
Placement instruction	Use 2+1 erasure coding

2. Sort the retention diagram by **Time period** and **Storage pool** to confirm the placement instruction.

- The **Time period** for the rule is **Day 0 - forever**. **Day 0** means that the rule is applied when data is tiered from ONTAP. **Forever** means that StorageGRID ILM will not delete data that has been tiered from ONTAP.
- The **Storage pool** for the rule is the storage pool you selected. **EC 2+1** means the data will be stored using 2+1 erasure coding. Each object will be saved as two data fragments and one parity fragment. The three fragments for each object will be saved to different Storage Nodes at a single site.

3. Select **Create and Continue** to create this rule and to go to the ILM policy step.

Step 7 of 9: Review and activate ILM policy

After the FabricPool setup wizard creates the ILM rule for FabricPool use, it creates an ILM policy. You must carefully simulate and review this policy before activating it.

For details about this step, see [Create ILM policy](#) and [Best practices for using ILM with FabricPool data](#).



When you activate a new ILM policy, StorageGRID uses that policy to manage the placement, duration, and data protection of all objects in the grid, including existing objects and newly ingested objects. In some cases, activating a new policy can cause existing objects to be moved to new locations.



To avoid data loss, do not use an ILM rule that will expire or delete FabricPool cloud tier data. Set the retention period to **forever** to ensure that FabricPool objects aren't deleted by StorageGRID ILM.

Steps

1. Optionally, update the system-generated **Policy name**. By default, the system appends "+ FabricPool" to the name of your active or inactive policy, but you can provide your own name.
2. Review the list of rules in the inactive policy.
 - If your grid doesn't have an inactive ILM policy, the wizard creates an inactive policy by cloning your active policy and adding the new rule to the top.

- If your grid already has an inactive ILM policy and that policy uses the same rules and same order as the active ILM policy, the wizard adds the new rule to the top of the inactive policy.
- If your inactive policy contains different rules or a different order than the active policy, the wizard creates a new inactive policy by cloning your active policy and adding the new rule to the top.

3. Review the order of the rules in the new inactive policy.

Because the FabricPool rule is the first rule, any objects in the FabricPool bucket are placed before the other rules in the policy are evaluated. Objects in any other buckets are placed by subsequent rules in the policy.

4. Review the retention diagram to learn how different objects will be retained.

- Select **Expand all** to see a retention diagram for each rule in the inactive policy.
- Select **Time period** and **Storage pool** to review the retention diagram. Confirm that any rules that apply to the FabricPool bucket or tenant retain objects **forever**.

5. When you have reviewed the inactive policy, select **Activate and continue** to activate the policy and go to the traffic classification step.



Errors in an ILM policy can cause irreparable data loss. Review the policy carefully before activating.

Step 8 of 9: Create traffic classification policy

As an option, the FabricPool setup wizard can create a traffic classification policy that you can use to monitor the FabricPool workload. The system-created policy uses a matching rule to identify all network traffic related to the bucket you created. This policy monitors traffic only; it does not limit traffic for FabricPool or any other clients.

For details about this step, see [Create a traffic classification policy for FabricPool](#).

Steps

1. Review the policy.
2. If you want to create this traffic classification policy, select **Create and continue**.

As soon as FabricPool begins tiering data to StorageGRID, you can go to the Traffic Classification Policies page to view network traffic metrics for this policy. Later, you can also add rules to limit other workloads and ensure that the FabricPool workload has most of the bandwidth.

3. Otherwise, select **Skip this step**.

Step 9 of 9: Review summary

The summary provides details about the items you configured, including the name of the load balancer, tenant, and bucket, the traffic classification policy, and the active ILM policy,

Steps

1. Review the summary.
2. Select **Finish**.

Next steps

After completing the FabricPool wizard, perform these additional steps.

Steps

1. Go to [Configure ONTAP System Manager](#) to enter the saved values and to complete the ONTAP side of the connection. You must add StorageGRID as a cloud tier, attach the cloud tier to a local tier to create a FabricPool, and set volume tiering policies.
2. Go to [Configure the DNS server](#) and ensure that the DNS includes a record to associate the StorageGRID server name (fully qualified domain name) to each StorageGRID IP address you will use.
3. Go to [Other best practices for StorageGRID and FabricPool](#) to learn the best practices for StorageGRID audit logs and other global configuration options.

Configure StorageGRID manually

Create a high availability (HA) group for FabricPool

When configuring StorageGRID for use with FabricPool, you can optionally create one or more high availability (HA) groups.

An HA group is a collection of nodes that each contain the StorageGRID Load Balancer service. An HA group can contain Gateway Nodes, Admin Nodes, or both.

You can use an HA group to help keep FabricPool data connections available. An HA group uses virtual IP addresses (VIPs) to provide highly available access to the Load Balancer service. If the active interface in the HA group fails, a backup interface can manage the workload with little impact to FabricPool operations.

For details about this task, see [Manage high availability groups](#). To use the FabricPool setup wizard to complete this task, go to [Access and complete the FabricPool setup wizard](#).

Before you begin

- You have reviewed the [best practices for high availability groups](#).
- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).
- If you plan to use a VLAN, you have created the VLAN interface. See [Configure VLAN interfaces](#).

Steps

1. Select **CONFIGURATION > Network > High availability groups**.
2. Select **Create**.
3. For the **Enter details** step, complete the following fields.

Field	Description
HA group name	A unique display name for this HA group.
Description (optional)	The description of this HA group.

4. For the **Add interfaces** step, select the node interfaces you want to use in this HA group.

Use the column headers to sort the rows, or enter a search term to locate interfaces more quickly.

You can select one or more nodes, but you can select only one interface for each node.

- For the **Prioritize interfaces** step, determine the Primary interface and any backup interfaces for this HA group.

Drag rows to change the values in the **Priority order** column.

The first interface in the list is the Primary interface. The Primary interface is the active interface unless a failure occurs.

If the HA group includes more than one interface and the active interface fails, the virtual IP (VIP) addresses move to the first backup interface in the priority order. If that interface fails, the VIP addresses move to the next backup interface, and so on. When failures are resolved, the VIP addresses move back to highest priority interface available.

- For the **Enter IP addresses** step, complete the following fields.

Field	Description
Subnet CIDR	The address of the VIP subnet in CIDR notation—an IPv4 address followed by a slash and the subnet length (0-32). The network address must not have any host bits set. For example, 192.16.0.0/22.
Gateway IP address (optional)	Optional. If the ONTAP IP addresses used to access StorageGRID aren't on the same subnet as the StorageGRID VIP addresses, enter the StorageGRID VIP local gateway IP address. The local gateway IP address must be within the VIP subnet.
Virtual IP address	Enter at least one and no more than ten VIP addresses for the active interface in the HA group. All VIP addresses must be within the VIP subnet. At least one address must be IPv4. Optionally, you can specify additional IPv4 and IPv6 addresses.

- Select **Create HA group** and then select **Finish**.

Create a load balancer endpoint for FabricPool

StorageGRID uses a load balancer to manage the workload from client applications, such as FabricPool. Load balancing maximizes speed and connection capacity across multiple Storage Nodes.

When configuring StorageGRID for use with FabricPool, you must configure a load balancer endpoint and upload or generate a load balancer endpoint certificate, which is used to secure the connection between ONTAP and StorageGRID.

To use the FabricPool setup wizard to complete this task, go to [Access and complete the FabricPool setup wizard](#).

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission](#).
- You have reviewed the general [considerations for load balancing](#) as well as the [best practices for load balancing for FabricPool](#).

Steps

1. Select **CONFIGURATION > Network > Load balancer endpoints**.
2. Select **Create**.
3. For the **Enter endpoint details** step, complete the following fields.

Field	Description
Name	A descriptive name for the endpoint.
Port	The StorageGRID port you want to use for load balancing. This field defaults to 10433 for the first endpoint you create, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes. These ports are reserved on Admin Nodes. Note: Ports used by other grid services aren't permitted. See the Network port reference. You will provide this number to ONTAP when you attach StorageGRID as a FabricPool cloud tier.
Client type	Select S3 .
Network protocol	Select HTTPS. Note: Communicating with StorageGRID without TLS encryption is supported but not recommended.

4. For the **Select binding mode** step, specify the binding mode. The binding mode controls how the endpoint is accessed using any IP address or using specific IP addresses and network interfaces.

Mode	Description
Global (default)	Clients can access the endpoint using the IP address of any Gateway Node or Admin Node, the virtual IP (VIP) address of any HA group on any network, or a corresponding FQDN. Use the Global setting (default) unless you need to restrict the accessibility of this endpoint.

Mode	Description
Virtual IPs of HA groups	Clients must use a virtual IP address (or corresponding FQDN) of an HA group to access this endpoint. Endpoints with this binding mode can all use the same port number, as long as the HA groups you select for the endpoints don't overlap.
Node interfaces	Clients must use the IP addresses (or corresponding FQDNs) of selected node interfaces to access this endpoint.
Node type	Based on the type of node you select, clients must use either the IP address (or corresponding FQDN) of any Admin Node or the IP address (or corresponding FQDN) of any Gateway Node to access this endpoint.

5. For the **Tenant access** step, select one of the following:

Field	Description
Allow all tenants (default)	All tenant accounts can use this endpoint to access their buckets. Allow all tenants is almost always the appropriate option for the load balancer endpoint used for FabricPool. You must select this option if you have not yet created any tenant accounts.
Allow selected tenants	Only the selected tenant accounts can use this endpoint to access their buckets.
Block selected tenants	The selected tenant accounts can't use this endpoint to access their buckets. All other tenants can use this endpoint.

6. For the **Attach certificate** step, select one of the following:

Field	Description
Upload certificate (recommended)	Use this option to upload a CA-signed server certificate, certificate private key, and optional CA bundle.
Generate certificate	Use this option to generate a self-signed certificate. See Configure load balancer endpoints for details of what to enter.
Use StorageGRID S3 certificate	This option is available only if you have already uploaded or generated a custom version of the StorageGRID global certificate. See Configure S3 API certificates for details.

7. Select **Create**.



Changes to an endpoint certificate can take up to 15 minutes to be applied to all nodes.

Create a tenant account for FabricPool

You must create a tenant account in the Grid Manager for FabricPool use.

Tenant accounts allow client applications to store and retrieve objects on StorageGRID. Each tenant account has its own account ID, authorized groups and users, buckets, and objects.

For details about this task, see [Create tenant account](#). To use the FabricPool setup wizard to complete this task, go to [Access and complete the FabricPool setup wizard](#).

Before you begin

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

Steps

1. Select **TENANTS**.
2. Select **Create**.
3. For the Enter details steps, enter the following information.

Field	Description
Name	A name for the tenant account. Tenant names don't need to be unique. When the tenant account is created, it receives a unique, numeric account ID.
Description (optional)	A description to help identify the tenant.
Client type	Must be S3 for FabricPool.
Storage quota (optional)	Leave this field blank for FabricPool.

4. For the Select permissions step:

- a. Don't select **Allow platform services**.

FabricPool tenants don't typically need to use platform services, such as CloudMirror replication.

- b. Optionally, select **Use own identity source**.

- c. Don't select **Allow S3 Select**.

FabricPool tenants don't typically need to use S3 Select.

- d. Optionally, select **Use grid federation connection** to allow the tenant to use a [grid federation connection](#) for account clone and cross-grid replication. Then, select the grid federation connection to use.

5. For the Define root access step, specify which user will have the initial Root access permission for the tenant account, based on whether your StorageGRID system uses [identity federation](#), [single sign-on \(SSO\)](#), or both.

Option	Do this
If identity federation is not enabled	Specify the password to use when signing into the tenant as the local root user.
If identity federation is enabled	1. Select an existing federated group to have Root access permission for the tenant. 2. Optionally, specify the password to use when signing in to the tenant as the local root user.
If both identity federation and single sign-on (SSO) are enabled	Select an existing federated group to have Root access permission for the tenant. No local users can sign in.

6. Select **Create tenant**.

Create an S3 bucket and obtain access keys

Before using StorageGRID with a FabricPool workload, you must create an S3 bucket for your FabricPool data. You also need to obtain an access key and secret access key for the tenant account you will use for FabricPool.

For details about this task, see [Create S3 bucket](#) and [Create your own S3 access keys](#). To use the FabricPool setup wizard to complete this task, go to [Access and complete the FabricPool setup wizard](#).

Before you begin

- You have created a tenant account for FabricPool use.
- You have Root access to the tenant account.

Steps

1. Sign in to the Tenant Manager.

You can do either of the following:

- From the Tenant Accounts page in the Grid Manager, select the **Sign in** link for the tenant, and enter your credentials.
- Enter the URL for the tenant account in a web browser, and enter your credentials.

2. Create an S3 bucket for FabricPool data.

You must create a unique bucket for each ONTAP cluster you plan to use.

- Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
- Select **Create bucket**.
- Enter the name of the StorageGRID bucket you want to use with FabricPool. For example, `fabricpool-bucket`.



You can't change the bucket name after creating the bucket.

- Select the region for this bucket.

By default, all buckets are created in the `us-east-1` region.

- e. Select **Continue**.
- f. Select **Create bucket**.



Don't select **Enable object versioning** for the FabricPool bucket. Similarly, don't edit a FabricPool bucket to use **Available** or a non-default consistency. The recommended bucket consistency for FabricPool buckets is **Read-after-new-write**, which is the default consistency for a new bucket.

- 3. Create an access key and a secret access key.
 - a. Select **STORAGE (S3) > My access keys**.
 - b. Select **Create key**.
 - c. Select **Create access key**.
 - d. Copy the access key ID and the secret access key to a safe location, or select **Download .csv** to save a spreadsheet file containing the access key ID and secret access key.

You will enter these values in ONTAP when you configure StorageGRID as a FabricPool cloud tier.



If you generate a new access key and secret access key in StorageGRID in the future, enter the new keys into ONTAP before deleting the old values from StorageGRID. Otherwise, ONTAP might temporarily lose its access to StorageGRID.

Configure ILM for FabricPool data

You can use this simple example policy as a starting point for your own ILM rules and policy.

This example assumes you are designing the ILM rules and an ILM policy for a StorageGRID system that has four Storage Nodes at a single data center in Denver, Colorado. The FabricPool data in this example uses a bucket named `fabricpool-bucket`.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate it to confirm it will work as intended to protect content from loss. To learn more, see [Manage objects with ILM](#).



To avoid data loss, do not use an ILM rule that will expire or delete FabricPool cloud tier data. Set the retention period to **forever** to ensure that FabricPool objects aren't deleted by StorageGRID ILM.

Before you begin

- You have reviewed the [best practices for using ILM with FabricPool data](#).
- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [ILM or Root access permission](#).
- If you upgraded to StorageGRID 11.9 from a previous StorageGRID version, you have configured the storage pool you will use. In general, you should create a storage pool for each StorageGRID site you will use to store data.



This prerequisite does not apply if you initially installed StorageGRID 11.7 or 11.8. When you initially install either of these versions, storage pools are automatically created for each site.

Steps

1. Create an ILM rule that applies only to the data in `fabricpool-bucket`. This example rule creates erasure-coded copies.

Rule definition	Example value
Rule name	2 + 1 erasure coding for FabricPool data
Bucket name	<code>fabricpool-bucket</code> You could also filter on the FabricPool tenant account.
Advanced filters	Object size greater than 0.2 MB. Note: FabricPool only writes 4 MB objects, but you must add an Object size filter because this rule uses erasure coding.
Reference time	Ingest time
Time period and placements	From Day 0 store forever Store objects by erasure coding using 2+1 EC scheme at Denver and retain those objects in StorageGRID forever.  To avoid data loss, do not use an ILM rule that will expire or delete FabricPool cloud tier data.
Ingest behavior	Balanced

2. Create a default ILM rule that will create two replicated copies of any objects not matched by the first rule. Don't select a basic filter (tenant account or bucket name) or any advanced filters.

Rule definition	Example value
Rule name	Two replicated copies
Bucket name	<i>none</i>
Advanced filters	<i>none</i>
Reference time	Ingest time

Rule definition	Example value
Time period and placements	From Day 0 store forever Store objects by replicating 2 copies at Denver.
Ingest behavior	Balanced

3. Create an ILM policy and select the two rules. Because the replication rule does not use any filters, it can be the default (last) rule for the policy.
4. Ingest test objects into the grid.
5. Simulate the policy with the test objects to verify the behavior.
6. Activate the policy.

When this policy is activated, StorageGRID places object data as follows:

- The data tiered from FabricPool in `fabricpool-bucket` will be erasure-coded using the 2+1 erasure-coding scheme. Two data fragments and one parity fragment will be placed on three different Storage Nodes.
- All objects in all other buckets will be replicated. Two copies will be created and placed on two different Storage Nodes.
- The copies will be maintained in StorageGRID forever. StorageGRID ILM won't delete these objects.

Create a traffic classification policy for FabricPool

You can optionally design a StorageGRID traffic classification policy to optimize quality of service for the FabricPool workload.

For details about this task, see [Manage traffic classification policies](#). To use the FabricPool setup wizard to complete this task, go to [Access and complete the FabricPool setup wizard](#).

Before you begin

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

About this task

The best practices for creating a traffic classification policy for FabricPool depend on the workload, as follows:

- If you plan to tier FabricPool primary workload data to StorageGRID, you should ensure that the FabricPool workload has most of the bandwidth. You can create a traffic classification policy to limit all other workloads.



In general, FabricPool read operations are more important to prioritize than write operations.

For example, if other S3 clients use this StorageGRID system, you should create a traffic classification policy. You can limit network traffic for the other buckets, tenants, IP subnets, or load balancer endpoints.

- Generally, you should not impose quality of service limits on any FabricPool workload; you should only limit the other workloads.

- The limits placed on other workloads should account for the behavior of those workloads. The limits imposed will also vary based on the sizing and capabilities of your grid and what the expected amount of utilization is.

Steps

1. Select **CONFIGURATION > Network > Traffic classification**.
2. Select **Create**.
3. Enter a name and a description (optional) for the policy and select **Continue**.
4. For the Add matching rules step, add at least one rule.
 - a. Select **Add rule**
 - b. For Type, select **Load balancer endpoint**, and select the load balancer endpoint you created for FabricPool.

You can also select the FabricPool tenant account or bucket.

- c. If you want this traffic policy to limit traffic for the other endpoints, select **Inverse match**.
5. Optionally, add one or more limits to control the network traffic matched by the rule.



StorageGRID collects metrics even if you don't add any limits, so you can understand traffic trends.

- a. Select **Add a limit**.
 - b. Select the type of traffic you want to limit and the limit to apply.
6. Select **Continue**.
 7. Read and review the Traffic classification policy. Use the **Previous** button to go back and make changes as required. When you are satisfied with the policy, select **Save and continue**.

After your finish

[View network traffic metrics](#) to verify that the policies are enforcing the traffic limits you expect.

Configure ONTAP System Manager

After you have obtained the required StorageGRID information, you can go to ONTAP to add StorageGRID as a cloud tier.

Before you begin

- If you completed the FabricPool setup wizard, you have the `ONTAP_FabricPool_settings_bucketname.txt` file you downloaded.
- If you configured StorageGRID manually, you have the fully qualified domain name (FQDN) you are using for StorageGRID or the virtual IP (VIP) address for the StorageGRID HA group, the port number for the load balancer endpoint, the load balancer certificate, the access key ID and secret key for the root user of the tenant account, and the name of the bucket ONTAP will use in that tenant.

Access ONTAP System Manager

These instructions describe how to use ONTAP System Manager to add StorageGRID as a cloud tier. You can complete the same configuration using the ONTAP CLI. For instructions, go to [ONTAP documentation for FabricPool](#).

Steps

1. Access System Manager for the ONTAP cluster you want to tier to StorageGRID.
2. Sign in as an administrator for the cluster.
3. Navigate to **STORAGE > Tiers > Add Cloud Tier**.
4. Select **StorageGRID** from the list of object store providers.

Enter StorageGRID values

See [ONTAP documentation for FabricPool](#) for more information.

Steps

1. Complete the Add Cloud Tier form, using the `ONTAP_FabricPool_settings_bucketname.txt` file or the values you obtained manually.

Field	Description
Name	Enter a unique name for this cloud tier. You can accept the default value.
URL style	If you configured S3 endpoint domain names, select Virtual Hosted-Style URL. Path-Style URL is the default for ONTAP, but using virtual hosted-style requests is recommended for StorageGRID. You must use Path-Style URL if you provide an IP address instead of a domain name for the Server name (FQDN) field.
Server name (FQDN)	Enter the fully qualified domain name (FQDN) you are using for StorageGRID or the virtual IP (VIP) address for the StorageGRID HA group. For example, <code>s3.storagegrid.company.com</code>. Note the following: • The IP address or domain name that you specify here must match the certificate you uploaded or generated for the StorageGRID load balancer endpoint. • If you provide a domain name, the DNS record must map to each IP address you will use to connect to StorageGRID. See Configure the DNS server.
SSL	Enabled (default).
Object store certificate	Paste the certificate PEM you are using for the StorageGRID load balancer endpoint, including: <pre>-----BEGIN CERTIFICATE----- and -----END CERTIFICATE-----.</pre> Note: If an intermediate CA issued the StorageGRID certificate, you must provide the intermediate CA certificate. If the StorageGRID certificate was issued directly by the Root CA, you must provide the Root CA certificate.

Field	Description
Port	Enter the port used by the StorageGRID load balancer endpoint. ONTAP will use this port when it connects to StorageGRID. For example, 10433.
Access key and secret key	Enter the access key ID and secret access key for the root user of the StorageGRID tenant account. Tip: If you generate a new access key and secret access key in StorageGRID in the future, enter the new keys into ONTAP before deleting the old values from StorageGRID. Otherwise, ONTAP might temporarily lose its access to StorageGRID.
Container name	Enter the name of the StorageGRID bucket you created for use with this ONTAP tier.

2. Complete the final FabricPool configuration in ONTAP.
 - a. Attach one or more aggregates to the cloud tier.
 - b. Optionally, create a volume tiering policy.

Configure the DNS server

After configuring high availability groups, load balancer endpoints, and S3 endpoint domain names, you must ensure that the DNS includes the necessary entries for StorageGRID. You must include a DNS entry for each name in the security certificate and for each IP address you might use.

See [Considerations for load balancing](#).

DNS entries for StorageGRID server name

Add DNS entries to associate the StorageGRID server name (fully qualified domain name) to each StorageGRID IP address you will use.

The IP addresses you enter in the DNS depend on whether you are using an HA group of load-balancing nodes:

- If you have configured an HA group, ONTAP will connect to the virtual IP addresses of that HA group.
- If you aren't using an HA group, ONTAP can connect to the StorageGRID Load Balancer service using the IP address of any Gateway Node or Admin Node.
- If the server name resolves to more than one IP address, ONTAP establishes client connections with all IP addresses (up to a maximum of 16 IP addresses). The IP addresses are picked up in a round-robin method when connections are established.

DNS entries for virtual hosted-style requests

If you have defined [S3 endpoint domain names](#) and you will use virtual hosted-style requests, add DNS entries for all required S3 endpoint domain names, including any wildcard names.

StorageGRID best practices for FabricPool

Best practices for high availability (HA) groups

Before attaching StorageGRID as a FabricPool cloud tier, learn about StorageGRID high availability (HA) groups and review the best practices for using HA groups with FabricPool.

What is an HA group?

A high availability (HA) group is a collection of interfaces from multiple StorageGRID Gateway Nodes, Admin Nodes, or both. An HA group helps to keep client data connections available. If the active interface in the HA group fails, a backup interface can manage the workload with little impact on FabricPool operations.

Each HA group provides highly available access to the shared services on the associated nodes. For example, an HA group that consists of interfaces only on Gateway Nodes or on both Admin Nodes and Gateway Nodes provides highly available access to the shared Load Balancer service.

To learn more about high availability groups, see [Manage high availability \(HA\) groups](#).

Using HA groups

The best practices for creating a StorageGRID HA group for FabricPool depend on the workload.

- If you plan to use FabricPool with primary workload data, you must create an HA group that includes at least two load-balancing nodes to prevent data retrieval interruption.
- If you plan to use the FabricPool snapshot-only volume tiering policy or non-primary local performance tiers (for example, disaster recovery locations or NetApp SnapMirror® destinations), you can configure an HA group with only one node.

These instructions describe setting up an HA group for Active-Backup HA (one node is active and one node is backup). However, you might prefer to use DNS Round Robin or Active-Active HA. To learn the benefits of these other HA configurations, see [Configuration options for HA groups](#).

Best practices for load balancing for FabricPool

Before attaching StorageGRID as a FabricPool cloud tier, review the best practices for using load balancers with FabricPool.

To learn general information about the StorageGRID load balancer and the load balancer certificate, see [Considerations for load balancing](#).

Best practices for tenant access to the load balancer endpoint used for FabricPool

You can control which tenants can use a specific load balancer endpoint to access their buckets. You can allow all tenants, allow some tenants, or block some tenants. When creating a load balance endpoint for FabricPool use, select **Allow all tenants**. ONTAP encrypts the data that is placed in StorageGRID buckets, so little additional security would be provided by this extra security layer.

Best practices for the security certificate

When you create a StorageGRID load balancer endpoint for FabricPool use, you provide the security certificate that will allow ONTAP to authenticate with StorageGRID.

In most cases, the connection between ONTAP and StorageGRID should use Transport Layer Security (TLS) encryption. Using FabricPool without TLS encryption is supported but not recommended. When you select the network protocol for the StorageGRID load balancer endpoint, select **HTTPS**. Then provide the security certificate that will allow ONTAP to authenticate with StorageGRID.

To learn more about the server certificate for a load balancing endpoint:

- [Manage security certificates](#)
- [Considerations for load balancing](#)
- [Hardening guidelines for server certificates](#)

Add certificate to ONTAP

When you add StorageGRID as a FabricPool cloud tier, you must install the same certificate on the ONTAP cluster, including the root and any subordinate certificate authority (CA) certificates.

Manage certificate expiration



If the certificate used to secure the connection between ONTAP and StorageGRID expires, FabricPool will temporarily stop working and ONTAP will temporarily lose access to data tiered to StorageGRID.

To avoid certificate expiration issues, follow these best practices:

- Carefully monitor any alerts that warn of approaching certificate expiration dates, such as the **Expiration of load balancer endpoint certificate** and **Expiration of global server certificate for S3 API** alerts.
- Always keep the StorageGRID and ONTAP versions of the certificate in sync. If you replace or renew the certificate used for a load balancer endpoint, you must replace or renew the equivalent certificate used by ONTAP for the cloud tier.
- Use a publicly signed CA certificate. If you use a certificate signed by a CA, you can use the Grid Management API to automate certificate rotation. This allows you to replace soon-to-expire certificates nondisruptively.
- If you have generated a self-signed StorageGRID certificate and that certificate is about to expire, you must manually replace the certificate in both StorageGRID and in ONTAP before the existing certificate expires. If a self-signed certificate has already expired, turn off certificate validation in ONTAP to prevent access loss.

See [NetApp Knowledge Base: How to configure a new StorageGRID self-signed server certificate on an existing ONTAP FabricPool deployment](#) for instructions.

Best practices for using ILM with FabricPool data

If you are using FabricPool to tier data to StorageGRID, you must understand the requirements for using StorageGRID information lifecycle management (ILM) with FabricPool data.



FabricPool has no knowledge of StorageGRID ILM rules or policies. Data loss can occur if the StorageGRID ILM policy is misconfigured. For detailed information, see [Use ILM rules to manage objects](#) and [Create ILM policies](#).

Guidelines for using ILM with FabricPool

When you use the FabricPool setup wizard, the wizard automatically creates a new ILM rule for each S3 bucket you create and adds that rule to an inactive policy. You are prompted to activate the policy. The automatically created rule follows the recommended best practices: it uses 2+1 erasure coding at a single site.

If you are configuring StorageGRID manually instead of using the FabricPool setup wizard, review these guidelines to ensure that your ILM rules and ILM policy are suitable for FabricPool data and your business requirements. You might need to create new rules and update your active ILM policies to meet these guidelines.

- You can use any combination of replication and erasure-coding rules to protect cloud tier data.

The recommended best practice is to use 2+1 erasure coding within a site for cost-efficient data protection. Erasure coding uses more CPU, but offers significantly less storage capacity, than replication. The 4+1 and 6+1 schemes use less capacity than the 2+1 scheme. However, the 4+1 and 6+1 schemes are less flexible if you need to add Storage Nodes during grid expansion. For details, see [Add storage capacity for erasure-coded objects](#).

- Each rule applied to FabricPool data must either use erasure coding or it must create at least two replicated copies.



An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.

- If you need to [remove FabricPool data from StorageGRID](#), use ONTAP to retrieve all data for the FabricPool volume and promote it to the performance tier.



To avoid data loss, do not use an ILM rule that will expire or delete FabricPool cloud tier data. Set the retention period in each ILM rule to **forever** to ensure that FabricPool objects aren't deleted by StorageGRID ILM.

- Don't create rules that will move FabricPool cloud tier data out of the bucket to another location. You can't use a Cloud Storage Pool to move FabricPool data to another object store.



Using Cloud Storage Pools with FabricPool is not supported because of the added latency to retrieve an object from the Cloud Storage Pool target.

- Starting with ONTAP 9.8, you can optionally create object tags to help classify and sort tiered data for easier management. For example, you can set tags only on FabricPool volumes attached to StorageGRID. Then, when you create ILM rules in StorageGRID, you can use the Object Tag advanced filter to select and place this data.

Other best practices for StorageGRID and FabricPool

When configuring a StorageGRID system for use with FabricPool, you might need to change other StorageGRID options. Before changing a global setting, consider how the change will affect other S3 applications.

Audit message and log destinations

FabricPool workloads often have a high rate of read operations, which can generate a high volume of audit messages.

- If you don't require a record of client read operations for FabricPool or any other S3 application, optionally go to **CONFIGURATION > Monitoring > Audit and syslog server**. Change the **Client Reads** setting to **Error** to decrease the number of audit messages recorded in the audit log. See [Configure audit messages and log destinations](#) for details.
- If you have a large grid, use multiple types of S3 applications, or want to retain all audit data, configure an external syslog server and save audit information remotely. Using an external server minimizes the performance impact of audit message logging without reducing the completeness of of audit data. See [Considerations for external syslog server](#) for details.

Object encryption

When configuring StorageGRID, you can optionally enable the [global option for stored object encryption](#) if data encryption is required for other StorageGRID clients. The data that is tiered from FabricPool to StorageGRID is already encrypted, so enabling the StorageGRID setting is not required. Client-side encryption keys are owned by ONTAP.

Object compression

When configuring StorageGRID, don't enable the [global option to compress stored objects](#). The data that is tiered from FabricPool to StorageGRID is already compressed. Using the StorageGRID option will not further reduce an object's size.

Bucket consistency

For FabricPool buckets, the recommended bucket consistency is **Read-after-new-write**, which is the default consistency for a new bucket. Don't edit FabricPool buckets to use **Available** or **Strong-site**.

FabricPool tiering

If a StorageGRID node uses storage assigned from a NetApp ONTAP system, confirm that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running on a VMware host, ensure the volume backing the datastore for the StorageGRID node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

Remove FabricPool data from StorageGRID

If you need to remove the FabricPool data that is currently stored in StorageGRID, you must use ONTAP to retrieve all data for the FabricPool volume and promote it to the performance tier.

Before you begin

- You have reviewed the instructions and considerations in [Promote data to the performance tier](#).

- You are using ONTAP 9.8 or later.
- You are using a [supported web browser](#).
- You belong to a StorageGRID user group for the FabricPool tenant account that has the [Manage all buckets or Root access permission](#).

About this task

These instructions explain how to move data from StorageGRID back to FabricPool. You perform this procedure using ONTAP and StorageGRID Tenant Manager.

Steps

1. From ONTAP, issue the `volume modify` command.

Set `tiering-policy` to `none` to stop new tiering and set `cloud-retrieval-policy` to `promote` to return all data that was previously tiered to StorageGRID.

See [Promote all data from a FabricPool volume to the performance tier](#).

2. Wait for the operation to complete.

You can use the `volume object-store` command with the `tiering` option to [check the status of the performance tier promotion](#).

3. When the promote operation is complete, sign in to StorageGRID Tenant Manager for the FabricPool tenant account.
4. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
5. Confirm that the FabricPool bucket is now empty.
6. If the bucket is empty, [delete the bucket](#).

After you finish

When you delete the bucket, tiering from FabricPool to StorageGRID can no longer continue. However, because the local tier is still attached to the StorageGRID cloud tier, ONTAP System Manager will return error messages indicating that the bucket is inaccessible.

To prevent these error messages, do either of the following:

- Use FabricPool Mirror to attach a different cloud tier to the aggregate.
- Move the data from the FabricPool aggregate to a non-FabricPool aggregate and then delete the unused aggregate.

See the [ONTAP documentation for FabricPool](#) for instructions.

Use StorageGRID tenants and clients

Use a tenant account

Use a tenant account

A tenant account allows you to use either the Simple Storage Service (S3) REST API or the Swift REST API to store and retrieve objects in a StorageGRID system.

What is a tenant account?

Each tenant account has its own federated or local groups, users, S3 buckets or Swift containers, and objects.

Tenant accounts can be used to segregate stored objects by different entities. For example, multiple tenant accounts can be used for either of these use cases:

- **Enterprise use case:** If the StorageGRID system is being used within an enterprise, the grid's object storage might be segregated by the different departments in the organization. For example, there might be tenant accounts for the Marketing department, the Customer Support department, the Human Resources department, and so on.



If you use the S3 client protocol, you can also use S3 buckets and bucket policies to segregate objects between the departments in an enterprise. You don't need to create separate tenant accounts. See instructions for implementing [S3 buckets and bucket policies](#) for more information.

- **Service provider use case:** If the StorageGRID system is being used by a service provider, the grid's object storage might be segregated by the different entities that lease the storage. For example, there might be tenant accounts for Company A, Company B, Company C, and so on.

How to create a tenant account

Tenant accounts are created by a [StorageGRID grid administrator using the Grid Manager](#). When creating a tenant account, the grid administrator specifies the following:

- Basic information including the tenant name, client type (S3) and optional storage quota.
- Permissions for the tenant account, such as whether the tenant account can use S3 platform services, configure its own identity source, use S3 Select, or use a grid federation connection.
- The initial root access for the tenant, based on whether the StorageGRID system uses local groups and users, identity federation, or single sign-on (SSO).

In addition, grid administrators can enable the S3 Object Lock setting for the StorageGRID system if S3 tenant accounts need to comply with regulatory requirements. When S3 Object Lock is enabled, all S3 tenant accounts can create and manage compliant buckets.

Configure S3 tenants

After an [S3 tenant account is created](#), you can access the Tenant Manager to perform tasks such as the following:

- Set up identity federation (unless the identity source is shared with the grid)

- Manage groups and users
- Use grid federation for account clone and cross-grid replication
- Manage S3 access keys
- Create and manage S3 buckets
- Use S3 platform services
- Use S3 Select
- Monitor storage usage



Although you can create and manage S3 buckets with the Tenant Manager, you must use an [S3 client](#) or [S3 Console](#) to ingest and manage objects.

How to sign in and sign out

Sign in to Tenant Manager

You access the Tenant Manager by entering the URL for the tenant into the address bar of a [supported web browser](#).

Before you begin

- You have your login credentials.
- You have a URL for accessing the Tenant Manager, as supplied by your grid administrator. The URL will look like one of these examples:

```
https://FQDN_or_Admin_Node_IP/
```

```
https://FQDN_or_Admin_Node_IP:port/
```

```
https://FQDN_or_Admin_Node_IP/?accountId=20-digit-account-id
```

```
https://FQDN_or_Admin_Node_IP:port/?accountId=20-digit-account-id
```

The URL always includes a fully qualified domain name (FQDN), the IP address of an Admin Node, or the virtual IP address of an HA group of Admin Nodes. It might also include a port number, the 20-digit tenant account ID, or both.

- If the URL does not include the tenant's 20-digit account ID, you have this account ID.
- You are using a [supported web browser](#).
- Cookies are enabled in your web browser.
- You belong to a user group that has [specific access permissions](#).

Steps

1. Launch a [supported web browser](#).
2. In the browser's address bar, enter the URL for accessing Tenant Manager.
3. If you are prompted with a security alert, install the certificate using the browser's installation wizard.
4. Sign in to the Tenant Manager.

The sign-in screen that appears depends on the URL you entered and whether single sign-on (SSO) has been configured for StorageGRID.

Not using SSO

If StorageGRID is not using SSO, one of the following screens appears:

- The Grid Manager sign-in page. Select the **Tenant sign-in** link.



NetApp StorageGRID®

Grid Manager

Username

Password

[Sign in](#)

[Tenant sign in](#) | [NetApp support](#) | [NetApp.com](#)

- The Tenant Manager sign-in page. The **Account** field might already be completed, as shown below.

The screenshot shows the NetApp StorageGRID Tenant Manager login interface. At the top is the NetApp StorageGRID logo. Below it is the title 'Tenant Manager'. The form includes a 'Recent' dropdown menu currently showing '-- Optional --'. Under the 'Account' section, a text box contains the 20-digit account ID '64600207336181242061'. The 'Username' section has an empty text box with a cursor. The 'Password' section has an empty password field. A blue 'Sign in' button is located below the password field. At the bottom, there are links for 'NetApp support' and 'NetApp.com'.

- a. If the tenant's 20-digit account ID is not shown, select the name of the tenant account if it appears in the list of recent accounts, or enter the account ID.
- b. Enter your username and password.
- c. Select **Sign in**.

The Tenant Manager dashboard appears.

- d. If you received an initial password from someone else, select **username > Change password** to secure your account.

Using SSO

If StorageGRID is using SSO, one of the following screens appears:

- Your organization's SSO page. For example:

Sign in with your organizational account

Sign in

Enter your standard SSO credentials, and select **Sign in**.

- The Tenant Manager SSO sign-in page.

NetApp StorageGRID®

Tenant Manager

Recent

S3 tenant ▼

Account

62984032838045582045

Sign in

[NetApp support](#) | [NetApp.com](#)

- If the tenant's 20-digit account ID is not shown, select the name of the tenant account if it appears in the list of recent accounts, or enter the account ID.
- Select **Sign in**.
- Sign in with your standard SSO credentials on your organization's SSO sign-in page.

The Tenant Manager dashboard appears.


Sign out of Tenant Manager

When you are done working with the Tenant Manager, you must sign out to ensure that unauthorized users can't access the StorageGRID system. Closing your browser might

not sign you out of the system, based on browser cookie settings.

Steps

1. Locate the username drop-down in the top-right corner of the user interface.



2. Select the username and then select **Sign out**.

- If SSO is not in use:

You are signed out of the Admin Node. The Tenant Manager sign in page is displayed.



If you signed into more than one Admin Node, you must sign out of each node.

- If SSO is enabled:

You are signed out of all Admin Nodes you were accessing. The StorageGRID Sign in page is displayed. The name of the tenant account you just accessed is listed as the default in the **Recent Accounts** drop-down, and the tenant's **Account ID** is shown.



If SSO is enabled and you are also signed in to the Grid Manager, you must also sign out of the Grid Manager to sign out of SSO.

Understand Tenant Manager dashboard

The Tenant Manager dashboard provides an overview of a tenant account's configuration and the amount of space used by objects in the tenant's buckets (S3) or containers (Swift). If the tenant has a quota, the dashboard shows how much of the quota is used and how much is remaining. If there are any errors related to the tenant account, the errors are shown on the dashboard.



The Space used values are estimates. These estimates are affected by the timing of ingests, network connectivity, and node status.

When objects have been uploaded, the dashboard looks like the following example:

Dashboard

16**Buckets**[View buckets](#)**2****Platform services****endpoints**
[View endpoints](#)**0****Groups**[View groups](#)**1****User**[View users](#)

Storage usage ?

6.5 TB of 7.2 TB used

0.7 TB (10.1%) remaining



Bucket name	Space used	Number of objects
Bucket-15	969.2 GB	913,425
Bucket-04	937.2 GB	576,806
Bucket-13	815.2 GB	957,389
Bucket-06	812.5 GB	193,843
Bucket-10	473.9 GB	583,245
Bucket-03	403.2 GB	981,226
Bucket-07	362.5 GB	420,726
Bucket-05	294.4 GB	785,190
8 other buckets	1.4 TB	3,007,036

Top buckets by capacity limit usage ?

Bucket name	Usage
Bucket-10	82%
Bucket-03	57%
Bucket-15	20%

Tenant details ?

Name: Tenant02

ID: 3341 1240 0546 8283 2208

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

Tenant account information

The top of the dashboard displays the number of configured buckets or containers, groups, and users. It also displays the number of platform services endpoints, if any have been configured. Select the links to view the details.

Depending on the [tenant management permissions](#) you have and the options you've configured, the remainder of the dashboard displays various combinations of guidelines, storage usage, object information, and tenant details.

Storage and quota usage

The Storage usage panel contains the following information:

- The amount of object data for the tenant.

This value indicates the total amount of object data uploaded and does not represent the space used to store copies of those objects and their metadata.

- If a quota is set, the total amount of space available for object data and the amount and percentage of space remaining. The quota limits the amount of object data that can be ingested.



Quota usage is based on internal estimates and might be exceeded in some cases. For example, StorageGRID checks the quota when a tenant starts uploading objects and rejects new ingests if the tenant has exceeded the quota. However, StorageGRID does not take into account the size of the current upload when determining if the quota has been exceeded. If objects are deleted, a tenant might be temporarily prevented from uploading new objects until the quota usage is recalculated. Quota usage calculations can take 10 minutes or longer.

- A bar chart that represents the relative sizes of the largest buckets or containers.

You can place your cursor over any of the chart segments to view the total space consumed by that bucket or container.



- To correspond with the bar chart, a list of the largest buckets or containers, including the total amount of object data and the number of objects for each bucket or container.

Bucket name	Space used	Number of objects
 Bucket-02	944.7 GB	7,575
 Bucket-09	899.6 GB	589,677
 Bucket-15	889.6 GB	623,542
 Bucket-06	846.4 GB	648,619
 Bucket-07	730.8 GB	808,655
 Bucket-04	700.8 GB	420,493
 Bucket-11	663.5 GB	993,729
 Bucket-03	656.9 GB	379,329
 9 other buckets	2.3 TB	5,171,588

If the tenant has more than nine buckets or containers, all other buckets or containers are combined into a single entry at the bottom of the list.



To change units for the storage values displayed in the Tenant Manager, select the user drop-down in the upper right of the Tenant Manager, then select **User preferences**.

Quota usage alerts

If quota usage alerts have been enabled in the Grid Manager, these alerts will appear in the Tenant Manager when the quota is low or exceeded, as follows:

- If 90% or more of a tenant's quota has been used, the **Tenant quota usage high** alert is triggered.

Consider asking your grid administrator to increase the quota.

- If you exceed your quota, a notification tells you that you can't upload new objects.

Capacity limit usage

If you've set a capacity limit for your buckets, the Tenant Manager dashboard displays a list of top buckets by capacity limit usage.

If no limit is set for a bucket, its capacity is unlimited. However, if your tenant account has a total storage quota and that quota is reached, you won't be able to ingest more objects regardless of the remaining capacity limit on a bucket.

Endpoint errors

If you have used the Grid Manager to configure one or more endpoints for use with platform services, the Tenant Manager dashboard displays an alert if any endpoint errors have occurred within the past seven days.

 One or more endpoints have experienced an error and might not be functioning properly. Go to the [Endpoints](#) page to view the error details. The last error occurred 2 hours ago.

To see details about [platform services endpoint errors](#), select **Endpoints** to display the Endpoints page.

Tenant Management API

Understand Tenant Management API

You can perform system management tasks using the Tenant Management REST API instead of the Tenant Manager user interface. For example, you might want to use the API to automate operations or to create multiple entities, such as users, more quickly.

The Tenant Management API:

- Uses the Swagger open source API platform. Swagger provides an intuitive user interface that allows developers and non-developers to interact with the API. The Swagger user interface provides complete details and documentation for each API operation.
- Uses [versioning to support non-disruptive upgrades](#).

To access the Swagger documentation for the Tenant Management API:

1. Sign in to the Tenant Manager.
2. From the top of the Tenant Manager, select the help icon and select **API documentation**.

API operations

The Tenant Management API organizes the available API operations into the following sections:

- **account**: Operations on the current tenant account, including getting storage usage information.
- **auth**: Operations to perform user session authentication.

The Tenant Management API supports the Bearer Token Authentication Scheme. For a tenant login, you provide a username, password, and accountId in the JSON body of the authentication request (that is, `POST /api/v3/authorize`). If the user is successfully authenticated, a security token is returned. This token must be provided in the header of subsequent API requests ("Authorization: Bearer token").

For information about improving authentication security, see [Protect against Cross-Site Request Forgery](#).



If single sign-on (SSO) is enabled for the StorageGRID system, you must perform different steps to authenticate. See the [instructions for using the Grid Management API](#).

- **config:** Operations related to the product release and versions of the Tenant Management API. You can list the product release version and the major versions of the API supported by that release.
- **containers:** Operations on S3 buckets or Swift containers.
- **deactivated-features:** Operations to view features that might have been deactivated.
- **endpoints:** Operations to manage an endpoint. Endpoints allow an S3 bucket to use an external service for StorageGRID CloudMirror replication, notifications, or search integration.
- **grid-federation-connections:** Operations on grid federation connections and cross-grid replication.
- **groups:** Operations to manage local tenant groups and to retrieve federated tenant groups from an external identity source.
- **identity-source:** Operations to configure an external identity source and to manually synchronize federated group and user information.
- **ilm:** Operations on information lifecycle management (ILM) settings.
- **regions:** Operations to determine which regions have been configured for the StorageGRID system.
- **s3:** Operations to manage S3 access keys for tenant users.
- **s3-object-lock:** Operations on global S3 Object Lock settings, used to support regulatory compliance.
- **users:** Operations to view and manage tenant users.

Operation details

When you expand each API operation, you can see its HTTP action, endpoint URL, a list of any required or optional parameters, an example of the request body (when required), and the possible responses.

groups
Operations on groups

GET
/org/groups
Lists Tenant User Groups

Parameters
Try it out

Name	Description
type string (query)	filter by group type
limit integer (query)	maximum number of results
marker string (query)	marker-style pagination offset (value is Group's URN)
includeMarker boolean (query)	if set, the marker element is also returned
order string (query)	pagination order (desc requires marker)

Responses
Response content type
application/json

Code	Description
200	Example Value Model <pre>{ "responseTime": "2018-02-01T16:22:31.066Z", "status": "success", "apiVersion": "2.1" }</pre>

Issue API requests



Any API operations you perform using the API Documentation webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Steps

1. Select the HTTP action to see the request details.
2. Determine if the request requires additional parameters, such as a group or user ID. Then, obtain these values. You might need to issue a different API request first to get the information you need.
3. Determine if you need to modify the example request body. If so, you can select **Model** to learn the requirements for each field.
4. Select **Try it out**.

5. Provide any required parameters, or modify the request body as required.
6. Select **Execute**.
7. Review the response code to determine if the request was successful.

Tenant Management API versioning

The Tenant Management API uses versioning to support non-disruptive upgrades.

For example, this Request URL specifies version 4 of the API.

```
https://hostname_or_ip_address/api/v4/authorize
```

The major version of the API is bumped when changes are made that are *not compatible* with older versions. The minor version of the API is bumped when changes are made that *are compatible* with older versions. Compatible changes include the addition of new endpoints or new properties.

The following example illustrates how the API version is bumped based on the type of changes made.

Type of change to API	Old version	New version
Compatible with older versions	2.1	2.2
Not compatible with older versions	2.1	3.0
	3.0	4.0

When you install StorageGRID software for the first time, only the most recent version of the API is enabled. However, when you upgrade to a new feature release of StorageGRID, you continue to have access to the older API version for at least one StorageGRID feature release.



You can configure the supported versions. See the **config** section of the Swagger API documentation for the [Grid Management API](#) for more information. You should deactivate support for the older version after updating all API clients to use the newer version.

Outdated requests are marked as deprecated in the following ways:

- The response header is "Deprecated: true"
- The JSON response body includes "deprecated": true
- A deprecated warning is added to nms.log. For example:

```
Received call to deprecated v2 API at POST "/api/v2/authorize"
```

Determine which API versions are supported in the current release

Use the `GET /versions` API request to return a list of the supported API major versions. This request is located in the **config** section of the Swagger API documentation.

```
GET https://{{IP-Address}}/api/versions
{
  "responseTime": "2023-06-27T22:13:50.750Z",
  "status": "success",
  "apiVersion": "4.0",
  "data": [
    2,
    3,
    4
  ]
}
```

Specify an API version for a request

You can specify the API version using a path parameter (/api/v4) or a header (Api-Version: 4). If you provide both values, the header value overrides the path value.

```
curl https://[IP-Address]/api/v4/grid/accounts

curl -H "Api-Version: 4" https://[IP-Address]/api/grid/accounts
```

Protect against Cross-Site Request Forgery (CSRF)

You can help protect against Cross-Site Request Forgery (CSRF) attacks against StorageGRID by using CSRF tokens to enhance authentication that uses cookies. The Grid Manager and Tenant Manager automatically enable this security feature; other API clients can choose whether to enable it when they sign in.

An attacker that can trigger a request to a different site (such as with an HTTP form POST) can cause certain requests to be made using the signed-in user's cookies.

StorageGRID helps protect against CSRF attacks by using CSRF tokens. When enabled, the contents of a specific cookie must match the contents of either a specific header or a specific POST body parameter.

To enable the feature, set the `csrfToken` parameter to `true` during authentication. The default is `false`.

```
curl -X POST --header "Content-Type: application/json" --header "Accept:
application/json" -d "{
  \"username\": \"MyUserName\",
  \"password\": \"MyPassword\",
  \"cookie\": true,
  \"csrfToken\": true
}" "https://example.com/api/v3/authorize"
```

When true, a `GridCsrfToken` cookie is set with a random value for sign-ins to the Grid Manager, and the

`AccountCsrfToken` cookie is set with a random value for sign-ins to the Tenant Manager.

If the cookie is present, all requests that can modify the state of the system (POST, PUT, PATCH, DELETE) must include one of the following:

- The `X-Csrf-Token` header, with the value of the header set to the value of the CSRF token cookie.
- For endpoints that accept a form-encoded body: A `csrfToken` form-encoded request body parameter.

To configure CSRF protection, use the [Grid Management API](#) or [Tenant Management API](#).



Requests that have a CSRF token cookie set will also enforce the "Content-Type: application/json" header for any request that expects a JSON request body as an additional protection against CSRF attacks.

Use grid federation connections

Clone tenant groups and users


If a tenant was created or edited to use a grid federation connection, that tenant is replicated from one StorageGRID system (the source tenant) to another StorageGRID system (the replica tenant). After the tenant has been replicated, any groups and users added to the source tenant are cloned to the replica tenant.

The StorageGRID system where the tenant is originally created is the tenant's *source grid*. The StorageGRID system where the tenant is replicated is the tenant's *destination grid*. Both tenant accounts have the same account ID, name, description, storage quota, and assigned permissions, but the destination tenant does not initially have a root user password. For details, see [What is account clone](#) and [Manage permitted tenants](#).

The cloning of tenant account information is required for [cross-grid replication](#) of bucket objects. Having the same tenant groups and users on both grids ensures you can access the corresponding buckets and objects on either grid.

Tenant workflow for account clone

If your tenant account has the **Use grid federation connection** permission, review the workflow diagram to see the steps you will perform to clone groups, users, and S3 access keys.



These are the primary steps in the workflow:

1

Sign in to tenant

Sign in to the tenant account on the source grid (the grid where the tenant was initially created.)

2

Optionally, configure identity federation

If your tenant account has the **Use own identity source** permission to use federated groups and users, configure the same identity source (with the same settings) for both the source and destination tenant accounts. Federated groups and users can't be cloned unless both grids are using the same identity source. For instructions, see [Use identity federation](#).

3

Create groups and users

When creating groups and users, always start from the tenant's source grid. When you add a new group, StorageGRID automatically clones it to the destination grid.

- If identity federation is configured for the entire StorageGRID system or for your tenant account, [create new tenant groups](#) by importing federated groups from the identity source.
- If you aren't using identity federation, [create new local groups](#) and then [create local users](#).

4

Create S3 access keys

You can [create your own access keys](#) or to [create another user's access keys](#) on either the source grid or the destination grid to access buckets on that grid.

5

Optionally, clone S3 access keys

If you need to access buckets with the same access keys on both grids, create the access keys on the source grid and then use the Tenant Manager API to manually clone them to the destination grid. For instructions, see [Clone S3 access keys using the API](#).


How are groups, users, and S3 access keys cloned?

Review this section to understand how groups, users, and S3 access keys are cloned between the tenant source grid and the tenant destination grid.

Local groups created on source grid are cloned

After a tenant account is created and replicated to the destination grid, StorageGRID automatically clones any local groups you add to the tenant's source grid to the tenant's destination grid.

Both the original group and its clone have the same access mode, group permissions, and S3 group policy. For instructions, see [Create groups for S3 tenant](#).




Any users you select when you create a local group on the source grid aren't included when the group is cloned to the destination grid. For this reason, don't select users when you create the group. Instead, select the group when you create the users.

Local users created on source grid are cloned

When you create a new local user on the source grid, StorageGRID automatically clones that user to the destination grid. Both the original user and its clone have the same full name, username, and **Deny access** setting. Both users also belong to the same groups. For instructions, see [Manage local users](#).

For security reasons, local user passwords aren't cloned to the destination grid. If a local user needs to access Tenant Manager on the destination grid, the root user for the tenant account must add a password for that user on the destination grid. For instructions, see [Manage local users](#).




Federated groups created on source grid are cloned

Assuming the requirements for using account clone with [single sign-on](#) and [identity federation](#) have been met, federated groups that you create (import) for the tenant on the source grid are automatically cloned to the tenant on the destination grid.

Both groups have the same access mode, group permissions and S3 group policy.

After federated groups are created for the source tenant and cloned to the destination tenant, federated users can sign in to the tenant on either grid.




S3 access keys can be manually cloned

StorageGRID does not automatically clone S3 access keys because security is improved by having different keys on each grid.

To manage access keys on the two grids, you can do either of the following:


- If you don't need to use the same keys for each grid, you can [create your own access keys](#) or [create another user's access keys](#) on each grid.
- If you need to use the same keys on both grids, you can create keys on the source grid and then use the Tenant Manager API to manually [clone the keys](#) to the destination grid.



When you clone S3 access keys for a federated user, both the user and the S3 access keys are cloned to the destination tenant.

Groups and users added to destination grid aren't cloned


Cloning occurs only from the tenant's source grid to the tenant's destination grid. If you create or import groups and users on the tenant's destination grid, StorageGRID will not clone these items back the tenant's source grid.



Edited or deleted groups, users, and access keys aren't cloned

Cloning occurs only when you create new groups and users.

If you edit or delete groups, users, or access keys on either grid, your changes will not be cloned to the other grid.



Clone S3 access keys using the API

If your tenant account has the **Use grid federation connection** permission, you can use the Tenant Management API to manually clone S3 access keys from the tenant on the source grid to the tenant on the destination grid.

Before you begin

- The tenant account has the **Use grid federation connection** permission.
- The grid federation connection has a **Connection status** of **Connected**.
- You are signed in to the Tenant Manager on the tenant's source grid using a [supported web browser](#).
- You belong to a user group that has the [Manage your own S3 credentials or Root access permission](#).
- If you are cloning access keys for a local user, the user already exists on both grids.



When you clone S3 access keys for a federated user, both the user and the S3 access keys are added to the destination tenant.

Clone your own access keys

You can clone your own access keys if you need to access the same buckets on both grids.

Steps

1. Using the Tenant Manager on the source grid, [create your own access keys](#) and download the `.csv` file.
2. From the top of the Tenant Manager, select the help icon and select **API documentation**.
3. In the **s3** section, select the following endpoint:

```
POST /org/users/current-user/replicate-s3-access-key
```

POST

`/org/users/current-user/replicate-s3-access-key` Clone the current user's S3 key to the other grids.



4. Select **Try it out**.
5. In the **body** text box, replace the example entries for **accessKey** and **secretAccessKey** with the values from the `.csv` file you downloaded.

Be sure to retain the double quotes around each string.

body * required

(body)

Edit Value | Model

```
{
  "accessKey": "AKIAIOSFODNN7EXAMPLE",
  "secretAccessKey": "wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY",
  "expires": "2028-09-04T00:00:00.000Z"
}
```

- If the key will expire, replace the example entry for **expires** with the expiration date and time as a string in ISO 8601 data-time format (for example, 2024-02-28T22:46:33-08:00). If the key will not expire, enter **null** as the value for the **expires** entry (or remove the **Expires** line and the preceding comma).
- Select **Execute**.
- Confirm that the server response code is **204**, indicating that the key was successfully cloned to the destination grid.

Clone another user's access keys

You can clone another user's access keys if they need to access the same buckets on both grids.

Steps

- Using the Tenant Manager on the source grid, [create the other user's S3 access keys](#) and download the `.csv` file.
- From the top of the Tenant Manager, select the help icon and select **API documentation**.
- Obtain the user ID. You will need this value to clone the other user's access keys.
 - From the **users** section, select the following endpoint:
- In the **s3** section, select the following endpoint:

```
GET /org/users
```

- Select **Try it out**.
- Specify any parameters you want to use when looking up users.
- Select **Execute**.
- Find the user whose keys you want to clone, and copy the number in the **id** field.

```
POST /org/users/{userId}/replicate-s3-access-key
```

POST /org/users/{userId}/replicate-s3-access-key Clone an S3 key to the other grids. 

- Select **Try it out**.
- In the **userId** text box, paste the user ID you copied.
- In the **body** text box, replace the example entries for **example access key** and **secret access key** with the values from the `.csv` file for that user.

Be sure to retain the double quotes around the string.

- If the key will expire, replace the example entry for **expires** with the expiration date and time as a string in

ISO 8601 data-time format (for example, 2023-02-28T22:46:33-08:00). If the key will not expire, enter **null** as the value for the **expires** entry (or remove the **Expires** line and the preceding comma).


9. Select **Execute**.
10. Confirm that the server response code is **204**, indicating that the key was successfully cloned to the destination grid.

Manage cross-grid replication

If your tenant account was assigned the **Use grid federation connection** permission when it was created, you can use cross-grid replication to automatically replicate objects between buckets on the tenant's source grid and buckets on the tenant's destination grid. Cross-grid replication can occur in one or both directions.

Workflow for cross-grid replication

The workflow diagram summarize the steps you will perform to configure cross-grid replication between buckets on two grids. These steps are described in more detail below.



Configure cross-grid replication

Before you can use cross-grid replication, you must sign in to the corresponding tenant accounts on each grid and create identical buckets. Then, you can enable cross-grid replication on either or both buckets.

Before you begin

- You have reviewed the requirements for cross-grid replication. See [What is cross-grid replication](#).
- You are using a [supported web browser](#).
- The tenant account has the **Use grid federation connection** permission, and identical tenant accounts exist on both grids. See [Manage the permitted tenants for grid federation connection](#).
- The tenant user you will be signing in as already exists on both grids and belongs to a user group that has the [Root access permission](#).
- If you will be signing in to the tenant's destination grid as a local user, the root user for the tenant account has set a password for your user account on that grid.

Create two identical buckets

As a first step, sign in to the corresponding tenant accounts on each grid and create identical buckets.

Steps

1. Starting from either grid in the grid federation connection, create a new bucket:
 - a. Sign in to the tenant account using the credentials of a tenant user who exists on both grids.
- 
- If you are unable to sign in to the tenant's destination grid as a local user, confirm that the root user for the tenant account has set a password for your user account.
- b. Follow the instructions to [create an S3 bucket](#).
 - c. On the **Manage object settings** tab, select **Enable object versioning**.
 - d. If S3 Object Lock is enabled for your StorageGRID system, don't enable S3 Object Lock for the bucket.
 - e. Select **Create bucket**.
 - f. Select **Finish**.
 2. Repeat these steps to create an identical bucket for the same tenant account on the other grid in the grid federation connection.




As required, each bucket can use a different region.

Enable cross-grid replication

You must perform these steps before adding any objects to either bucket.

Steps

1. Starting from a grid whose objects you want to replicate, enable [cross-grid replication in one direction](#):
 - a. Sign in to the tenant account for the bucket.
 - b. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
 - c. Select the bucket name from the table to access the bucket details page.
 - d. Select the **Cross-grid replication** tab.
 - e. Select **Enable**, and review the list of requirements.
 - f. If all requirements have been met, select the grid federation connection you want to use.
 - g. Optionally, change the setting of **Replicate delete markers** to determine what happens on the destination grid if an S3 client issues a delete request to the source grid that doesn't include a version ID:
 - **Yes** (default): A delete marker is added to the source bucket and replicated to the destination bucket.
 - **No**: A delete marker is added to the source bucket but is not replicated to the destination bucket.



If the delete request includes a version ID, that object version is permanently removed from the source bucket. StorageGRID does not replicate delete requests that include a version ID, so the same object version is not deleted from the destination.

See [What is cross-grid replication](#) for details.

- h. Optionally, change the setting of the **Cross-grid replication** audit category to manage the volume of audit messages:
 - **Error** (default): Only failed cross-grid replication requests are included in the audit output.
 - **Normal**: All cross-grid replication requests are included, which significantly increases the volume of the audit output.
- i. Review your selections. You aren't able to change these settings unless both buckets are empty.
- j. Select **Enable and test**.

After a few moments, a success message appears. Objects added to this bucket will now be automatically replicated to the other grid. **Cross-grid replication** is shown as an enabled feature on the bucket details page.

2. Optionally, go to the corresponding bucket on the other grid and [enable cross-grid replication in both directions](#).

Test replication between grids

If cross-grid replication is enabled for a bucket, you might need to verify that the connection and cross-grid replication are working correctly and that the source and destination buckets still meet all requirements (for example, versioning is still enabled).

Before you begin

- You are using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#).

Steps

1. Sign in to the tenant account for the bucket.

2. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
3. Select the bucket name from the table to access the bucket details page.
4. Select the **Cross-grid replication** tab.
5. Select **Test connection**.

If the connection is healthy, a success banner appears. Otherwise, an error message appears, which you and the grid admin can use to resolve the issue. For details, see [Troubleshoot grid federation errors](#).

6. If cross-grid replication is configured to occur in both directions, go to the corresponding bucket on the other grid and select **Test connection** to verify that cross-grid replication is working in the other direction.

Disable cross-grid replication

You can permanently stop cross-grid replication if you no longer want to copy objects to the other grid.

Before disabling cross-grid replication, note the following:

- Disabling cross-grid replication does not remove any objects that have already been copied between grids. For example, objects in `my-bucket` on Grid 1 that have been copied to `my-bucket` on Grid 2 aren't removed if you disable cross-grid replication for that bucket. If you want to delete these objects, you must remove them manually.
- If cross-grid replication was enabled for each of the buckets (that is, if replication occurs in both directions), you can disable cross-grid replication for either or both buckets. For example, you might want to disable replicating objects from `my-bucket` on Grid 1 to `my-bucket` on Grid 2, while continuing to replicate objects from `my-bucket` on Grid 2 to `my-bucket` on Grid 1.
- You must disable cross-grid replication before you can remove a tenant's permission to use the grid federation connection. See [Manage permitted tenants](#).
- If you disable cross-grid replication for a bucket that contains objects, you will not be able to reenabling cross-grid replication unless you delete all objects from both the source and destination buckets.



You can't reenabling replication unless both buckets are empty.

Before you begin

- You are using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#).

Steps

1. Starting from the grid whose objects you no longer want to replicate, stop cross-grid replication for the bucket:
 - a. Sign in to the tenant account for the bucket.
 - b. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
 - c. Select the bucket name from the table to access the bucket details page.
 - d. Select the **Cross-grid replication** tab.
 - e. Select **Disable replication**.
 - f. If you are sure you want to disable cross-grid replication for this bucket, type **Yes** in the text box, and select **Disable**.

After a few moments, a success message appears. New objects added to this bucket can no longer be automatically replicated to the other grid. **Cross-grid replication** is no longer shown as a Enabled feature on the Buckets page.

2. If cross-grid replication was configured to occur in both directions, go to the corresponding bucket on the other grid and stop cross-grid replication in the other direction.

View grid federation connections

If your tenant account has the **Use grid federation connection** permission, you can view the allowed connections.

Before you begin

- The tenant account has the **Use grid federation connection** permission.
- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#).

Steps

1. Select **STORAGE (S3) > Grid federation connections**.

The Grid federation connection page appears and includes a table that summarizes the following information:

Column	Description
Connection name	The grid federation connections this tenant has permission to use.
Buckets with cross-grid replication	For each grid federation connection, the tenant buckets that have cross-grid replication enabled. Objects added to these buckets will be replicated to the other grid in the connection.
Last error	For each grid federation connection, the most recent error to occur, if any, when data was being replicated to the other grid. See Clear the last error .

2. Optionally, select a bucket name to [view bucket details](#).

Clear the last error

An error might appear in the **Last error** column for one of these reasons:

- The source object version was not found.
- The source bucket was not found.
- The destination bucket was deleted.
- The destination bucket was re-created by a different account.
- The destination bucket has versioning suspended.
- The destination bucket was re-created by the same account but is now unversioned.



This column only shows the last cross-grid replication error to occur; previous errors that might have occurred will not be shown.

Steps

1. If a message appears in the **Last error** column, view the message text.

For example, this error indicates that the destination bucket for cross-grid replication was in an invalid state, possibly because versioning was suspended or S3 Object Lock was enabled.

Grid federation connections

Clear error

Search...

Q

Displaying one result

Connection name	Buckets with cross-grid replication	Last error
Grid 1-Grid 2	my-cgr-bucket	2022-12-07 16:02:20 MST Cross-grid replication has encountered an error. Failed to send cross-grid replication request from source bucket 'my-cgr-bucket' to destination bucket 'my-cgr-bucket'. Error code: DestinationRequestError. Detail: InvalidBucketState. Confirm that the source and destination buckets have object versioning enabled and S3 Object Lock disabled. (logID 4791585492825418592)

2. Perform any recommended actions. For example, if versioning was suspended on the destination bucket for cross-grid replication, reenable versioning for that bucket.
3. Select the connection from the table.
4. Select **Clear error**.
5. Select **Yes** to clear the message and update the system's status.
6. Wait 5-6 minutes and then ingest a new object into the bucket. Confirm that the error message does not reappear.



To ensure the error message is cleared, wait at least 5 minutes after the timestamp in the message before ingesting a new object.

7. To determine if any objects failed to be replicated because of the bucket error, see [Identify and retry failed replication operations](#).

Manage groups and users

Use identity federation

Using identity federation makes setting up tenant groups and users faster, and it allows tenant users to sign in to the tenant account using familiar credentials.

Configure identity federation for Tenant Manager

You can configure identity federation for the Tenant Manager if you want tenant groups and users to be managed in another system such as Active Directory, Azure Active Directory (Azure AD), OpenLDAP, or Oracle Directory Server.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#).
- You are using Active Directory, Azure AD, OpenLDAP, or Oracle Directory Server as the identity provider.



If you want to use an LDAP v3 service that is not listed, contact technical support.

- If you plan to use OpenLDAP, you must configure the OpenLDAP server. See [Guidelines for configuring OpenLDAP server](#).
- If you plan to use Transport Layer Security (TLS) for communications with the LDAP server, the identity provider must be using TLS 1.2 or 1.3. See [Supported ciphers for outgoing TLS connections](#).

About this task

Whether you can configure an identity federation service for your tenant depends on how your tenant account was set up. Your tenant might share the identity federation service that was configured for the Grid Manager. If you see this message when you access the Identity Federation page, you can't configure a separate federated identity source for this tenant.



This tenant account uses the LDAP server that is configured for the Grid Manager.
Contact the grid administrator for information or to change this setting.

Enter configuration

When you configure identity federation, you provide the values StorageGRID needs to connect to an LDAP service.

Steps

1. Select **ACCESS MANAGEMENT > Identity federation**.
2. Select **Enable identity federation**.
3. In the LDAP service type section, select the type of LDAP service you want to configure.

LDAP service type

Select the type of LDAP service you want to configure.

Active Directory	Azure	OpenLDAP	Other

Select **Other** to configure values for an LDAP server that uses Oracle Directory Server.

4. If you selected **Other**, complete the fields in the LDAP Attributes section. Otherwise, go to the next step.
 - **User Unique Name:** The name of the attribute that contains the unique identifier of an LDAP user. This attribute is equivalent to `sAMAccountName` for Active Directory and `uid` for OpenLDAP. If you are configuring Oracle Directory Server, enter `uid`.
 - **User UUID:** The name of the attribute that contains the permanent unique identifier of an LDAP user. This attribute is equivalent to `objectGUID` for Active Directory and `entryUUID` for OpenLDAP. If you are configuring Oracle Directory Server, enter `nsuniqueid`. Each user's value for the specified

attribute must be a 32-digit hexadecimal number in either 16-byte or string format, where hyphens are ignored.

- **Group Unique Name:** The name of the attribute that contains the unique identifier of an LDAP group. This attribute is equivalent to `sAMAccountName` for Active Directory and `cn` for OpenLDAP. If you are configuring Oracle Directory Server, enter `cn`.
- **Group UUID:** The name of the attribute that contains the permanent unique identifier of an LDAP group. This attribute is equivalent to `objectGUID` for Active Directory and `entryUUID` for OpenLDAP. If you are configuring Oracle Directory Server, enter `nsuniqueid`. Each group's value for the specified attribute must be a 32-digit hexadecimal number in either 16-byte or string format, where hyphens are ignored.

5. For all LDAP service types, enter the required LDAP server and network connection information in the Configure LDAP server section.

- **Hostname:** The fully qualified domain name (FQDN) or IP address of the LDAP server.
- **Port:** The port used to connect to the LDAP server.



The default port for STARTTLS is 389, and the default port for LDAPS is 636. However, you can use any port as long as your firewall is configured correctly.

- **Username:** The full path of the distinguished name (DN) for the user that will connect to the LDAP server.

For Active Directory, you can also specify the Down-Level Logon Name or the User Principal Name.

The specified user must have permission to list groups and users and to access the following attributes:

- `sAMAccountName` or `uid`
 - `objectGUID`, `entryUUID`, or `nsuniqueid`
 - `cn`
 - `memberOf` or `isMemberOf`
 - **Active Directory:** `objectSid`, `primaryGroupID`, `userAccountControl`, and `userPrincipalName`
 - **Azure:** `accountEnabled` and `userPrincipalName`
- **Password:** The password associated with the username.



If you change the password in the future, you must update it on this page.

- **Group Base DN:** The full path of the distinguished name (DN) for an LDAP subtree you want to search for groups. In the Active Directory example (below), all groups whose Distinguished Name is relative to the base DN (`DC=storagegrid,DC=example,DC=com`) can be used as federated groups.



The **Group unique name** values must be unique within the **Group Base DN** they belong to.

- **User Base DN:** The full path of the distinguished name (DN) of an LDAP subtree you want to search for users.



The **User unique name** values must be unique within the **User Base DN** they belong to.

- **Bind username format** (optional): The default username pattern StorageGRID should use if the pattern can't be determined automatically.

Providing **Bind username format** is recommended because it can allow users to sign in if StorageGRID is unable to bind with the service account.

Enter one of these patterns:

- **UserPrincipalName pattern (Active Directory and Azure):** `[USERNAME]@example.com`
- **Down-level logon name pattern (Active Directory and Azure):** `example\[USERNAME]`
- **Distinguished name pattern:** `CN=[USERNAME],CN=Users,DC=example,DC=com`

Include **[USERNAME]** exactly as written.

6. In the Transport Layer Security (TLS) section, select a security setting.

- **Use STARTTLS:** Use STARTTLS to secure communications with the LDAP server. This is the recommended option for Active Directory, OpenLDAP, or Other, but this option is not supported for Azure.
- **Use LDAPS:** The LDAPS (LDAP over SSL) option uses TLS to establish a connection to the LDAP server. You must select this option for Azure.
- **Do not use TLS:** The network traffic between the StorageGRID system and the LDAP server will not be secured. This option is not supported for Azure.



Using the **Do not use TLS** option is not supported if your Active Directory server enforces LDAP signing. You must use STARTTLS or LDAPS.

7. If you selected STARTTLS or LDAPS, choose the certificate used to secure the connection.

- **Use operating system CA certificate:** Use the default Grid CA certificate installed on the operating system to secure connections.
- **Use custom CA certificate:** Use a custom security certificate.

If you select this setting, copy and paste the custom security certificate into the CA certificate text box.

Test the connection and save the configuration

After entering all values, you must test the connection before you can save the configuration. StorageGRID verifies the connection settings for the LDAP server and the bind username format, if you provided one.

Steps

1. Select **Test connection**.
2. If you did not provide a bind username format:
 - A "Test connection successful" message appears if the connection settings are valid. Select **Save** to save the configuration.
 - A "test connection could not be established" message appears if the connection settings are invalid. Select **Close**. Then, resolve any issues and test the connection again.

3. If you provided a bind username format, enter the username and password of a valid federated user.

For example, enter your own username and password. Don't include any special characters in the username, such as @ or /.

Test Connection ✕

To test the connection and the bind username format, enter the username and password of a federated user. For example, enter your own federated username and password. The test values are not saved.

Test username

The username of a federated user.

Test password

 👁

[Cancel](#) [Test Connection](#)

- A "Test connection successful" message appears if the connection settings are valid. Select **Save** to save the configuration.
- An error message appears if the connection settings, bind username format, or test username and password are invalid. Resolve any issues and test the connection again.

Force synchronization with identity source

The StorageGRID system periodically synchronizes federated groups and users from the identity source. You can force synchronization to start if you want to enable or restrict user permissions as quickly as possible.

Steps

1. Go to the Identity federation page.
2. Select **Sync server** at the top of the page.

The synchronization process might take some time depending on your environment.



The **Identity federation synchronization failure** alert is triggered if there is an issue synchronizing federated groups and users from the identity source.

Disable identity federation

You can temporarily or permanently disable identity federation for groups and users. When identity federation is disabled, there is no communication between StorageGRID and the identity source. However, any settings you have configured are retained, allowing you to easily reenable identity federation in the future.

About this task

Before you disable identity federation, you should be aware of the following:

- Federated users will be unable to sign in.

- Federated users who are currently signed in will retain access to the StorageGRID system until their session expires, but they will be unable to sign in after their session expires.
- Synchronization between the StorageGRID system and the identity source will not occur, and alerts will not be raised for accounts that have not been synchronized.
- The **Enable identity federation** checkbox is disabled if single sign-on (SSO) is set to **Enabled** or **Sandbox Mode**. The SSO Status on the Single Sign-on page must be **Disabled** before you can disable identity federation. See [Disable single sign-on](#).

Steps

1. Go to the Identity federation page.
2. Uncheck the **Enable identity federation** checkbox.

Guidelines for configuring OpenLDAP server

If you want to use an OpenLDAP server for identity federation, you must configure specific settings on the OpenLDAP server.



For identity sources that aren't ActiveDirectory or Azure, StorageGRID will not automatically block S3 access to users who are disabled externally. To block S3 access, delete any S3 keys for the user or remove the user from all groups.

Memberof and refint overlays

The memberof and refint overlays should be enabled. For more information, see the instructions for reverse group membership maintenance in the [OpenLDAP documentation: Version 2.4 Administrator's Guide](#).

Indexing

You must configure the following OpenLDAP attributes with the specified index keywords:

- `olcDbIndex: objectClass eq`
- `olcDbIndex: uid eq,pres,sub`
- `olcDbIndex: cn eq,pres,sub`
- `olcDbIndex: entryUUID eq`

In addition, ensure the fields mentioned in the help for Username are indexed for optimal performance.

See the information about reverse group membership maintenance in the [OpenLDAP documentation: Version 2.4 Administrator's Guide](#).

Manage tenant groups

Create groups for an S3 tenant

You can manage permissions for S3 user groups by importing federated groups or creating local groups.

Before you begin


- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#).
- If you plan to import a federated group, you have [configured identity federation](#), and the federated group already exists in the configured identity source.
- If your tenant account has the **Use grid federation connection** permission, you have reviewed the workflow and considerations for [cloning tenant groups and users](#), and you are signed in to the tenant's source grid.

Access the Create group wizard

As your first step, access the Create group wizard.

Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. If your tenant account has the **Use grid federation connection** permission, confirm that a blue banner appears, indicating that new groups created on this grid will be cloned to the same tenant on the other grid in the connection. If this banner does not appear, you might be signed in to the tenant's destination grid.



3. Select **Create group**.

Choose a group type

You can create a local group or import a federated group.

Steps

1. Select the **Local group** tab to create a local group, or select the **Federated group** tab to import a group from the previously configured identity source.

If single sign-on (SSO) is enabled for your StorageGRID system, users belonging to local groups will not be able to sign in to the Tenant Manager, although they can use client applications to manage the tenant's resources, based on group permissions.

2. Enter the group's name.
 - **Local group:** Enter both a display name and a unique name. You can edit the display name later.



If your tenant account has the **Use grid federation connection** permission, a cloning error will occur if the same **Unique name** already exists for the tenant on the destination grid.

- **Federated group:** Enter the unique name. For Active Directory, the unique name is the name associated with the `sAMAccountName` attribute. For OpenLDAP, the unique name is the name associated with the `uid` attribute.

3. Select **Continue**.

Manage group permissions

Group permissions control which tasks users can perform in the Tenant Manager and Tenant Management API.

Steps

1. For **Access mode**, select one of the following:

- **Read-write** (default): Users can sign in to Tenant Manager and manage the tenant configuration.
- **Read-only:** Users can only view settings and features. They can't make any changes or perform any operations in the Tenant Manager or Tenant Management API. Local read-only users can change their own passwords.



If a user belongs to multiple groups and any group is set to Read-only, the user will have read-only access to all selected settings and features.

2. Select one or more permissions for this group.

See [Tenant management permissions](#).

3. Select **Continue**.

Set S3 group policy

The group policy determines which S3 access permissions users will have.

Steps

1. Select the policy you want to use for this group.

Group policy	Description
No S3 Access	Default. Users in this group don't have access to S3 resources, unless access is granted with a bucket policy. If you select this option, only the root user will have access to S3 resources by default.
Read Only Access	Users in this group have read-only access to S3 resources. For example, users in this group can list objects and read object data, metadata, and tags. When you select this option, the JSON string for a read-only group policy appears in the text box. You can't edit this string.

Group policy	Description
Full Access	Users in this group have full access to S3 resources, including buckets. When you select this option, the JSON string for a full-access group policy appears in the text box. You can't edit this string.
Ransomware Mitigation	This example policy applies to all buckets for this tenant. Users in this group can perform common actions, but can't permanently delete objects from buckets that have object versioning enabled. Tenant Manager users who have the Manage all buckets permission can override this group policy. Limit the Manage all buckets permission to trusted users, and use Multi-Factor Authentication (MFA) where available.
Custom	Users in the group are granted the permissions you specify in the text box.

- If you selected **Custom**, enter the group policy. Each group policy has a size limit of 5,120 bytes. You must enter a valid JSON formatted string.

For detailed information about group policies, including language syntax and examples, see [Example group policies](#).

- If you are creating a local group, select **Continue**. If you are creating a federated group, select **Create group** and **Finish**.

Add users (local groups only)

You can save the group without adding users, or you can optionally add any local users that already exist.



If your tenant account has the **Use grid federation connection** permission, any users you select when you create a local group on the source grid aren't included when the group is cloned to the destination grid. For this reason, don't select users when you create the group. Instead, select the group when you create the users.

Steps

- Optionally, select one or more local users for this group.
- Select **Create group** and **Finish**.

The group you created appears in the list of groups.

If your tenant account has the **Use grid federation connection** permission and you are on the tenant's source grid, the new group is cloned to the tenant's destination grid. **Success** appears as the **Cloning status** in the Overview section of the group's detail page.

Create groups for a Swift tenant

You can manage access permissions for a Swift tenant account by importing federated groups or creating local groups. At least one group must have the Swift Administrator permission, which is required to manage the containers and objects for a Swift tenant

account.



Support for Swift client applications has been deprecated and will be removed in a future release.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#).
- If you plan to import a federated group, you have [configured identity federation](#), and the federated group already exists in the configured identity source.

Access the Create group wizard

Steps

As your first step, access the Create group wizard.

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select **Create group**.

Choose a group type

You can create a local group or import a federated group.

Steps

1. Select the **Local group** tab to create a local group, or select the **Federated group** tab to import a group from the previously configured identity source.

If single sign-on (SSO) is enabled for your StorageGRID system, users belonging to local groups will not be able to sign in to the Tenant Manager, although they can use client applications to manage the tenant's resources, based on group permissions.

2. Enter the group's name.
 - **Local group**: Enter both a display name and a unique name. You can edit the display name later.
 - **Federated group**: Enter the unique name. For Active Directory, the unique name is the name associated with the `sAMAccountName` attribute. For OpenLDAP, the unique name is the name associated with the `uid` attribute.
3. Select **Continue**.

Manage group permissions

Group permissions control which tasks users can perform in the Tenant Manager and Tenant Management API.

Steps

1. For **Access mode**, select one of the following:
 - **Read-write** (default): Users can sign in to Tenant Manager and manage the tenant configuration.
 - **Read-only**: Users can only view settings and features. They can't make any changes or perform any operations in the Tenant Manager or Tenant Management API. Local read-only users can change their own passwords.



If a user belongs to multiple groups and any group is set to Read-only, the user will have read-only access to all selected settings and features.

2. Select the **Root access** checkbox if group users need to sign in to the Tenant Manager or Tenant Management API.
3. Select **Continue**.

Set Swift group policy

Swift users need administrator permission to authenticate into the Swift REST API to create containers and ingest objects.

1. Select the **Swift administrator** checkbox if group users need to use the Swift REST API to manage containers and objects.
2. If you are creating a local group, select **Continue**. If you are creating a federated group, select **Create group** and **Finish**.

Add users (local groups only)

You can save the group without adding users, or you can optionally add any local users that already exist.

Steps

1. Optionally, select one or more local users for this group.

If you have not yet created local users, you can add this group to the user on the Users page. See [Manage local users](#).

2. Select **Create group** and **Finish**.

The group you created appears in the list of groups.

Tenant management permissions

Before you create a tenant group, consider which permissions you want to assign to that group. Tenant management permissions determine which tasks users can perform using the Tenant Manager or the Tenant Management API. A user can belong to one or more groups. Permissions are cumulative if a user belongs to multiple groups.

To sign in to the Tenant Manager or to use the Tenant Management API, users must belong to a group that has at least one permission. All users who can sign in can perform the following tasks:

- View the dashboard
- Change their own password (for local users)

For all permissions, the group's Access mode setting determines whether users can change settings and perform operations or whether they can only view the related settings and features.



If a user belongs to multiple groups and any group is set to Read-only, the user will have read-only access to all selected settings and features.

You can assign the following permissions to a group. Note that S3 tenants and Swift tenants have different

group permissions.

Permission	Description	Details
Root access	Provides full access to the Tenant Manager and the Tenant Management API.	Swift users must have Root access permission to sign in to the tenant account.
Administrator	Swift tenants only. Provides full access to the Swift containers and objects for this tenant account	Swift users must have the Swift Administrator permission to perform any operations with the Swift REST API.
Manage your own S3 credentials	Allows users to create and remove their own S3 access keys.	Users who don't have this permission don't see the STORAGE (S3) > My S3 access keys menu option.
View all buckets	S3 tenants: Allows users to view all buckets and bucket configurations. Swift tenants: Allows Swift users to view all containers and container configurations using the Tenant Management API.	Users who don't have either the View all buckets or the Manage all buckets permission don't see the Buckets menu option. This permission is superseded by the Manage all buckets permission. It does not affect S3 bucket or group policies used by S3 clients or S3 Console. You can only assign this permission to Swift groups from the Tenant Management API. You can't assign this permission to Swift groups using the Tenant Manager.
Manage all buckets	S3 tenants: Allows users to use the Tenant Manager and the Tenant Management API to create and delete S3 buckets and to manage the settings for all S3 buckets in the tenant account, regardless of S3 bucket or group policies. Swift tenants: Allows Swift users to control the consistency for Swift containers using the Tenant Management API.	Users who don't have either the View all buckets or the Manage all buckets permission don't see the Buckets menu option. This permission supersedes the View all buckets permission. It does not affect S3 bucket or group policies used by S3 clients or S3 Console. You can only assign this permission to Swift groups from the Tenant Management API. You can't assign this permission to Swift groups using the Tenant Manager.
Manage endpoints	Allows users to use the Tenant Manager or the Tenant Management API to create or edit platform service endpoints, which are used as the destination for StorageGRID platform services.	Users who don't have this permission don't see the Platform services endpoints menu option.

Permission	Description	Details
Use S3 Console tab	When combined with the View all buckets or Manage all buckets permission, allows users to view and manage objects from the S3 Console tab on the details page for a bucket.	

Manage groups

Manage your tenant groups as needed to view, edit, or duplicate a group, and more.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#).

View or edit group

You can view and edit the basic information and details for each group.


Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Review the information provided on the Groups page, which lists basic information for all local and federated groups for this tenant account.

If the tenant account has the **Use grid federation connection** permission and you are viewing groups on the tenant's source grid:

- A banner message indicates that if you edit or remove a group, your changes will not be synced to the other grid.
 - As needed, a banner message indicates if groups were not cloned to the tenant on the destination grid. You can [retry a group clone](#) that failed.
3. If you want to change the group's name:
 - a. Select the checkbox for the group.
 - b. Select **Actions > Edit group name**.
 - c. Enter the new name.
 - d. Select **Save changes**.
 4. If you want to view more details or make additional edits, do either of the following:
 - Select the group name.
 - Select the checkbox for the group, and select **Actions > View group details**.
 5. Review the Overview section, which shows the following information for each group:
 - Display name
 - Unique name
 - Type
 - Access mode
 - Permissions

- S3 Policy
 - Number of users in this group
 - Additional fields if the tenant account has the **Use grid federation connection** permission and you are viewing the group on the tenant's source grid:
 - Cloning status, either **Success** or **Failure**
 - A blue banner indicating that if you edit or delete this group, your changes will not be synced to the other grid.
6. Edit group settings as needed. See [Create groups for an S3 tenant](#) and [Create groups for a Swift tenant](#) for details about what to enter.
- a. In the Overview section, change the display name by selecting the name or the edit icon .
 - b. On the **Group permissions** tab, update the permissions, and select **Save changes**.
 - c. On the **Group policy** tab, make any changes, and select **Save changes**.
 - If you are editing an S3 group, optionally select a different S3 group policy or enter the JSON string for a custom policy, as required.
 - If you are editing a Swift group, optionally select or clear the **Swift Administrator** checkbox.
7. To add one or more existing local users to the group:
- a. Select the Users tab.



Manage users

You can add users to this group or remove users from this group.

Add users **Remove Users** Search Groups...  Displaying 1 results

Username	Full Name	Denied
User_02	User_02_Managers	

- b. Select **Add users**.
 - c. Select the existing users you want to add, and select **Add users**.
- A success message appears in the upper right.
8. To remove local users from the group:
- a. Select the Users tab.
 - b. Select **Remove users**.
 - c. Select the users you want to remove, and select **Remove users**.
- A success message appears in the upper right.
9. Confirm that you selected **Save changes** for each section you changed.

Duplicate group

You can duplicate an existing group to create new groups more quickly.



If your tenant account has the **Use grid federation connection** permission and you duplicate a group from the tenant's source grid, the duplicated group will be cloned to the tenant's destination grid.

Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select the checkbox for the group you want to duplicate.
3. Select **Actions > Duplicate group**.
4. See [Create groups for an S3 tenant](#) or [Create groups for a Swift tenant](#) for details about what to enter.
5. Select **Create group**.

Retry group clone

To retry a clone that failed:

1. Select each group that indicates *(Cloning failed)* below the group name.
2. Select **Actions > Clone groups**.
3. View the status of the clone operation from the details page of each group you're cloning.

For additional information, see [Clone tenant groups and users](#).

Delete one or more groups

You can delete one or more groups. Any users who belong only to a group that is deleted will no longer be able to sign in to the Tenant Manager or use the tenant account.



If your tenant account has the **Use grid federation connection** permission and you delete a group, StorageGRID will not delete the corresponding group on the other grid. If you need to keep this information in sync, you must delete the same group from both grids.

Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select the checkbox for each group you want to delete.
3. Select **Actions > Delete group** or **Actions > Delete groups**.

A confirmation dialog box appears.

4. Select **Delete group** or **Delete groups**.

Manage local users

You can create local users and assign them to local groups to determine which features these users can access. The Tenant Manager includes one predefined local user, named "root." Although you can add and remove local users, you can't remove the root user.



If single sign-on (SSO) is enabled for your StorageGRID system, local users will not be able to sign in to the Tenant Manager or the Tenant Management API, although they can use client applications to access the tenant's resources, based on group permissions.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#).
- If your tenant account has the **Use grid federation connection** permission, you have reviewed the workflow and considerations for [cloning tenant groups and users](#), and you are signed in to the tenant's source grid.

Create a local user

You can create a local user and assign them to one or more local groups to control their access permissions.

S3 users who don't belong to any groups don't have management permissions or S3 group policies applied to them. These users might have S3 bucket access granted through a bucket policy.

Swift users who don't belong to any groups don't have management permissions or Swift container access.

Access the Create user wizard

Steps

1. Select **ACCESS MANAGEMENT > Users**.

If your tenant account has the **Use grid federation connection** permission, a blue banner indicates that this is the tenant's source grid. Any local users you create on this grid will be cloned to the other grid in the connection.

2. Select **Create user**.

Enter credentials

Steps

1. For the **Enter user credentials** step, complete the following fields.

Field	Description
Full name	The full name for this user, for example, the first name and last name of a person or the name of an application.
Username	The name this user will use to sign in. Usernames must be unique and can't be changed. Note: If your tenant account has the Use grid federation connection permission, a cloning error will occur if the same Username already exists for the tenant on the destination grid.
Password and Confirm password	The password the user will initially use when signing in.
Deny access	Select Yes to prevent this user from signing in to the tenant account, even though they might still belong to one or more groups. For example, select Yes to temporarily suspend a user's ability to sign in.

2. Select **Continue**.

Assign to groups

Steps

1. Assign the user to one or more local groups to determine which tasks they can perform.

Assigning a user to groups is optional. If you'd prefer, you can select users when you create or edit groups.

Users who don't belong to any groups will have no management permissions. Permissions are cumulative. Users will have all permissions for all groups they belong to. See [Tenant management permissions](#).

2. Select **Create user**.

If your tenant account has the **Use grid federation connection** permission and you are on the tenant's source grid, the new local user is cloned to the tenant's destination grid. **Success** appears as the **Cloning status** in the Overview section of the user's detail page.

3. Select **Finish** to return to the Users page.

View or edit local user

Steps

1. Select **ACCESS MANAGEMENT > Users**.
2. Review the information provided on the Users page, which lists basic information for all local and federated users for this tenant account.

If the tenant account has the **Use grid federation connection** permission and you are viewing the user on the tenant's source grid:

- A banner message indicates that if you edit or remove a user, your changes will not be synced to the other grid.

- As needed, a banner message indicates if users were not cloned to the tenant on the destination grid. You can [retry a user clone that failed](#).
3. If you want to change the user's full name:
 - a. Select the checkbox for the user.
 - b. Select **Actions > Edit full name**.
 - c. Enter the new name.
 - d. Select **Save changes**.
 4. If you want to view more details or make additional edits, do either of the following:
 - Select the username.
 - Select the checkbox for the user, and select **Actions > View user details**.
 5. Review the Overview section, which shows the following information for each user:
 - Full name
 - Username
 - User type
 - Denied access
 - Access mode
 - Group membership
 - Additional fields if the tenant account has the **Use grid federation connection** permission and you are viewing the user on the tenant's source grid:
 - Cloning status, either **Success** or **Failure**
 - A blue banner indicating that if you edit this user, your changes will not be synced to the other grid.
 6. Edit user settings as needed. See [Create local user](#) for details about what to enter.
 - a. In the Overview section, change the full name by selecting the name or the edit icon .

You can't change the username.

 - b. On the **Password** tab, change the user's password, and select **Save changes**.
 - c. On the **Access** tab, select **No** to allow the user to sign in or select **Yes** to prevent the user from signing in. Then, select **Save changes**.
 - d. On the **Access keys** tab, select **Create key** and follow the instructions for [creating another user's S3 access keys](#).
 - e. On the **Groups** tab, select **Edit groups** to add the user to groups or remove the user from groups. Then, select **Save changes**.
 7. Confirm that you selected **Save changes** for each section you changed.

Duplicate local user

You can duplicate a local user to create a new user more quickly.



If your tenant account has the **Use grid federation connection** permission and you duplicate a user from the tenant's source grid, the duplicated user will be cloned to the tenant's destination grid.

Steps

1. Select **ACCESS MANAGEMENT > Users**.
2. Select the checkbox for the user you want to duplicate.
3. Select **Actions > Duplicate user**.
4. See [Create local user](#) for details about what to enter.
5. Select **Create user**.

Retry user clone

To retry a clone that failed:

1. Select each user that indicates (*Cloning failed*) below the user name.
2. Select **Actions > Clone users**.
3. View the status of the clone operation from the details page of each user you're cloning.

For additional information, see [Clone tenant groups and users](#).

Delete one or more local users

You can permanently delete one or more local users who no longer need to access the StorageGRID tenant account.



If your tenant account has the **Use grid federation connection** permission and you delete a local user, StorageGRID will not delete the corresponding user on the other grid. If you need to keep this information in sync, you must delete the same user from both grids.



You must use the federated identity source to delete federated users.

Steps

1. Select **ACCESS MANAGEMENT > Users**.
2. Select the checkbox for each user you want to delete.
3. Select **Actions > Delete user** or **Actions > Delete users**.

A confirmation dialog box appears.

4. Select **Delete user** or **Delete users**.

Manage S3 access keys

Manage S3 access keys

Each user of an S3 tenant account must have an access key to store and retrieve objects in the StorageGRID system. An access key consists of an access key ID and a secret access key.

S3 access keys can be managed as follows:

- Users who have the **Manage your own S3 credentials** permission can create or remove their own S3 access keys.

- Users who have the **Root access** permission can manage the access keys for the S3 root account and all other users. Root access keys provide full access to all buckets and objects for the tenant unless explicitly disabled by a bucket policy.

StorageGRID supports Signature Version 2 and Signature Version 4 authentication. Cross-account access is not permitted unless explicitly enabled by a bucket policy.

Create your own S3 access keys

If you are using an S3 tenant and you have the appropriate permission, you can create your own S3 access keys. You must have an access key to access your buckets and objects.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Manage your own S3 credentials or Root access permission](#).

About this task

You can create one or more S3 access keys that allow you to create and manage buckets for your tenant account. After you create a new access key, update the application with your new access key ID and secret access key. For security, don't create more keys than you need, and delete the keys you aren't using. If you have only one key and it is about to expire, create a new key before the old one expires, and then delete the old one.

Each key can have a specific expiration time or no expiration. Follow these guidelines for expiration time:

- Set an expiration time for your keys to limit your access to a certain time period. Setting a short expiration time can help reduce your risk if your access key ID and secret access key are accidentally exposed. Expired keys are removed automatically.
- If the security risk in your environment is low and you don't need to periodically create new keys, you don't have to set an expiration time for your keys. If you decide later to create new keys, delete the old keys manually.



The S3 buckets and objects belonging to your account can be accessed using the access key ID and secret access key displayed for your account in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from your account, and never share them with other users.

Steps

1. Select **STORAGE (S3) > My access keys**.

The My access keys page appears and lists any existing access keys.

2. Select **Create key**.
3. Do one of the following:
 - Select **Do not set an expiration time** to create a key that will not expire. (Default)
 - Select **Set an expiration time**, and set the expiration date and time.



The expiration date can be a maximum of five years from the current date. The expiration time can be a minimum of one minute from the current time.

4. Select **Create access key**.

The Download access key dialog box appears, listing your access key ID and secret access key.

5. Copy the access key ID and the secret access key to a safe location, or select **Download .csv** to save a spreadsheet file containing the access key ID and secret access key.



Don't close this dialog box until you have copied or downloaded this information. You can't copy or download keys after the dialog box has been closed.

6. Select **Finish**.

The new key is listed on the My access keys page.

7. If your tenant account has the **Use grid federation connection** permission, optionally use the Tenant Management API to manually clone S3 access keys from the tenant on the source grid to the tenant on the destination grid. See [Clone S3 access keys using the API](#).

View your S3 access keys

If you are using an S3 tenant and you have the [appropriate permission](#), you can view a list of your S3 access keys. You can sort the list by expiration time, so you can determine which keys will expire soon. As needed, you can [create new keys](#) or [delete keys](#) that you are no longer using.



The S3 buckets and objects belonging to your account can be accessed using the access key ID and secret access key displayed for your account in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from your account, and never share them with other users.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the Manage your own S3 credentials [permission](#).

Steps

1. Select **STORAGE (S3) > My access keys**.
2. From the My access keys page, sort any existing access keys by **Expiration time** or **Access key ID**.
3. As needed, create new keys or delete any keys that you are no longer using.

If you create new keys before the existing keys expire, you can begin using the new keys without temporarily losing access to the objects in the account.

Expired keys are removed automatically.

Delete your own S3 access keys

If you are using an S3 tenant and you have the appropriate permission, you can delete your own S3 access keys. After an access key is deleted, it can no longer be used to access the objects and buckets in the tenant account.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You have the [Manage your own S3 credentials permission](#).



The S3 buckets and objects belonging to your account can be accessed using the access key ID and secret access key displayed for your account in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from your account, and never share them with other users.

Steps

1. Select **STORAGE (S3) > My access keys**.
2. From the My access keys page, select the checkbox for each access key you want to remove.
3. Select **Delete key**.
4. From the confirmation dialog box, select **Delete key**.

A confirmation message appears in the upper right corner of the page.

Create another user's S3 access keys

If you are using an S3 tenant and you have the appropriate permission, you can create S3 access keys for other users, such as applications that need access to buckets and objects.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#).

About this task

You can create one or more S3 access keys for other users so they can create and manage buckets for their tenant account. After you create a new access key, update the application with the new access key ID and secret access key. For security, don't create more keys than the user needs, and delete the keys that aren't being used. If you have only one key and it is about to expire, create a new key before the old one expires, and then delete the old one.

Each key can have a specific expiration time or no expiration. Follow these guidelines for expiration time:

- Set an expiration time for the keys to limit the user's access to a certain time period. Setting a short expiration time can help reduce risk if the access key ID and secret access key are accidentally exposed. Expired keys are removed automatically.
- If the security risk in your environment is low and you don't need to periodically create new keys, you don't have to set an expiration time for the keys. If you decide later to create new keys, delete the old keys manually.



The S3 buckets and objects belonging to a user can be accessed using the access key ID and secret access key displayed for that user in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from the account, and never share them with other users.

Steps

1. Select **ACCESS MANAGEMENT > Users**.
2. Select the user whose S3 access keys you want to manage.

The user detail page appears.

3. Select **Access keys**, then select **Create key**.
4. Do one of the following:
 - Select **Don't set an expiration time** to create a key that does not expire. (Default)
 - Select **Set an expiration time**, and set the expiration date and time.



The expiration date can be a maximum of five years from the current date. The expiration time can be a minimum of one minute from the current time.

5. Select **Create access key**.

The Download access key dialog box appears, listing the access key ID and secret access key.

6. Copy the access key ID and the secret access key to a safe location, or select **Download .csv** to save a spreadsheet file containing the access key ID and secret access key.



Don't close this dialog box until you have copied or downloaded this information. You can't copy or download keys after the dialog box has been closed.

7. Select **Finish**.

The new key is listed on the Access keys tab of the user details page.

8. If your tenant account has the **Use grid federation connection** permission, optionally use the Tenant Management API to manually clone S3 access keys from the tenant on the source grid to the tenant on the destination grid. See [Clone S3 access keys using the API](#).

View another user's S3 access keys

If you are using an S3 tenant and you have appropriate permissions, you can view another user's S3 access keys. You can sort the list by expiration time so you can determine which keys will expire soon. As needed, you can create new keys and delete keys that are no longer in use.

Before you begin

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



The S3 buckets and objects belonging to a user can be accessed using the access key ID and secret access key displayed for that user in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from the account, and never share them with other users.

Steps

1. Select **ACCESS MANAGEMENT > Users**.

2. From the Users page, select the user whose S3 access keys you want to view.
3. From the User details page, select **Access keys**.
4. Sort the keys by **Expiration time** or **Access key ID**.
5. As needed, create new keys and manually delete keys that the are no longer in use.

If you create new keys before the existing keys expire, the user can begin using the new keys without temporarily losing access to the objects in the account.

Expired keys are removed automatically.

Related information

- [Create another user's S3 access keys](#)
- [Delete another user's S3 access keys](#)

Delete another user's S3 access keys

If you are using an S3 tenant and you have appropriate permissions, you can delete another user's S3 access keys. After an access key is deleted, it can no longer be used to access the objects and buckets in the tenant account.

Before you begin

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



The S3 buckets and objects belonging to a user can be accessed using the access key ID and secret access key displayed for that user in the Tenant Manager. For this reason, protect access keys as you would a password. Rotate access keys on a regular basis, remove any unused keys from the account, and never share them with other users.

Steps

1. Select **ACCESS MANAGEMENT > Users**.
2. From the Users page, select the user whose S3 access keys you want to manage.
3. From the User details page, select **Access keys**, and then select the checkbox for each access key you want to delete.
4. Select **Actions > Delete selected key**.
5. From the confirmation dialog box, select **Delete key**.

A confirmation message appears in the upper right corner of the page.

Manage S3 buckets

Create an S3 bucket

You can use the Tenant Manager to create S3 buckets for object data.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the Root access or Manage all buckets [permission](#). These permissions override the permissions settings in group or bucket policies.



Permissions to set or modify S3 Object Lock properties of buckets or objects can be granted by [bucket policy](#) or [group policy](#).

- If you plan to enable S3 Object Lock for a bucket, a grid admin has enabled the global S3 Object Lock setting for the StorageGRID system, and you have reviewed the requirements for S3 Object Lock buckets and objects.
- If each tenant will have 5,000 buckets, each Storage Node in the grid has a minimum of 64 GB of RAM.



Each grid can have a maximum of 100,000 buckets.

Access the wizard

Steps

1. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
2. Select **Create bucket**.

Enter details

Steps

1. Enter details for the bucket.

Field	Description
Bucket name	A name for the bucket that complies with these rules: • Must be unique across each StorageGRID system (not just unique within the tenant account). • Must be DNS compliant. • Must contain at least 3 and no more than 63 characters. • Each label must start and end with a lowercase letter or a number and can only use lowercase letters, numbers, and hyphens. • Must not contain periods in virtual hosted style requests. Periods will cause problems with server wildcard certificate verification. For more information, see the Amazon Web Services (AWS) documentation on bucket naming rules. Note: You can't change the bucket name after creating the bucket.

Field	Description
Region	The bucket's region. Your StorageGRID administrator manages the available regions. A bucket's region can affect the data-protection policy applied to objects. By default, all buckets are created in the <code>us-east-1</code> region. Note: You can't change the region after creating the bucket.

2. Select **Continue**.

Manage settings

Steps

1. Optionally, enable object versioning for the bucket.

Enable object versioning if you want to store every version of each object in this bucket. You can then retrieve previous versions of an object as needed. You must enable object versioning if the bucket will be used for cross-grid replication.

2. If the global S3 Object Lock setting is enabled, optionally enable S3 Object Lock for the bucket to store objects using a write-once-read-many (WORM) model.

Enable S3 Object Lock for a bucket only if you need to keep objects for fixed amount of time, for example, to meet certain regulatory requirements. S3 Object Lock is a permanent setting that helps you prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely.



After the S3 Object Lock setting is enabled for a bucket, it can't be disabled. Anyone with the correct permissions can add objects to this bucket that can't be changed. You might not be able to delete these objects or the bucket itself.

If you enable S3 Object Lock for a bucket, bucket versioning is enabled automatically.

3. If you selected **Enable S3 Object Lock**, optionally enable **Default retention** for this bucket.



Your grid administrator must give you permission to [use specific features of S3 Object Lock](#).

When **Default retention** is enabled, new objects added to the bucket will be automatically protected from being deleted or overwritten. The **Default retention** setting does not apply to objects that have their own retention periods.

- a. If **Default retention** is enabled, specify a **Default retention mode** for the bucket.

Default retention mode	Description
Governance	• Users with the <code>s3:BypassGovernanceRetention</code> permission can use the <code>x-amz-bypass-governance-retention: true</code> request header to bypass retention settings. • These users can delete an object version before its retain-until-date is reached. • These users can increase, decrease, or remove an object's retain-until-date.
Compliance	• The object can't be deleted until its retain-until-date is reached. • The object's retain-until-date can be increased, but it can't be decreased. • The object's retain-until-date can't be removed until that date is reached. Note: Your grid administrator must allow you to use compliance mode.

b. If **Default retention** is enabled, specify the **Default retention period** for the bucket.

The **Default retention period** indicates how long new objects added to this bucket should be retained, starting from the time they are ingested. Specify a value that is less than or equal to the maximum retention period for the tenant, as set by the grid administrator.

A *maximum* retention period, which can be a value from 1 day to 100 years, is set when the grid administrator creates the tenant. When you set a *default* retention period, it can't exceed the value set for the maximum retention period. If needed, ask your grid administrator to increase or decrease the maximum retention period.

4. Optionally, select **Enable capacity limit**.

Capacity limit is the maximum capacity available for this bucket's objects. This value represents a logical amount (object size), not a physical amount (size on disk).

If no limit is set, the capacity for this bucket is unlimited. Refer to [Capacity limit usage](#) for more information.

5. Select **Create bucket**.

The bucket is created and added to the table on the Buckets page.

6. Optionally, select **Go to bucket details page** to [view bucket details](#) and perform additional configuration.

View bucket details

You can view the buckets in your tenant account.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).

- You belong to a user group that has the [Root access, Manage all buckets, or View all buckets permission](#). These permissions override the permission settings in group or bucket policies.

Steps

1. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.

The Buckets page appears.

2. Review the summary table for each bucket.

As required, you can sort the information by any column, or you can page forward and back through the list.



The Object Count, Space Used, and Usage values displayed are estimates. These estimates are affected by the timing of ingests, network connectivity, and node status. If buckets have versioning enabled, deleted object versions are included in the object count.

Name

The bucket's unique name, which can't be changed.

Enabled features

The list of features that are enabled for the bucket.

S3 Object Lock

Whether S3 Object Lock is enabled for the bucket.

This column appears only if S3 Object Lock is enabled for the grid. This column also shows information for any legacy Compliant buckets.

Region

The bucket's region, which can't be changed. This column is hidden by default.

Object count

The number of objects in this bucket. If buckets have versioning enabled, non-current object versions are included in this value.

When objects are added or deleted, this value might not update immediately.

Space used

The logical size of all objects in the bucket. The logical size does not include the actual space required for replicated or erasure-coded copies or for object metadata.

This value can take up to 10 minutes to update.

Usage

The percentage used of the bucket's capacity limit, if one has been set.

The usage value is based on internal estimates and might be exceeded in some cases. For example, StorageGRID checks capacity limit (if set) when a tenant starts uploading objects and rejects new ingests to this bucket if the tenant has exceeded the capacity limit. However, StorageGRID does not take into account the size of the current upload when determining if the capacity limit has been exceeded. If objects are deleted, a tenant might be temporarily prevented from uploading new objects to this bucket until the capacity limit usage is recalculated. The calculations can take 10 minutes or longer.

This value indicates logical size, not physical size needed to store the objects and their metadata.

Capacity

If set, the capacity limit for the bucket.

Date created

The date and time the bucket was created. This column is hidden by default.

3. To view details for a specific bucket, select the bucket name from the table.
 - a. View the summary information at the top of the web page to confirm the details for the bucket, such as Region and Object count.
 - b. View the Capacity limit usage bar. If the usage is 100% or near 100%, consider increasing the limit or deleting some objects.
 - c. As needed, select **Delete objects in bucket** and **Delete bucket**.



Pay close attention to the cautions that appear when you select each of these options. For more information, refer to:

- [Delete all objects in a bucket](#)
- [Delete a bucket](#) (bucket must be empty)

- d. View or change settings for the bucket in each of the tabs as needed.
 - **S3 Console:** View the objects for the bucket. For more information, refer to [Use S3 Console](#).
 - **Bucket options:** View or change option settings. Some settings, such as S3 Object Lock, can't be changed after the bucket is created.
 - [Manage bucket consistency](#)
 - [Last access time updates](#)
 - [Capacity limit](#)
 - [Object versioning](#)
 - [S3 Object Lock](#)
 - [Default bucket retention](#)
 - [Manage cross-grid replication](#) (if allowed for the tenant)
 - **Platform services:** [Manage platform services](#) (if allowed for the tenant)
 - **Bucket access:** View or change option settings. You must have specific access permissions.
 - Configure [Cross-Origin Resource Sharing \(CORS\)](#) so the bucket and objects in the bucket will be accessible to web applications in other domains.
 - [Control user access](#) for an S3 bucket and objects in that bucket.

Apply an ILM policy tag to a bucket

Choose an ILM policy tag to apply to a bucket based on your object storage requirements.

The ILM policy controls where the object data is stored and whether it is deleted after a certain time period. Your grid administrator creates ILM policies and assigns them to ILM policy tags when using multiple active

policies.



Avoid frequently reassigning a bucket's policy tag. Otherwise, performance issues might occur.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access](#), [Manage all buckets](#), or [View all buckets permission](#). These permissions override the permission settings in group or bucket policies.

Steps

1. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.

The Buckets page appears. As required, you can sort the information by any column, or you can page forward and back through the list.

2. Select the name of the bucket you want to assign an ILM policy tag to.

You can also change the ILM policy tag assignment for a bucket that already has a tag assigned.



The Object Count and Space Used values displayed are estimates. These estimates are affected by the timing of ingests, network connectivity, and node status. If buckets have versioning enabled, deleted object versions are included in the object count.

3. In the Bucket options tab, expand the ILM policy tag accordion. This accordion only appears if your grid administrator has enabled the use of custom policy tags.
4. Read the description of each policy tag to determine which tag should be applied to the bucket.



Changing the ILM policy tag for a bucket will trigger ILM reevaluation of all objects in the bucket. If the new policy retains objects for a limited time, older objects will be deleted.

5. Select the radio button for the tag you want to assign to the bucket.
6. Select **Save changes**. A new S3 bucket tag will be set on the bucket with the key `NTAP-SG-ILM-BUCKET-TAG` and the value of the ILM policy tag name.



Ensure that your S3 applications do not accidentally override or delete the new bucket tag. If this tag is omitted when applying a new TagSet to the bucket, objects in the bucket will revert to being evaluated against the default ILM policy.



Set and modify ILM policy tags using only the Tenant Manager or Tenant Manager API where the ILM policy tag is validated. Do not modify the `NTAP-SG-ILM-BUCKET-TAG` ILM policy tag using the S3 PutBucketTagging API or the S3 DeleteBucketTagging API.



Changing the policy tag assigned to a bucket has a temporary performance impact while objects are being reevaluated using the new ILM policy.

Manage bucket policy

You can control user access for an S3 bucket and the objects in that bucket.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#). The View all buckets and Manage all buckets permissions only allow viewing.
- You've verified that the required number of Storage Nodes and sites are available. If two or more Storage Nodes are not available within any site, or if a site is not available, changes to these settings might not be available.

Steps

1. Select **Buckets**, then select the bucket you want to manage.
2. On the bucket details page, select **Bucket access** > **Bucket policy**.
3. Do one of the following:
 - Enter a bucket policy by selecting the **Enable policy** checkbox. Then enter a valid JSON formatted string.

Each bucket policy has a size limit of 20,480 bytes.

- Modify an existing policy by editing the string.
- Disable a policy by unselecting **Enable policy**.

For detailed information about bucket policies, including language syntax and examples, see [Example bucket policies](#).

Manage bucket consistency

Consistency values can be used to specify the availability of bucket setting changes as well as to provide a balance between the availability of the objects within a bucket and the consistency of those objects across different Storage Nodes and sites. You can change the consistency values to be different from the default values so that client applications can meet their operational needs.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Manage all buckets or Root access permission](#). These permissions override the permissions settings in group or bucket policies.

Bucket consistency guidelines

The bucket consistency is used to determine the consistency for client applications affecting objects within that S3 bucket. In general, you should use the **Read-after-new-write** consistency for your buckets.

Change bucket consistency

If the **Read-after-new-write** consistency does not meet the client application's requirements, you can change the consistency by setting the bucket consistency or by using the `Consistency-Control` header. The `Consistency-Control` header overrides the bucket consistency.



When you change a bucket's consistency, only those objects that are ingested after the change are guaranteed to meet the revised setting.

Steps

1. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
2. Select the bucket name from the table.

The bucket details page appears.

3. From the **Bucket options** tab, select the ** accordion.
4. Select a consistency for operations performed on the objects in this bucket.
 - **All**: Provides the highest level of consistency. All nodes receive the data immediately, or the request will fail.
 - **Strong-global**: Guarantees read-after-write consistency for all client requests across all sites.
 - **Strong-site**: Guarantees read-after-write consistency for all client requests within a site.
 - **Read-after-new-write** (default): Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.
 - **Available**: Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that don't exist). Not supported for S3 FabricPool buckets.
5. Select **Save changes**.

What happens when you change bucket settings

Buckets have multiple settings that affect the behavior of the buckets and the objects within those buckets.

The following bucket settings use **strong** consistency by default. If two or more Storage Nodes are not available within any site, or if a site is not available, any changes to these settings might not be available.

- [Background empty bucket deletion](#)
- [Last Access Time](#)
- [Bucket lifecycle](#)
- [Bucket policy](#)
- [Bucket tagging](#)
- [Bucket versioning](#)
- [S3 Object Lock](#)
- [Bucket encryption](#)



The consistency value for bucket versioning, S3 Object Lock, and bucket encryption cannot be set to a value that is not strongly consistent.

The following bucket settings do not use strong consistency and have higher availability for changes. Changes to these settings might take some time before having an effect.

- [Platform services configuration: Notification, Replication, or Search integration](#)
- [CORS configuration](#)
- [Change bucket consistency](#)



If the default consistency used when changing bucket settings does not meet the client application's requirements, you can change the consistency by using the `Consistency-Control` header for the [S3 REST API](#) or by using the `reducedConsistency` or `force` options in the [Tenant Management API](#).

Enable or disable last access time updates

When grid administrators create the information lifecycle management (ILM) rules for a StorageGRID system, they can optionally specify that an object's last access time be used to determine whether to move that object to a different storage location. If you are using an S3 tenant, you can take advantage of such rules by enabling last access time updates for the objects in an S3 bucket.

These instructions only apply to StorageGRID systems that include at least one ILM rule that uses the **Last access time** option as an advanced filter or as a reference time. You can ignore these instructions if your StorageGRID system does not include such a rule. See [Use Last access time in ILM rules](#) for details.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Manage all buckets or Root access permission](#). These permissions override the permissions settings in group or bucket policies.

About this task

Last access time is one of the options available for the **Reference time** placement instruction for an ILM rule. Setting the Reference time for a rule to Last access time lets grid administrators specify that objects be placed in certain storage locations based on when those objects were last retrieved (read or viewed).

For example, to ensure that recently viewed objects remain on faster storage, a grid administrator can create an ILM rule specifying the following:

- Objects that have been retrieved in the past month should remain on local Storage Nodes.
- Objects that have not been retrieved in the past month should be moved to an off-site location.

By default, updates to last access time are disabled. If your StorageGRID system includes an ILM rule that uses the **Last access time** option and you want this option to apply to objects in this bucket, you must enable updates to last access time for the S3 buckets specified in that rule.



Updating the last access time when an object is retrieved can reduce StorageGRID performance, especially for small objects.

A performance impact occurs with last access time updates because StorageGRID must perform these additional steps every time objects are retrieved:

- Update the objects with new timestamps
- Add the objects to the ILM queue, so they can be reevaluated against current ILM rules and policy

The table summarizes the behavior applied to all objects in the bucket when last access time is disabled or enabled.

Type of request	Behavior if last access time is disabled (default)		Behavior if last access time is enabled	
	Last access time updated?	Object added to ILM evaluation queue?	Last access time updated?	Object added to ILM evaluation queue?
Request to retrieve an object, its access control list, or its metadata	No	No	Yes	Yes
Request to update an object's metadata	Yes	Yes	Yes	Yes
Request to list objects or object versions	No	No	No	No
Request to copy an object from one bucket to another	- No, for the source copy - Yes, for the destination copy	- No, for the source copy - Yes, for the destination copy	- Yes, for the source copy - Yes, for the destination copy	- Yes, for the source copy - Yes, for the destination copy
Request to complete a multipart upload	Yes, for the assembled object	Yes, for the assembled object	Yes, for the assembled object	Yes, for the assembled object

Steps

1. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
2. Select the bucket name from the table.

The bucket details page appears.
3. From the **Bucket options** tab, select the **Last access time updates** accordion.
4. Enable or disable last access time updates.
5. Select **Save changes**.

Change object versioning for a bucket

If you are using an S3 tenant, you can change the versioning state for S3 buckets.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Manage all buckets or Root access permission](#). These permissions override the permissions settings in group or bucket policies.
- You've verified that the required number of Storage Nodes and sites are available. If two or more Storage Nodes are not available within any site, or if a site is not available, changes to these settings might not be available.

About this task

You can enable or suspend object versioning for a bucket. After you enable versioning for a bucket, it can't return to an unversioned state. However, you can suspend versioning for the bucket.

- Disabled: Versioning has never been enabled
- Enabled: Versioning is enabled
- Suspended: Versioning was previously enabled and is suspended

For more information, see the following:

- [Object versioning](#)
- [ILM rules and policies for S3 versioned objects \(Example 4\)](#)
- [How objects are deleted](#)

Steps

1. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
2. Select the bucket name from the table.

The bucket details page appears.

3. From the **Bucket options** tab, select the **Object versioning** accordion.
4. Select a versioning state for the objects in this bucket.

Object versioning must remain enabled for a bucket used for cross-grid replication. If S3 Object Lock or legacy compliance is enabled, the **Object versioning** options are disabled.

Option	Description
Enable versioning	Enable object versioning if you want to store every version of each object in this bucket. You can then retrieve previous versions of an object as needed. Objects that were already in the bucket will be versioned when they are modified by a user.
Suspend versioning	Suspend object versioning if you no longer want new object versions to be created. You can still retrieve any existing object versions.

5. Select **Save changes**.

Use S3 Object Lock to retain objects

You can use S3 Object Lock if buckets and objects must comply with regulatory requirements for retention.



Your grid administrator must give you permission to use specific features of S3 Object Lock.

What is S3 Object Lock?

The StorageGRID S3 Object Lock feature is an object-protection solution that is equivalent to S3 Object Lock in Amazon Simple Storage Service (Amazon S3).

When the global S3 Object Lock setting is enabled for a StorageGRID system, an S3 tenant account can create buckets with or without S3 Object Lock enabled. If a bucket has S3 Object Lock enabled, bucket versioning is required and is enabled automatically.

A bucket without S3 Object Lock can only have objects without retention settings specified. No ingested objects will have retention settings.

A bucket with S3 Object Lock can have objects with and without retention settings specified by S3 client applications. Some objects ingested will have retention settings.

A bucket with S3 Object Lock and default retention configured can have uploaded objects with retention settings specified and new objects without retention settings. The new objects use the default setting, because the retention setting hasn't been configured at the object-level.

Effectively, all newly ingested objects have retention settings when default retention is configured. Existing objects without object retention settings remain unaffected.

Retention modes

The StorageGRID S3 Object Lock feature supports two retention modes to apply different levels of protection to objects. These modes are equivalent to the Amazon S3 retention modes.

- In compliance mode:
 - The object can't be deleted until its retain-until-date is reached.
 - The object's retain-until-date can be increased, but it can't be decreased.
 - The object's retain-until-date can't be removed until that date is reached.
- In governance mode:
 - Users with special permission can use a bypass header in requests to modify certain retention settings.
 - These users can delete an object version before its retain-until-date is reached.
 - These users can increase, decrease, or remove an object's retain-until-date.

Retention settings for object versions

If a bucket is created with S3 Object Lock enabled, users can use the S3 client application to optionally specify the following retention settings for each object that is added to the bucket:

- **Retention mode:** Either compliance or governance.
- **Retain-until-date:** If an object version's retain-until-date is in the future, the object can be retrieved, but it can't be deleted.
- **Legal hold:** Applying a legal hold to an object version immediately locks that object. For example, you might need to put a legal hold on an object that is related to an investigation or legal dispute. A legal hold has no expiration date, but remains in place until it is explicitly removed. Legal holds are independent of the retain-until-date.



If an object is under a legal hold, no one can delete the object, regardless of its retention mode.

For details on the object settings, see [Use S3 REST API to configure S3 Object Lock](#).

Default retention setting for buckets

If a bucket is created with S3 Object Lock enabled, users can optionally specify the following default settings for the bucket:

- **Default retention mode:** Either compliance or governance.
- **Default retention period:** How long new object versions added to this bucket should be retained, starting from the day they are added.

The default bucket settings apply only to new objects that don't have their own retention settings. Existing bucket objects aren't affected when you add or change these default settings.

See [Create an S3 bucket](#) and [Update S3 Object Lock default retention](#).

S3 Object Lock tasks

The following lists for grid administrators and tenant users contain the high-level tasks for using the S3 Object Lock feature.

Grid administrator

- Enable global S3 Object Lock setting for entire StorageGRID system.
- Ensure that information lifecycle management (ILM) policies are *compliant*; that is, they meet the [requirements of buckets with S3 Object Lock enabled](#).
- As needed, allow a tenant to use Compliance as the retention mode. Otherwise, only Governance mode is allowed.
- As needed, set a maximum retention period for a tenant.

Tenant user

- Review considerations for buckets and objects with S3 Object Lock.
- As needed, contact grid administrator to enable global S3 Object Lock setting and set permissions.
- Create buckets with S3 Object Lock enabled.
- Optionally, configure default retention settings for a bucket:
 - Default retention mode: Governance or Compliance, if allowed by the grid administrator.
 - Default retention period: Must be less than or equal to maximum retention period set by grid administrator.
- Use the S3 client application to add objects and optionally set object-specific retention:
 - Retention mode. Governance or Compliance, if allowed by the grid administrator.
 - Retain Until Date: Must be less than or equal to what is allowed by the maximum retention period set by grid administrator.

Requirements for buckets with S3 Object Lock enabled

- If the global S3 Object Lock setting is enabled for the StorageGRID system, you can use the Tenant Manager, the Tenant Management API, or the S3 REST API to create buckets with S3 Object Lock enabled.
- If you plan to use S3 Object Lock, you must enable S3 Object Lock when you create the bucket. You can't enable S3 Object Lock for an existing bucket.
- When S3 Object Lock is enabled for a bucket, StorageGRID automatically enables versioning for that

bucket. You can't disable S3 Object Lock or suspend versioning for the bucket.

- Optionally, you can specify a default retention mode and retention period for each bucket using the Tenant Manager, the Tenant Management API, or the S3 REST API. The bucket's default retention settings apply only to new objects added to the bucket that don't have their own retention settings. You can override these default settings by specifying a retention mode and retain-until-date for each object version when it is uploaded.
- Bucket lifecycle configuration is supported for buckets with S3 Object Lock enabled.
- CloudMirror replication is not supported for buckets with S3 Object Lock enabled.

Requirements for objects in buckets with S3 Object Lock enabled

- To protect an object version, you can specify default retention settings for the bucket, or you can specify retention settings for each object version. Object-level retention settings can be specified using the S3 client application or the S3 REST API.
- Retention settings apply to individual object versions. An object version can have both a retain-until-date and a legal hold setting, one but not the other, or neither. Specifying a retain-until-date or a legal hold setting for an object protects only the version specified in the request. You can create new versions of the object, while the previous version of the object remains locked.

Lifecycle of objects in buckets with S3 Object Lock enabled

Each object that is saved in a bucket with S3 Object Lock enabled goes through these stages:

1. Object ingest

When an object version is added to bucket that has S3 Object Lock enabled, retention settings are applied as follows:

- If retention settings are specified for the object, the object-level settings are applied. Any default bucket settings are ignored.
- If no retention settings are specified for the object, the default bucket settings are applied, if they exist.
- If no retention settings are specified for the object or the bucket, the object is not protected by S3 Object Lock.

If retention settings are applied, both the object and any S3 user-defined metadata are protected.

2. Object retention and deletion

Multiple copies of each protected object are stored by StorageGRID for the specified retention period. The exact number and type of object copies and the storage locations are determined by the compliant rules in the active ILM policies. Whether a protected object can be deleted before its retain-until-date is reached depends on its retention mode.

- If an object is under a legal hold, no one can delete the object, regardless of its retention mode.

Can I still manage legacy Compliant buckets?

The S3 Object Lock feature replaces the Compliance feature that was available in previous StorageGRID versions. If you created compliant buckets using a previous version of StorageGRID, you can continue to manage the settings of these buckets; however, you can no longer create new compliant buckets. For instructions, see

[NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5.](#)

Update S3 Object Lock default retention

If you enabled S3 Object Lock when you created the bucket, you can edit the bucket to change the default retention settings. You can enable (or disable) default retention and set a default retention mode and retention period.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Manage all buckets or Root access permission](#). These permissions override the permissions settings in group or bucket policies.
- S3 Object Lock is enabled globally for your StorageGRID system, and you enabled S3 Object Lock when you created the bucket. See [Use S3 Object Lock to retain objects](#).

Steps

1. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
2. Select the bucket name from the table.

The bucket details page appears.

3. From the **Bucket options** tab, select the **S3 Object Lock** accordion.
4. Optionally, enable or disable **Default retention** for this bucket.

Changes to this setting don't apply to objects already in the bucket or to any objects that might have their own retention periods.

5. If **Default retention** is enabled, specify a **Default retention mode** for the bucket.

Default retention mode	Description
Governance	• Users with the <code>s3:BypassGovernanceRetention</code> permission can use the <code>x-amz-bypass-governance-retention: true</code> request header to bypass retention settings. • These users can delete an object version before its retain-until-date is reached. • These users can increase, decrease, or remove an object's retain-until-date.
Compliance	• The object can't be deleted until its retain-until-date is reached. • The object's retain-until-date can be increased, but it can't be decreased. • The object's retain-until-date can't be removed until that date is reached. Note: Your grid administrator must allow you to use compliance mode.

6. If **Default retention** is enabled, specify the **Default retention period** for the bucket.

The **Default retention period** indicates how long new objects added to this bucket should be retained,

starting from the time they are ingested. Specify a value that is less than or equal to the maximum retention period for the tenant, as set by the grid administrator.

A *maximum* retention period, which can be a value from 1 day to 100 years, is set when the grid administrator creates the tenant. When you set a *default* retention period, it can't exceed the value set for the maximum retention period. If needed, ask your grid administrator to increase or decrease the maximum retention period.

7. Select **Save changes**.

Configure cross-origin resource sharing (CORS)

You can configure cross-origin resource sharing (CORS) for an S3 bucket if you want that bucket and objects in that bucket to be accessible to web applications in other domains.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- For GET CORS configuration requests, you belong to a user group that has the [Manage all buckets or View all buckets permission](#). These permissions override the permissions settings in group or bucket policies.
- For PUT CORS configuration requests, you belong to a user group that has the [Manage all buckets permission](#). This permission overrides the permissions settings in group or bucket policies.
- The [Root access permission](#) provides access to all CORS configuration requests.

About this task

Cross-origin resource sharing (CORS) is a security mechanism that allows client web applications in one domain to access resources in a different domain. For example, suppose you use an S3 bucket named `Images` to store graphics. By configuring CORS for the `Images` bucket, you can allow the images in that bucket to be displayed on the website `http://www.example.com`.

Enable CORS for a bucket

Steps

1. Use a text editor to create the required XML. This example shows the XML used to enable CORS for an S3 bucket. Specifically:
 - Allows any domain to send GET requests to the bucket
 - Only allows the `http://www.example.com` domain to send GET, POST, and DELETE requests
 - All request headers are allowed

```

<CORSConfiguration
  xmlns="http://s3.amazonaws.com/doc/2020-10-22/">
  <CORSRule>
    <AllowedOrigin>*</AllowedOrigin>
    <AllowedMethod>GET</AllowedMethod>
    <AllowedHeader>*</AllowedHeader>
  </CORSRule>
  <CORSRule>
    <AllowedOrigin>http://www.example.com</AllowedOrigin>
    <AllowedMethod>GET</AllowedMethod>
    <AllowedMethod>POST</AllowedMethod>
    <AllowedMethod>DELETE</AllowedMethod>
    <AllowedHeader>*</AllowedHeader>
  </CORSRule>
</CORSConfiguration>

```

For more information about the CORS configuration XML, see [Amazon Web Services \(AWS\) Documentation: Amazon Simple Storage Service User Guide](#).

2. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
3. Select the bucket name from the table.

The bucket details page appears.

4. From the **Bucket access** tab, select the **Cross-Origin Resource Sharing (CORS)** accordion.
5. Select the **Enable CORS** checkbox.
6. Paste the CORS configuration XML into the text box.
7. Select **Save changes**.

Modify CORS setting

Steps

1. Update the CORS configuration XML in the text box, or select **Clear** to start over.
2. Select **Save changes**.

Disable CORS setting

Steps

1. Clear the **Enable CORS** checkbox.
2. Select **Save changes**.

Delete objects in bucket

You can use the Tenant Manager to delete the objects in one or more buckets.

Considerations and requirements

Before performing these steps, note the following:

- When you delete the objects in a bucket, StorageGRID permanently removes all objects and all object versions in each selected bucket from all nodes and sites in your StorageGRID system. StorageGRID also removes any related object metadata. You will not be able to recover this information.
- Deleting all of the objects in a bucket might take minutes, days, or even weeks, based on the number of objects, object copies, and concurrent operations.
- If a bucket has [S3 Object Lock enabled](#), it might remain in the **Deleting objects: read-only** state for years.



A bucket that uses S3 Object Lock will remain in the **Deleting objects: read-only** state until the retention date is reached for all objects and any legal holds are removed.

- While objects are being deleted, the bucket's state is **Deleting objects: read-only**. In this state, you can't add new objects to the bucket.
- When all objects have been deleted, the bucket remains in the read-only state. You can do one of the following:
 - Return the bucket to write mode and reuse it for new objects
 - Delete the bucket
 - Keep the bucket in read-only mode to reserve its name for future use
- If a bucket has object versioning enabled, delete markers that were created in StorageGRID 11.8 or later can be removed using the Delete objects in bucket operations.
- If a bucket has object versioning enabled, the delete objects operation will not remove delete markers that were created in StorageGRID 11.7 or earlier. See information about deleting objects in a bucket in [How S3 versioned objects are deleted](#).
- If you use [cross-grid replication](#), note the following:
 - Using this option does not delete any objects from the bucket on the other grid.
 - If you select this option for the source bucket, the **Cross-grid replication failure** alert will be triggered if you add objects to the destination bucket on the other grid. If you can't guarantee no one will add objects to the bucket on the other grid, [disable cross-grid replication](#) for that bucket before deleting all bucket objects.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Root access permission](#). This permission overrides the permissions settings in group or bucket policies.

Steps

1. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.

The Buckets page appears and shows all existing S3 buckets.

2. Use the **Actions** menu or the details page for a specific bucket.

Actions menu

- a. Select the checkbox for each bucket you want to delete objects from.
- b. Select **Actions > Delete objects in bucket**.

Details page

- a. Select a bucket name to display its details.
- b. Select **Delete objects in bucket**.

3. When the confirmation dialog box appears, review the details, enter **Yes**, and select **OK**.
4. Wait for the delete operation to begin.

After a few minutes:

- A yellow status banner appears on the bucket details page. The progress bar represents what percentage of objects have been deleted.
- **(read-only)** appears after the bucket's name on the bucket details page.
- **(Deleting objects: read-only)** appears next to the bucket's name on the Buckets page.

Buckets > my-bucket

my-bucket (read-only)

Region: us-east-1

Date created: 2022-12-14 10:09:50 MST

Object count: 3

[View bucket contents in Experimental S3 Console](#)

Delete bucket

⚠ All bucket objects are being deleted

StorageGRID is deleting all copies of the objects in this bucket, which might take days or weeks. While objects are being deleted, the bucket is read only. To stop the operation, select **Stop deleting objects**. You cannot restore objects that have already been deleted.

0% (0 of 3 objects deleted)


Stop deleting objects

5. As required while the operation is running, select **Stop deleting objects** to halt the process. Then, optionally, select **Delete objects in bucket** to resume the process.

When you select **Stop deleting objects**, the bucket is returned to write mode; however, you can't access or restore any objects that have been deleted.

6. Wait for the operation to complete.

When the bucket is empty, the status banner is updated, but the bucket remains read only.



7. Do one of the following:

- Exit the page to keep the bucket in read-only mode. For example, you might keep an empty bucket in read-only mode to reserve the bucket name for future use.
- Delete the bucket. You can select **Delete bucket** to delete a single bucket or return the Buckets page and select **Actions > Delete** buckets to remove more than one bucket.



If you are unable to delete a versioned bucket after all objects were deleted, delete markers might remain. To delete the bucket, you must remove all remaining delete markers.

- Return the bucket to write mode and optionally reuse it for new objects. You can select **Stop deleting objects** for a single bucket or return to the Buckets page and select **Action > Stop deleting objects** for more than one bucket.

Delete S3 bucket

You can use the Tenant Manager to delete one or more S3 buckets that are empty.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Manage all buckets or Root access permission](#). These permissions override the permissions settings in group or bucket policies.
- The buckets you want to delete are empty. If buckets you want to delete are *not* empty, [delete objects from the bucket](#).

About this task

These instructions describe how to delete an S3 bucket using the Tenant Manager. You can also delete S3 buckets using the [Tenant Management API](#) or the [S3 REST API](#).

You can't delete an S3 bucket if it contains objects, noncurrent object versions, or delete markers. For information about how S3 versioned objects are deleted, see [How objects are deleted](#).

Steps

1. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.

The Buckets page appears and shows all existing S3 buckets.

2. Use the **Actions** menu or the details page for a specific bucket.

Actions menu

- a. Select the checkbox for each bucket you want to delete.
- b. Select **Actions > Delete buckets**.

Details page

- a. Select a bucket name to display its details.
- b. Select **Delete bucket**.

3. When the confirmation dialog box appears, select **Yes**.

StorageGRID confirms that each bucket is empty and then deletes each bucket. This operation might take a few minutes.

If a bucket is not empty, an error message appears. You must [delete all objects and any delete markers in the bucket](#) before you can delete the bucket.

Use S3 Console

You can use S3 Console to view and manage the objects in an S3 bucket.

S3 Console allows you to:

- Upload, download, rename, copy, move, and delete objects
- View, revert, download, and delete object versions
- Search for objects by prefix
- Manage object tags
- View object metadata
- View, create, rename, copy, move, and delete folders

S3 Console provides an improved user experience for the most common cases. It is not designed to replace CLI or API operations in all situations.



If using S3 Console results in operations taking too long (for example, minutes or hours), consider:

- Reducing the number of selected objects
- Using non-graphical (API or CLI) methods to access your data

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- If you want to manage objects, you belong to a user group that has the Root access permission. Alternatively, you belong to a user group that has the Use S3 Console tab permission and either the View all buckets permission or Manage all buckets permission. See [Tenant management permissions](#).
- An S3 Group or Bucket policy has been configured for the user. See [Use bucket and group access policies](#).
- You know the user's access key ID and secret access key. Optionally, you have a `.csv` file containing this information. See the [instructions for creating access keys](#).

Steps

1. Select **STORAGE** > **Buckets** > *bucket name*.
2. Select the S3 Console tab.
3. Paste the access key ID and secret access key into the fields. Otherwise, select **Upload access keys** and select your `.csv` file.
4. Select **Sign in**.
5. The table of bucket objects appears. You can manage objects as needed.

Additional information

- **Search by prefix:** The prefix search feature only searches for objects that begin with a specific word relative to the current folder. The search does not include objects that contain the word elsewhere. This rule also applies to objects within folders. For example, a search for `folder1/folder2/somefile-` would return objects that are within the `folder1/folder2/` folder and begin with the word `somefile-`.
- **Drag and drop:** You can drag and drop files from your computer's file manager to S3 Console. However, you cannot upload folders.
- **Operations on folders:** When you move, copy, or rename a folder, all objects in the folder are updated one at a time, which might take time.
- **Permanent deletion when bucket versioning is disabled:** When you overwrite or delete an object in a bucket with versioning disabled, the operation is permanent. See [Change object versioning for a bucket](#).

Manage S3 platform services

S3 platform services

Platform services overview and considerations

Before implementing platform services, review the overview and considerations for using these services.

For information about S3, see [Use S3 REST API](#).

Overview of platform services

StorageGRID platform services can help you implement a hybrid cloud strategy by allowing you to send event notifications and copies of S3 objects and object metadata to external destinations.

Because the target location for platform services is typically external to your StorageGRID deployment, platform services give you the power and flexibility that comes from using external storage resources, notification services, and search or analysis services for your data.

Any combination of platform services can be configured for a single S3 bucket. For example, you could configure both the [CloudMirror service](#) and [notifications](#) on a StorageGRID S3 bucket so that you can mirror specific objects to the Amazon Simple Storage Service (S3), while sending a notification about each such object to a third party monitoring application to help you track your AWS expenses.



The use of platform services must be enabled for each tenant account by a StorageGRID administrator using the Grid Manager or the Grid Management API.

How platform services are configured

Platform services communicate with external endpoints that you configure using the [Tenant Manager](#) or the [Tenant Management API](#). Each endpoint represents an external destination, such as a StorageGRID S3 bucket, an Amazon Web Services bucket, an Amazon SNS topic, or an Elasticsearch cluster hosted locally, on AWS, or elsewhere.

After you create an external endpoint, you can enable a platform service for a bucket by adding XML configuration to the bucket. The XML configuration identifies the objects that the bucket should act on, the action that the bucket should take, and the endpoint that the bucket should use for the service.

You must add separate XML configurations for each platform service that you want to configure. For example:

- If you want all objects whose keys start with `/images` to be replicated to an Amazon S3 bucket, you must add a replication configuration to the source bucket.
- If you also want to send notifications when these objects are stored to the bucket, you must add a notifications configuration.
- If you want to index the metadata for these objects, you must add the metadata notification configuration that is used to implement search integration.

The format for the configuration XML is governed by the S3 REST APIs used to implement StorageGRID platform services:

Platform service	S3 REST API	Refer to
CloudMirror replication	• GetBucketReplication • PutBucketReplication	• CloudMirror replication • Operations on buckets

Platform service	S3 REST API	Refer to
Notifications	• <code>GetBucketNotificationConfiguration</code> • <code>PutBucketNotificationConfiguration</code>	• Notifications • Operations on buckets
Search integration	• GET Bucket metadata notification configuration • PUT Bucket metadata notification configuration	• Search integration • StorageGRID custom operations

Considerations for using platform services

Consideration	Details
Destination endpoint monitoring	You must monitor the availability of each destination endpoint. If connectivity to the destination endpoint is lost for an extended period of time and a large backlog of requests exists, additional client requests (such as PUT requests) to StorageGRID will fail. You must retry these failed requests when the endpoint becomes reachable.
Destination endpoint throttling	StorageGRID software might throttle incoming S3 requests for a bucket if the rate at which the requests are being sent exceeds the rate at which the destination endpoint can receive the requests. Throttling only occurs when there is a backlog of requests waiting to be sent to the destination endpoint. The only visible effect is that the incoming S3 requests will take longer to execute. If you start to detect significantly slower performance, you should reduce the ingest rate or use an endpoint with higher capacity. If the backlog of requests continues to grow, client S3 operations (such as PUT requests) will eventually fail. CloudMirror requests are more likely to be affected by the performance of the destination endpoint because these requests typically involve more data transfer than search integration or event notification requests.
Ordering guarantees	StorageGRID guarantees ordering of operations on an object within a site. As long as all operations against an object are within the same site, the final object state (for replication) will always equal the state in StorageGRID. StorageGRID makes a best effort attempt to order requests when operations are made across StorageGRID sites. For example, if you write an object initially to site A and then later overwrite the same object at site B, the final object replicated by CloudMirror to the destination bucket is not guaranteed to be the newer object.

Consideration	Details
ILM-driven object deletions	To match the deletion behavior of the AWS CRR and Amazon Simple Notification Service, CloudMirror and event notification requests aren't sent when an object in the source bucket is deleted because of StorageGRID ILM rules. For example, no CloudMirror or event notifications requests are sent if an ILM rule deletes an object after 14 days. In contrast, search integration requests are sent when objects are deleted because of ILM.
Using Kafka endpoints	For Kafka endpoints, Mutual TLS is not supported. As a result, if you have <code>ssl.client.auth</code> set to <code>required</code> in your Kafka broker configuration, it might cause Kafka endpoint configuration issues. The authentication of Kafka endpoints uses the following authentication types. These types are different from those used for the authentication of other endpoints, such as Amazon SNS, and require username and password credentials. • SASL/PLAIN • SASL/SCRAM-SHA-256 • SASL/SCRAM-SHA-512 Note: Configured storage proxy settings do not apply to Kafka platform services endpoints.

Considerations for using CloudMirror replication service

Consideration	Details
Replication status	StorageGRID does not support the <code>x-amz-replication-status</code> header.
Object size	The maximum size for objects that can be replicated to a destination bucket by the CloudMirror replication service is 5 TiB, which is the same as the maximum <i>supported</i> object size. Note: The maximum <i>recommended</i> size for a single PutObject operation is 5 GiB (5,368,709,120 bytes). If you have objects that are larger than 5 GiB, use multipart upload instead.
Bucket versioning and version IDs	If the source S3 bucket in StorageGRID has versioning enabled, you should also enable versioning for the destination bucket. When using versioning, note that the ordering of object versions in the destination bucket is best effort and not guaranteed by the CloudMirror service, due to limitations in the S3 protocol. Note: Version IDs for the source bucket in StorageGRID aren't related to the version IDs for the destination bucket.

Consideration	Details
Tagging for object versions	The CloudMirror service does not replicate any PutObjectTagging or DeleteObjectTagging requests that supply a version ID, due to limitations in the S3 protocol. Because version IDs for the source and destination aren't related, there is no way to ensure that a tag update to a specific version ID will be replicated. In contrast, the CloudMirror service does replicate PutObjectTagging requests or DeleteObjectTagging requests that don't specify a version ID. These requests update the tags for the latest key (or the latest version if the bucket is versioned). Normal ingests with tags (not tagging updates) are also replicated.
Multipart uploads and ETag values	When mirroring objects that were uploaded using a multipart upload, the CloudMirror service does not preserve the parts. As a result, the ETag value for the mirrored object will be different than the ETag value of the original object.
Objects encrypted with SSE-C (server-side encryption with customer-provided keys)	The CloudMirror service does not support objects that are encrypted with SSE-C. If you attempt to ingest an object into the source bucket for CloudMirror replication and the request includes the SSE-C request headers, the operation fails.
Bucket with S3 Object Lock enabled	Replication is not supported for source or destination buckets with S3 Object Lock enabled.

Understand CloudMirror replication service

You can enable CloudMirror replication for an S3 bucket if you want StorageGRID to replicate specified objects added to the bucket to one or more external destination buckets.

For example, you might use CloudMirror replication to mirror specific customer records into Amazon S3 and then leverage AWS services to perform analytics on your data.



CloudMirror replication is not supported if the source bucket has S3 Object Lock enabled.

CloudMirror and ILM

CloudMirror replication operates independently of the grid's active ILM policies. The CloudMirror service replicates objects as they are stored to the source bucket and delivers them to the destination bucket as soon as possible. Delivery of replicated objects is triggered when object ingest succeeds.

CloudMirror and cross-grid replication

CloudMirror replication has important similarities and differences with the cross-grid replication feature. Refer to [Compare cross-grid replication and CloudMirror replication](#).

CloudMirror and S3 buckets

CloudMirror replication is typically configured to use an external S3 bucket as a destination. However, you can also configure replication to use another StorageGRID deployment or any S3-compatible service.

Existing buckets

When you enable CloudMirror replication for an existing bucket, only the new objects added to that bucket are replicated. Any existing objects in the bucket aren't replicated. To force the replication of existing objects, you can update the existing object's metadata by performing an object copy.



If you are using CloudMirror replication to copy objects to an Amazon S3 destination, be aware that Amazon S3 limits the size of user-defined metadata within each PUT request header to 2 KB. If an object has user-defined metadata greater than 2 KB, that object will not be replicated.

Multiple destination buckets

To replicate objects in a single bucket to multiple destination buckets, specify the destination for each rule in the replication configuration XML. You can't replicate an object to more than one bucket at the same time.

Versioned or unversioned buckets

You can configure CloudMirror replication on versioned or unversioned buckets. The destination buckets can be versioned or unversioned. You can use any combination of versioned and unversioned buckets. For example, you could specify a versioned bucket as the destination for an unversioned source bucket, or vice versa. You can also replicate between unversioned buckets.

Deletion, replication loops, and events

Deletion behavior

Is the same as the deletion behavior of the Amazon S3 service, Cross-Region Replication (CRR). Deleting an object in a source bucket never deletes a replicated object in the destination. If both source and destination buckets are versioned, the delete marker is replicated. If the destination bucket is not versioned, deleting an object in the source bucket doesn't replicate the delete marker to the destination bucket or delete the destination object.

Protection from replication loops

As objects are replicated to the destination bucket, StorageGRID marks them as "replicas." A destination StorageGRID bucket won't replicate objects marked as replicas again, protecting you from accidental replication loops. This replica marking is internal to StorageGRID and doesn't prevent you from leveraging AWS CRR when using an Amazon S3 bucket as the destination.



The custom header used to mark a replica is `x-ntap-sg-replica`. This marking prevents a cascading mirror. StorageGRID does support a bidirectional CloudMirror between two grids.

Events in the destination bucket

The uniqueness and ordering of events in the destination bucket aren't guaranteed. More than one identical copy of a source object might be delivered to the destination as a result of operations taken to guarantee delivery success. In rare cases, when the same object is updated simultaneously from two or more different StorageGRID sites, the ordering of operations on the destination bucket might not match the ordering of events on the source bucket.

Understand notifications for buckets

You can enable event notification for an S3 bucket if you want StorageGRID to send notifications about specified events to a destination Kafka cluster or Amazon Simple Notification Service.

For example, you could configure alerts to be sent to administrators about each object added to a bucket, where the objects represent log files associated with a critical system event.

Event notifications are created at the source bucket as specified in the notification configuration and are delivered to the destination. If an event associated with an object succeeds, a notification about that event is created and queued for delivery.

The uniqueness and ordering of notifications aren't guaranteed. More than one notification of an event might be delivered to the destination as a result of operations taken to guarantee delivery success. And because delivery is asynchronous, the time ordering of notifications at the destination is not guaranteed to match the ordering of events on the source bucket, particularly for operations that originate from different StorageGRID sites. You can use the `sequencer` key in the event message to determine the order of events for a particular object, as described in Amazon S3 documentation.

StorageGRID event notifications follow the Amazon S3 API with some limitations.

- The following event types are supported:
 - `s3:ObjectCreated:`
 - `s3:ObjectCreated:Put`
 - `s3:ObjectCreated:Post`
 - `s3:ObjectCreated:Copy`
 - `s3:ObjectCreated:CompleteMultipartUpload`
 - `s3:ObjectRemoved:`
 - `s3:ObjectRemoved:Delete`
 - `s3:ObjectRemoved:DeleteMarkerCreated`
 - `s3:ObjectRestore:Post`
- Event notifications sent from StorageGRID use the standard JSON format but don't include some keys and use specific values for others, as shown in the table:

Key name	StorageGRID value
<code>eventSource</code>	<code>sgws:s3</code>
<code>awsRegion</code>	<i>not included</i>
<code>x-amz-id-2</code>	<i>not included</i>
<code>arn</code>	<code>urn:sgws:s3:::bucket_name</code>

Understand search integration service

You can enable search integration for an S3 bucket if you want to use an external search and data analysis service for your object metadata.

The search integration service is a custom StorageGRID service that automatically and asynchronously sends S3 object metadata to a destination endpoint whenever an object is created or deleted, or its metadata or tags are updated. You can then use sophisticated search, data analysis, visualization, or machine learning tools

provided by the destination service to search, analyze, and gain insights from your object data.

For example, you could configure your buckets to send S3 object metadata to a remote Elasticsearch service. You could then use Elasticsearch to perform searches across buckets, and perform sophisticated analyses of patterns present in your object metadata.

Although Elasticsearch integration can be configured on a bucket with S3 Object Lock enabled, the S3 Object Lock metadata (including Retain Until Date and Legal Hold status) of the objects will not be included in the metadata sent to Elasticsearch.



Because the search integration service causes object metadata to be sent to a destination, its configuration XML is referred to as "*metadata* notification configuration XML." This configuration XML is different from the "notification configuration XML" used to enable *event* notifications.

Search integration and S3 buckets

You can enable the search integration service for any versioned or unversioned bucket. Search integration is configured by associating metadata notification configuration XML with the bucket that specifies which objects to act on and the destination for the object metadata.

Metadata notifications are generated in the form of a JSON document named with the bucket name, object name, and version ID, if any. Each metadata notification contains a standard set of system metadata for the object in addition to all of the object's tags and user metadata.



For tags and user metadata, StorageGRID passes dates and numbers to Elasticsearch as strings or as S3 event notifications. To configure Elasticsearch to interpret these strings as dates or numbers, follow the Elasticsearch instructions for dynamic field mapping and for mapping date formats. You must enable the dynamic field mappings on the index before you configure the search integration service. After a document is indexed, you can't edit the document's field types in the index.

Search notifications

Metadata notifications are generated and queued for delivery whenever:

- An object is created.
- An object is deleted, including when objects are deleted as a result of the operation of the grid's ILM policy.
- Object metadata or tags are added, updated, or deleted. The complete set of metadata and tags is always sent on update — not just the changed values.

After you add metadata notification configuration XML to a bucket, notifications are sent for any new objects that you create and for any objects that you modify by updating its data, user metadata, or tags. However, notifications aren't sent for any objects that were already in the bucket. To ensure that object metadata for all objects in the bucket is sent to the destination, you should do either of the following:

- Configure the search integration service immediately after creating the bucket and before adding any objects.
- Perform an action on all objects already in the bucket that will trigger a metadata notification message to be sent to the destination.

Search integration service and Elasticsearch

The StorageGRID search integration service supports an Elasticsearch cluster as a destination. As with the other platform services, the destination is specified in the endpoint whose URN is used in the configuration XML for the service. Use the [NetApp Interoperability Matrix Tool](#) to determine the supported versions of Elasticsearch.

Manage platform services endpoints

Configure platform services endpoints

Before you can configure a platform service for a bucket, you must configure at least one endpoint to be the destination for the platform service.

Access to platform services is enabled on a per-tenant basis by a StorageGRID administrator. To create or use a platform services endpoint, you must be a tenant user with Manage endpoints or Root access permission, in a grid whose networking has been configured to allow Storage Nodes to access external endpoint resources. For a single tenant, you can configure a maximum of 500 platform services endpoints. Contact your StorageGRID administrator for more information.

What is a platform services endpoint?

A platform services endpoint specifies the information that StorageGRID needs to access the external destination.

For example, if you want to replicate objects from a StorageGRID bucket to an Amazon S3 bucket, you create a platform services endpoint that includes the information and credentials StorageGRID needs to access the destination bucket on Amazon.

Each type of platform service requires its own endpoint, so you must configure at least one endpoint for each platform service you plan to use. After defining a platform services endpoint, you use the endpoint's URN as the destination in the configuration XML used to enable the service.

You can use the same endpoint as the destination for more than one source bucket. For example, you could configure several source buckets to send object metadata to the same search integration endpoint so that you can perform searches across multiple buckets. You can also configure a source bucket to use more than one endpoint as a target, which enables you to do things like send notifications about object creation to one Amazon Simple Notification Service (Amazon SNS) topic and notifications about object deletion to a second Amazon SNS topic.

Endpoints for CloudMirror replication

StorageGRID supports replication endpoints that represent S3 buckets. These buckets might be hosted on Amazon Web Services, the same or a remote StorageGRID deployment, or another service.

Endpoints for notifications

StorageGRID supports Amazon SNS and Kafka endpoints. Simple Queue Service (SQS) or AWS Lambda endpoints aren't supported.

For Kafka endpoints, Mutual TLS is not supported. As a result, if you have `ssl.client.auth` set to `required` in your Kafka broker configuration, it might cause Kafka endpoint configuration issues.

Endpoints for the search integration service

StorageGRID supports search integration endpoints that represent Elasticsearch clusters. These Elasticsearch clusters can be in a local data center or hosted in an AWS cloud or elsewhere.

The search integration endpoint refers to a specific Elasticsearch index and type. You must create the index in Elasticsearch before creating the endpoint in StorageGRID, or endpoint creation will fail. You don't need to create the type before creating the endpoint. StorageGRID will create the type if required when it sends object metadata to the endpoint.

Related information

[Administer StorageGRID](#)

Specify URN for platform services endpoint

When you create a platform services endpoint, you must specify a Unique Resource Name (URN). You will use the URN to reference the endpoint when you create a configuration XML for the platform service. The URN for each endpoint must be unique.

StorageGRID validates platform services endpoints as you create them. Before you create a platform services endpoint, confirm that the resource specified in the endpoint exists and that it can be reached.

URN elements

The URN for a platform services endpoint must start with either `arn:aws` or `urn:mysite`, as follows:

- If the service is hosted on Amazon Web Services (AWS), use `arn:aws`
- If the service is hosted on Google Cloud Platform (GCP), use `arn:aws`
- If the service is hosted locally, use `urn:mysite`

For example, if you are specifying the URN for a CloudMirror endpoint hosted on StorageGRID, the URN might begin with `urn:sgws`.

The next element of the URN specifies the type of platform service, as follows:

Service	Type
CloudMirror replication	s3
Notifications	sns or kafka
Search integration	es

For example, to continue specifying the URN for a CloudMirror endpoint hosted on StorageGRID, you would add `s3` to get `urn:sgws:s3`.

The final element of the URN identifies the specific target resource at the destination URI.

Service	Specific resource
CloudMirror replication	bucket-name
Notifications	sns-topic-name or kafka-topic-name
Search integration	domain-name/index-name/type-name Note: If the Elasticsearch cluster is not configured to create indexes automatically, you must create the index manually before you create the endpoint.

URNs for services hosted on AWS and GCP

For AWS and GCP entities, the complete URN is a valid AWS ARN. For example:

- CloudMirror replication:

```
arn:aws:s3:::bucket-name
```

- Notifications:

```
arn:aws:sns:region:account-id:topic-name
```

- Search integration:

```
arn:aws:es:region:account-id:domain/domain-name/index-name/type-name
```



For an AWS search integration endpoint, the domain-name must include the literal string domain/, as shown here.

URNs for locally-hosted services

When using locally-hosted services instead of cloud services, you can specify the URN in any way that creates a valid and unique URN, as long as the URN includes the required elements in the third and final positions. You can leave the elements indicated by optional blank, or you can specify them in any way that helps you identify the resource and make the URN unique. For example:

- CloudMirror replication:

```
urn:mysite:s3:optional:optional:bucket-name
```

For a CloudMirror endpoint hosted on StorageGRID, you can specify a valid URN that begins with urn:sgws:

```
urn:sgws:s3:optional:optional:bucket-name
```

- Notifications:

Specify an Amazon Simple Notification Service endpoint:

```
urn:mystore:sns:optional:optional:sns-topic-name
```

Specify a Kafka endpoint:

```
urn:mystore:kafka:optional:optional:kafka-topic-name
```

- Search integration:

```
urn:mystore:es:optional:optional:domain-name/index-name/type-name
```



For locally-hosted search integration endpoints, the `domain-name` element can be any string as long as the URN of the endpoint is unique.

Create platform services endpoint

You must create at least one endpoint of the correct type before you can enable a platform service.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- Platform services were enabled for your tenant account by a StorageGRID administrator.
- You belong to a user group that has the [Manage endpoints or Root access permission](#).
- The resource referenced by the platform services endpoint have been created:
 - CloudMirror replication: S3 bucket
 - Event notification: Amazon Simple Notification Service (Amazon SNS) or Kafka topic
 - Search notification: Elasticsearch index, if the destination cluster is not configured to automatically create indexes.
- You have the information about the destination resource:
 - Host and port for the Uniform Resource Identifier (URI)



If you plan to use a bucket hosted on a StorageGRID system as an endpoint for CloudMirror replication, contact the grid administrator to determine the values you need to enter.

- Unique Resource Name (URN)

Specify URN for platform services endpoint

- Authentication credentials (if required):

Search integration endpoints

For search integration endpoints, you can use the following credentials:

- Access Key: Access key ID and secret access key
- Basic HTTP: Username and password

CloudMirror replication endpoints

For CloudMirror replication endpoints, you can use the following credentials:

- Access Key: Access key ID and secret access key
- CAP (C2S Access Portal): Temporary credentials URL, server and client certificates, client keys, and an optional client private key passphrase.

Amazon SNS endpoints

For Amazon SNS endpoints, you can use the following credentials:

- Access Key: Access key ID and secret access key

Kafka endpoints

For Kafka endpoints, you can use the following credentials:

- SASL/PLAIN: Username and password
- SASL/SCRAM-SHA-256: Username and password
- SASL/SCRAM-SHA-512: Username and password

- Security certificate (if using a custom CA certificate)

- If the Elasticsearch security features are enabled, you have the monitor cluster privilege for connectivity testing, and either the write index privilege or both the index and delete index privileges for document updates.

Steps

1. Select **STORAGE (S3) > Platform services endpoints**. The Platform services endpoints page appears.
2. Select **Create endpoint**.
3. Enter a display name to briefly describe the endpoint and its purpose.

The type of platform service that the endpoint supports is shown beside the endpoint name when it is listed on the Endpoints page, so you don't need to include that information in the name.

4. In the **URI** field, specify the Unique Resource Identifier (URI) of the endpoint.

Use one of the following formats:

```
https://host:port  
http://host:port
```

If you don't specify a port, the following default ports are used:

- Port 443 for HTTPS URIs and port 80 for HTTP URIs (most endpoints)
- Port 9092 for HTTPS and HTTP URIs (Kafka endpoints only)

For example, the URI for a bucket hosted on StorageGRID might be:

```
https://s3.example.com:10443
```

In this example, `s3.example.com` represents the DNS entry for the virtual IP (VIP) of the StorageGRID high availability (HA) group, and `10443` represents the port defined in the load balancer endpoint.



Whenever possible, you should connect to an HA group of load-balancing nodes to avoid a single point of failure.

Similarly, the URI for a bucket hosted on AWS might be:

```
https://s3-aws-region.amazonaws.com
```



If the endpoint is used for the CloudMirror replication service, don't include the bucket name in the URI. You include the bucket name in the **URN** field.

5. Enter the Unique Resource Name (URN) for the endpoint.



You can't change an endpoint's URN after the endpoint has been created.

6. Select **Continue**.

7. Select a value for **Authentication type**.

Search integration endpoints

Enter or upload the credentials for a search integration endpoint.

The credentials that you supply must have write permissions for the destination resource.

Authentication type	Description	Credentials
Anonymous	Provides anonymous access to the destination. Only works for endpoints that have security disabled.	No authentication.
Access Key	Uses AWS-style credentials to authenticate connections with the destination.	• Access key ID • Secret access key
Basic HTTP	Uses a username and password to authenticate connections to the destination.	• Username • Password

CloudMirror replication endpoints

Enter or upload the credentials for a CloudMirror replication endpoint.

The credentials that you supply must have write permissions for the destination resource.

Authentication type	Description	Credentials
Anonymous	Provides anonymous access to the destination. Only works for endpoints that have security disabled.	No authentication.
Access Key	Uses AWS-style credentials to authenticate connections with the destination.	• Access key ID • Secret access key
CAP (C2S Access Portal)	Uses certificates and keys to authenticate connections to the destination.	• Temporary credentials URL • Server CA certificate (PEM file upload) • Client certificate (PEM file upload) • Client private key (PEM file upload, OpenSSL encrypted format or unencrypted private key format) • Client private key passphrase (optional)

Amazon SNS endpoints

Enter or upload the credentials for an Amazon SNS endpoint.

The credentials that you supply must have write permissions for the destination resource.

Authentication type	Description	Credentials
Anonymous	Provides anonymous access to the destination. Only works for endpoints that have security disabled.	No authentication.
Access Key	Uses AWS-style credentials to authenticate connections with the destination.	• Access key ID • Secret access key

Kafka endpoints

Enter or upload the credentials for a Kafka endpoint.

The credentials that you supply must have write permissions for the destination resource.

Authentication type	Description	Credentials
Anonymous	Provides anonymous access to the destination. Only works for endpoints that have security disabled.	No authentication.
SASL/PLAIN	Uses a username and password with plain text to authenticate connections to the destination.	• Username • Password
SASL/SCRAM-SHA-256	Uses a username and password using a challenge-response protocol and SHA-256 hashing to authenticate connections to the destination.	• Username • Password
SASL/SCRAM-SHA-512	Uses a username and password using a challenge-response protocol and SHA-512 hashing to authenticate connections to the destination.	• Username • Password

Select **Use delegation token authentication** if the username and password are derived from a delegation token that was obtained from a Kafka cluster.

8. Select **Continue**.

9. Select a radio button for **Verify server** to choose how TLS connection to the endpoint is verified.

Type of certificate verification	Description
Use custom CA certificate	Use a custom security certificate. If you select this setting, copy and paste the custom security certificate in the CA Certificate text box.

Type of certificate verification	Description
Use operating system CA certificate	Use the default Grid CA certificate installed on the operating system to secure connections.
Do not verify certificate	The certificate used for the TLS connection is not verified. This option is not secure.

10. Select **Test and create endpoint**.

- A success message appears if the endpoint can be reached using the specified credentials. The connection to the endpoint is validated from one node at each site.
- An error message appears if endpoint validation fails. If you need to modify the endpoint to correct the error, select **Return to endpoint details** and update the information. Then, select **Test and create endpoint**.



Endpoint creation fails if platform services aren't enabled for your tenant account. Contact your StorageGRID administrator.

After you have configured an endpoint, you can use its URN to configure a platform service.

Related information

- [Specify URN for platform services endpoint](#)
- [Configure CloudMirror replication](#)
- [Configure event notifications](#)
- [Configure search integration service](#)

Test connection for platform services endpoint

If the connection to a platform service has changed, you can test the connection for the endpoint to validate that the destination resource exists and that it can be reached using the credentials you specified.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Manage endpoints or Root access permission](#).

About this task

StorageGRID does not validate that the credentials have the correct permissions.

Steps

1. Select **STORAGE (S3) > Platform services endpoints**.

The Platform services endpoints page appears and shows the list of platform services endpoints that have already been configured.

2. Select the endpoint whose connection you want to test.

The endpoint details page appears.

3. Select **Test connection**.

- A success message appears if the endpoint can be reached using the specified credentials. The connection to the endpoint is validated from one node at each site.
- An error message appears if endpoint validation fails. If you need to modify the endpoint to correct the error, select **Configuration** and update the information. Then, select **Test and save changes**.

Edit platform services endpoint

You can edit the configuration for a platform services endpoint to change its name, URI, or other details. For example, you might need to update expired credentials or change the URI to point to a backup Elasticsearch index for failover. You can't change the URN for a platform services endpoint.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Manage endpoints or Root access permission](#).

Steps

1. Select **STORAGE (S3) > Platform services endpoints**.

The Platform services endpoints page appears and shows the list of platform services endpoints that have already been configured.

2. Select the endpoint you want to edit.

The endpoint details page appears.

3. Select **Configuration**.

4. As needed, change the configuration of the endpoint.



You can't change an endpoint's URN after the endpoint has been created.

a. To change the display name for the endpoint, select the edit icon .

b. As needed, change the URI.

c. As needed, change the authentication type.

- For Access Key authentication, change the key as necessary by selecting **Edit S3 key** and pasting a new access key ID and secret access key. If you need to cancel your changes, select **Revert S3 key edit**.
- For CAP (C2S Access Portal) authentication, change the temporary credentials URL or optional client private key passphrase and upload new certificate and key files as needed.



The Client private key must be in OpenSSL encrypted format or unencrypted private key format.

d. As needed, change the method for verifying the server.

5. Select **Test and save changes**.

- A success message appears if the endpoint can be reached using the specified credentials. The connection to the endpoint is verified from one node at each site.

- An error message appears if endpoint validation fails. Modify the endpoint to correct the error, and then select **Test and save changes**.

Delete platform services endpoint

You can delete an endpoint if you no longer want to use the associated platform service.

Before you begin

- You are signed in to the Tenant Manager using a [supported web browser](#).
- You belong to a user group that has the [Manage endpoints or Root access permission](#).

Steps

1. Select **STORAGE (S3) > Platform services endpoints**.

The Platform services endpoints page appears and shows the list of platform services endpoints that have already been configured.

2. Select the checkbox for each endpoint you want to delete.



If you delete a platform services endpoint that is in use, the associated platform service will be disabled for any buckets that use the endpoint. Any requests that have not yet been completed will be dropped. Any new requests will continue to be generated until you change your bucket configuration to no longer reference the deleted URN. StorageGRID will report these requests as unrecoverable errors.

3. Select **Actions > Delete endpoint**.

A confirmation message appears.

4. Select **Delete endpoint**.

Troubleshoot platform services endpoint errors

If an error occurs when StorageGRID attempts to communicate with a platform services endpoint, a message is displayed on the dashboard. On the Platform services endpoints page, the Last error column indicates how long ago the error occurred. No error is displayed if the permissions associated with an endpoint's credentials are incorrect.

Determine if error has occurred

If any platform services endpoint errors have occurred within the past 7 days, the Tenant Manager dashboard displays an alert message. You can go the Platform services endpoints page to see more details about the error.



One or more endpoints have experienced an error and might not be functioning properly. Go to the [Endpoints](#) page to view the error details. The last error occurred 2 hours ago.

The same error that appears on the dashboard also appears at the top of the Platform services endpoints page. To view a more detailed error message:

Steps

1. From the list of endpoints, select the endpoint that has the error.
2. On the endpoint details page, select **Connection**. This tab displays only the most recent error for an endpoint and indicates how long ago the error occurred. Errors that include the red X icon  occurred within the past 7 days.

Check if error is still current

Some errors might continue to be shown in the **Last error** column even after they are resolved. To see if an error is current or to force the removal of a resolved error from the table:

Steps

1. Select the endpoint.

The endpoint details page appears.

2. Select **Connection > Test connection**.

Selecting **Test connection** causes StorageGRID to validate that the platform services endpoint exists and that it can be reached with the current credentials. The connection to the endpoint is validated from one node at each site.

Resolve endpoint errors

You can use the **Last error** message on the endpoint details page to help determine what is causing the error. Some errors might require you to edit the endpoint to resolve the issue. For example, a CloudMirroring error can occur if StorageGRID is unable to access the destination S3 bucket because it does not have the correct access permissions or the access key has expired. The message is "Either the endpoint credentials or the destination access needs to be updated," and the details are "AccessDenied" or "InvalidAccessKeyId."

If you need to edit the endpoint to resolve an error, selecting **Test and save changes** causes StorageGRID to validate the updated endpoint and confirm that it can be reached with the current credentials. The connection to the endpoint is validated from one node at each site.

Steps

1. Select the endpoint.
2. On the endpoint details page, select **Configuration**.
3. Edit the endpoint configuration as needed.
4. Select **Connection > Test connection**.

Endpoint credentials with insufficient permissions

When StorageGRID validates a platform services endpoint, it confirms that the endpoint's credentials can be used to contact the destination resource and it does a basic permissions check. However, StorageGRID does not validate all of the permissions required for certain platform services operations. For this reason, if you receive an error when attempting to use a platform service (such as "403 Forbidden"), check the permissions associated with the endpoint's credentials.

Related information

- [Administer StorageGRID > Troubleshoot platform services](#)
- [Create platform services endpoint](#)

- [Test connection for platform services endpoint](#)
- [Edit platform services endpoint](#)

Configure CloudMirror replication

To enable CloudMirror replication for a bucket, you create and apply a valid bucket replication configuration XML.

Before you begin

- Platform services were enabled for your tenant account by a StorageGRID administrator.
- You have already created a bucket to act as the replication source.
- The endpoint that you intend to use as a destination for CloudMirror replication already exists, and you have its URN.
- You belong to a user group that has the [Manage all buckets or Root access permission](#). These permissions override the permission settings in group or bucket policies when configuring the bucket using the Tenant Manager.

About this task

CloudMirror replication copies objects from a source bucket to a destination bucket that is specified in an endpoint.

For general information about bucket replication and how to configure it, see [Amazon Simple Storage Service \(S3\) documentation: Replicating objects](#). For information about how StorageGRID implements GetBucketReplication, DeleteBucketReplication, and PutBucketReplication, see the [Operations on buckets](#).



CloudMirror replication has important similarities and differences with the cross-grid replication feature. To learn more, see [Compare cross-grid replication and CloudMirror replication](#).

Note the following requirements and characteristics when configuring CloudMirror replication:

- When you create and apply a valid bucket replication configuration XML, it must use the URN of an S3 bucket endpoint for each destination.
- Replication is not supported for source or destination buckets with S3 Object Lock enabled.
- If you enable CloudMirror replication on a bucket that contains objects, new objects added to the bucket are replicated, but the existing objects in the bucket aren't replicated. You must update existing objects to trigger replication.
- If you specify a storage class in the replication configuration XML, StorageGRID uses that class when performing operations against the destination S3 endpoint. The destination endpoint must also support the specified storage class. Be sure to follow any recommendations provided by the destination system vendor.

Steps

1. Enable replication for your source bucket:
 - Use a text editor to create the replication configuration XML required to enable replication, as specified in the S3 replication API.
 - When configuring the XML:
 - Note that StorageGRID only supports V1 of the replication configuration. This means that StorageGRID does not support the use of the `Filter` element for rules, and follows V1

conventions for deletion of object versions. See the Amazon documentation on replication configuration for details.

- Use the URN of an S3 bucket endpoint as the destination.
- Optionally add the `<StorageClass>` element, and specify one of the following:
 - `STANDARD`: The default storage class. If you don't specify a storage class when you upload an object, the `STANDARD` storage class is used.
 - `STANDARD_IA`: (Standard - infrequent access.) Use this storage class for data that is accessed less frequently, but that still requires rapid access when needed.
 - `REDUCED_REDUNDANCY`: Use this storage class for noncritical, reproducible data that can be stored with less redundancy than the `STANDARD` storage class.
- If you specify a `Role` in the configuration XML it will be ignored. This value is not used by StorageGRID.

```
<ReplicationConfiguration>
  <Role></Role>
  <Rule>
    <Status>Enabled</Status>
    <Prefix>2020</Prefix>
    <Destination>
      <Bucket>urn:sgws:s3:::2017-records</Bucket>
      <StorageClass>STANDARD</StorageClass>
    </Destination>
  </Rule>
</ReplicationConfiguration>
```

2. Select **View buckets** from the dashboard, or select **STORAGE (S3) > Buckets**.
3. Select the name of the source bucket.

The bucket details page appears.

4. Select **Platform services > Replication**.
5. Select the **Enable replication** checkbox.
6. Paste the replication configuration XML into the text box, and select **Save changes**.



Platform services must be enabled for each tenant account by a StorageGRID administrator using the Grid Manager or Grid Management API. Contact your StorageGRID administrator if an error occurs when you save the configuration XML.

7. Verify that replication is configured correctly:
 - a. Add an object to the source bucket that meets the requirements for replication as specified in the replication configuration.

In the example shown earlier, objects that match the prefix "2020" are replicated.

- b. Confirm that the object has been replicated to the destination bucket.

For small objects, replication happens quickly.

Related information

[Create platform services endpoint](#)

Configure event notifications

You enable notifications for a bucket by creating notification configuration XML and using the Tenant Manager to apply the XML to a bucket.

Before you begin

- Platform services were enabled for your tenant account by a StorageGRID administrator.
- You have already created a bucket to act as the source of notifications.
- The endpoint that you intend to use as a destination for event notifications already exists, and you have its URN.
- You belong to a user group that has the [Manage all buckets or Root access permission](#). These permissions override the permission settings in group or bucket policies when configuring the bucket using the Tenant Manager.

About this task

You configure event notifications by associating notification configuration XML with a source bucket. The notification configuration XML follows S3 conventions for configuring bucket notifications, with the destination Kafka or Amazon SNS topic specified as the URN of an endpoint.

For general information about event notifications and how to configure them, refer to the [Amazon documentation](#). For information about how StorageGRID implements the S3 bucket notification configuration API, refer to the [instructions for implementing S3 client applications](#).

Note the following requirements and characteristics when configuring event notifications for a bucket:

- When you create and apply valid notification configuration XML, it must use the URN of an event notifications endpoint for each destination.
- Although event notification can be configured on a bucket with S3 Object Lock enabled, the S3 Object Lock metadata (including Retain Until Date and Legal Hold status) of the objects won't be included in the notification messages.
- After you configure event notifications, whenever a specified event occurs for an object in the source bucket, a notification is generated and sent to the Amazon SNS or Kafka topic used as the destination endpoint.
- If you enable event notifications for a bucket that contains objects, notifications are sent only for actions that are performed after the notification configuration is saved.

Steps

1. Enable notifications for your source bucket:

- Use a text editor to create the notification configuration XML required to enable event notifications, as specified in the S3 notification API.
- When configuring the XML, use the URN of an event notifications endpoint as the destination topic.


```

<NotificationConfiguration>
  <TopicConfiguration>
    <Id>Image-created</Id>
    <Filter>
      <S3Key>
        <FilterRule>
          <Name>prefix</Name>
          <Value>images/</Value>
        </FilterRule>
      </S3Key>
    </Filter>
    <Topic>arn:aws:sns:us-east-1:050340950352:sgws-topic</Topic>
    <Event>s3:ObjectCreated:*</Event>
  </TopicConfiguration>
</NotificationConfiguration>

```

2. In the Tenant Manager, select **STORAGE (S3) > Buckets**.

3. Select the name of the source bucket.

The bucket details page appears.

4. Select **Platform services > Event notifications**.

5. Select the **Enable event notifications** checkbox.

6. Paste the notification configuration XML into the text box, and select **Save changes**.



Platform services must be enabled for each tenant account by a StorageGRID administrator using the Grid Manager or Grid Management API. Contact your StorageGRID administrator if an error occurs when you save the configuration XML.

7. Verify that event notifications are configured correctly:

- a. Perform an action on an object in the source bucket that meets the requirements for triggering a notification as configured in the configuration XML.

In the example, an event notification is sent whenever an object is created with the `images/` prefix.

- b. Confirm that a notification has been delivered to the destination Amazon SNS or Kafka topic.

For example, if your destination topic is hosted on the Amazon SNS, you could configure the service to send you an email when the notification is delivered.

```

{
  "Records": [
    {
      "eventVersion": "2.0",
      "eventSource": "sgws:s3",
      "eventTime": "2017-08-08T23:52:38Z",
      "eventName": "ObjectCreated:Put",
      "userIdentity": {
        "principalId": "11111111111111111111"
      },
      "requestParameters": {
        "sourceIPAddress": "193.51.100.20"
      },
      "responseElements": {
        "x-amz-request-id": "122047343"
      },
      "s3": {
        "s3SchemaVersion": "1.0",
        "configurationId": "Image-created",
        "bucket": {
          "name": "test1",
          "ownerIdentity": {
            "principalId": "11111111111111111111"
          },
          "arn": "arn:sgws:s3:::test1"
        },
        "object": {
          "key": "images/cat.jpg",
          "size": 0,
          "eTag": "d41d8cd98f00b204e9800998ecf8427e",
          "sequencer": "14D90402421461C7"
        }
      }
    }
  ]
}

```

If the notification is received at the destination topic, you have successfully configured your source bucket for StorageGRID notifications.

Related information

[Understand notifications for buckets](#)

[Use S3 REST API](#)

[Create platform services endpoint](#)

Configure the search integration service

You enable search integration for a bucket by creating search integration XML and using the Tenant Manager to apply the XML to the bucket.

Before you begin

- Platform services were enabled for your tenant account by a StorageGRID administrator.
- You have already created an S3 bucket whose contents you want to index.
- The endpoint that you intend to use as a destination for the search integration service already exists, and you have its URN.
- You belong to a user group that has the [Manage all buckets or Root access permission](#). These permissions override the permission settings in group or bucket policies when configuring the bucket using the Tenant Manager.

About this task

After you configure the search integration service for a source bucket, creating an object or updating an object's metadata or tags triggers object metadata to be sent to the destination endpoint.

If you enable the search integration service for a bucket that already contains objects, metadata notifications aren't automatically sent for existing objects. Update these existing objects to ensure that their metadata is added to the destination search index.

Steps

1. Enable search integration for a bucket:
 - Use a text editor to create the metadata notification XML required to enable search integration.
 - When configuring the XML, use the URN of a search integration endpoint as the destination.

Objects can be filtered on the prefix of the object name. For example, you could send metadata for objects with the prefix `images` to one destination, and metadata for objects with the prefix `videos` to another. Configurations that have overlapping prefixes aren't valid, and are rejected when they're submitted. For example, a configuration that includes one rule for objects with the prefix `test` and a second rule for objects with the prefix `test2` is not allowed.

As needed, refer to the [examples for the metadata configuration XML](#).

```
<MetadataNotificationConfiguration>
  <Rule>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>/Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

Elements in the metadata notification configuration XML:

Name	Description	Required
MetadataNotificationConfiguration	Container tag for rules used to specify the objects and destination for metadata notifications. Contains one or more Rule elements.	Yes
Rule	Container tag for a rule that identifies the objects whose metadata should be added to a specified index. Rules with overlapping prefixes are rejected. Included in the MetadataNotificationConfiguration element.	Yes
ID	Unique identifier for the rule. Included in the Rule element.	No
Status	Status can be 'Enabled' or 'Disabled'. No action is taken for rules that are disabled. Included in the Rule element.	Yes
Prefix	Objects that match the prefix are affected by the rule, and their metadata is sent to the specified destination. To match all objects, specify an empty prefix. Included in the Rule element.	Yes
Destination	Container tag for the destination of a rule. Included in the Rule element.	Yes

Name	Description	Required
Urn	URN of the destination where object metadata is sent. Must be the URN of a StorageGRID endpoint with the following properties: • <code>es</code> must be the third element. • The URN must end with the index and type where the metadata is stored, in the form <code>domain-name/myindex/mytype</code>. Endpoints are configured using the Tenant Manager or Tenant Management API. They take the following form: • <code>arn:aws:es:region:account-ID:domain/mydomain/myindex/mytype</code> • <code>urn:mysite:es:::mydomain/myindex/mytype</code> The endpoint must be configured before the configuration XML is submitted, or configuration will fail with a 404 error. URN is included in the Destination element.	Yes

2. In the Tenant Manager select **STORAGE (S3) > Buckets**.

3. Select the name of the source bucket.

The bucket details page appears.

4. Select **Platform services > Search integration**

5. Select the **Enable search integration** checkbox.

6. Paste the metadata notification configuration into the text box, and select **Save changes**.



Platform services must be enabled for each tenant account by a StorageGRID administrator using the Grid Manager or Management API. Contact your StorageGRID administrator if an error occurs when you save the configuration XML.

7. Verify that the search integration service is configured correctly:

- Add an object to the source bucket that meets the requirements for triggering a metadata notification as specified in the configuration XML.

In the example shown earlier, all objects added to the bucket trigger a metadata notification.

- Confirm that a JSON document that contains the object's metadata and tags was added to the search index specified in the endpoint.

After you finish

As necessary, you can disable search integration for a bucket using either of the following methods:

- Select **STORAGE (S3) > Buckets** and clear the **Enable search integration** checkbox.

- If you are using the S3 API directly, use a DELETE Bucket metadata notification request. See the instructions for implementing S3 client applications.

Example: Metadata notification configuration that applies to all objects

In this example, object metadata for all objects is sent to the same destination.

```
<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>urn:myes:es::sgws-notifications/test1/all</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

Example: Metadata notification configuration with two rules

In this example, object metadata for objects that match the prefix /images is sent to one destination, while object metadata for objects that match the prefix /videos is sent to a second destination.

```
<MetadataNotificationConfiguration>
  <Rule>
    <ID>Images-rule</ID>
    <Status>Enabled</Status>
    <Prefix>/images</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-1:33333333:domain/es-
domain/graphics/imagetype</Urn>
    </Destination>
  </Rule>
  <Rule>
    <ID>Videos-rule</ID>
    <Status>Enabled</Status>
    <Prefix>/videos</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-west-1:22222222:domain/es-
domain/graphics/videotype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

Metadata notification format

When you enable the search integration service for a bucket, a JSON document is generated and sent to the destination endpoint each time object metadata or tags are added, updated, or deleted.

This example shows an example of the JSON that could be generated when an object with the key `SGWS/Tagging.txt` is created in a bucket named `test`. The `test` bucket is not versioned, so the `versionId` tag is empty.

```
{
  "bucket": "test",
  "key": "SGWS/Tagging.txt",
  "versionId": "",
  "accountId": "86928401983529626822",
  "size": 38,
  "md5": "3d6c7634a85436eee06d43415012855",
  "region": "us-east-1",
  "metadata": {
    "age": "25"
  },
  "tags": {
    "color": "yellow"
  }
}
```

Fields included in the JSON document

The document name includes the bucket name, object name, and version ID if present.

Bucket and object information

`bucket`: Name of the bucket

`key`: Object key name

`versionID`: Object version, for objects in versioned buckets

`region`: Bucket region, for example `us-east-1`

System metadata

`size`: Object size (in bytes) as visible to an HTTP client

`md5`: Object hash

User metadata

`metadata`: All user metadata for the object, as key-value pairs

`key:value`

Tags

`tags`: All object tags defined for the object, as key-value pairs

`key:value`

How to view results in Elasticsearch

For tags and user metadata, StorageGRID passes dates and numbers to Elasticsearch as strings or as S3 event notifications. To configure Elasticsearch to interpret these strings as dates or numbers, follow the Elasticsearch instructions for dynamic field mapping and for mapping date formats. Enable the dynamic field mappings on the index before you configure the search integration service. After a document is indexed, you can't edit the document's field types in the index.

Use S3 REST API

S3 REST API supported versions and updates

StorageGRID supports the Simple Storage Service (S3) API, which is implemented as a set of Representational State Transfer (REST) web services.

Support for the S3 REST API enables you to connect service-oriented applications developed for S3 web services with on-premise object storage that uses the StorageGRID system. Minimal changes to a client application's current use of S3 REST API calls are required.

Supported versions

StorageGRID supports the following specific versions of S3 and HTTP.

Item	Version
S3 API specification	Amazon Web Services (AWS) Documentation: Amazon Simple Storage Service API Reference
HTTP	1.1 For more information about HTTP, see HTTP/1.1 (RFCs 7230-35). IETF RFC 2616: Hypertext Transfer Protocol (HTTP/1.1) Note: StorageGRID does not support HTTP/1.1 pipelining.

Updates to S3 REST API support

Release	Comments
11.9	- Added support for pre-calculated SHA-256 checksum values for the following requests and supported headers. You can use this feature to verify the integrity of uploaded objects: - CompleteMultipartUpload: <code>x-amz-checksum-sha256</code> - CreateMultipartUpload: <code>x-amz-checksum-algorithm</code> - GetObject: <code>x-amz-checksum-mode</code> - HeadObject: <code>x-amz-checksum-mode</code> - ListParts - PutObject: <code>x-amz-checksum-sha256</code> - UploadPart: <code>x-amz-checksum-sha256</code> - Added the ability for the grid administrator to control tenant-level retention and Compliance settings. These settings affect S3 Object Lock settings. - Bucket default retention mode and object retention mode: Governance or Compliance, if allowed by the grid administrator. - Bucket default retention period and object Retain Until Date: Must be less than or equal to what is allowed by the maximum retention period set by grid administrator. - Improved support for <code>aws-chunked</code> content encoding and streaming <code>x-amz-content-sha256</code> values. Limitations: - If present, <code>chunk-signature</code> is optional and not validated - If present, <code>x-amz-trailer</code> content is ignored
11.8	Updated the names of S3 operations to match the names used in the Amazon Web Services (AWS) Documentation: Amazon Simple Storage Service API Reference .
11.7	- Added Quick reference: Supported S3 API requests. - Added support for using GOVERNANCE mode with S3 Object Lock. - Added support for the StorageGRID-specific <code>x-ntap-sg-cgr-replication-status</code> response header for GET Object and HEAD Object requests. This header provides an object's replication status for cross-grid replication. - SelectObjectContent requests now support Parquet objects.

Release	Comments
11.6	- Added support for using the <code>partNumber</code> request parameter in GET Object and HEAD Object requests. - Added support for a default retention mode and a default retention period at the bucket level for S3 Object Lock. - Added support for the <code>s3:object-lock-remaining-retention-days</code> policy condition key to set the range of allowable retention periods for your objects. - Changed the maximum <i>recommended</i> size for a single PUT Object operation to 5 GiB (5,368,709,120 bytes). If you have objects that are larger than 5 GiB, use multipart upload instead.
11.5	- Added support for managing bucket encryption. - Added support for S3 Object Lock and deprecated legacy Compliance requests. - Added support for using DELETE Multiple Objects on versioned buckets. - The <code>Content-MD5</code> request header is now correctly supported.
11.4	- Added support for DELETE Bucket tagging, GET Bucket tagging, and PUT Bucket tagging. Cost allocation tags aren't supported. - For buckets created in StorageGRID 11.4, restricting object key names to meet performance best practices is no longer required. - Added support for bucket notifications on the <code>s3:ObjectRestore:Post</code> event type. - AWS size limits for multipart parts are now enforced. Each part in a multipart upload must be between 5 MiB and 5 GiB. The last part can be smaller than 5 MiB. - Added support for TLS 1.3
11.3	- Added support for server-side encryption of object data with customer-provided keys (SSE-C). - Added support for DELETE, GET, and PUT Bucket lifecycle operations (Expiration action only) and for the <code>x-amz-expiration</code> response header. - Updated PUT Object, PUT Object - Copy, and Multipart Upload to describe the impact of ILM rules that use synchronous placement at ingest. - TLS 1.1 ciphers are no longer supported.
11.2	Added support for POST Object restore for use with Cloud Storage Pools. Added support for using the AWS syntax for ARN, policy condition keys, and policy variables in group and bucket policies. Existing group and bucket policies that use the StorageGRID syntax will continue to be supported. Note: Uses of ARN/URN in other configuration JSON/XML, including those used in custom StorageGRID features, have not changed.
11.1	Added support for cross-origin resource sharing (CORS), HTTP for S3 client connections to grid nodes, and compliance settings on buckets.

Release	Comments
11.0	Added support for configuring platform services (CloudMirror replication, notifications, and Elasticsearch search integration) for buckets. Also added support for object tagging location constraints for buckets, and the Available consistency.
10.4	Added support for ILM scanning changes to versioning, Endpoint Domain Names page updates, conditions and variables in policies, policy examples, and the PutOverwriteObject permission.
10.3	Added support for versioning.
10.2	Added support for group and bucket access policies, and for multipart copy (Upload Part - Copy).
10.1	Added support for multipart upload, virtual hosted-style requests, and v4 authentication.
10.0	Initial support of the S3 REST API by the StorageGRID system. The currently supported version of the <i>Simple Storage Service API Reference</i> is 2006-03-01.

Quick reference: Supported S3 API requests

This page summarizes how StorageGRID supports Amazon Simple Storage Service (S3) APIs.

This page includes only the S3 operations that are supported by StorageGRID.



To see the AWS documentation for each operation, select the link in the heading.

Common URI query parameters and request headers

Unless noted, the following common URI query parameters are supported:

- `versionId` (as required for object operations)

Unless noted, the following common request headers are supported:

- `Authorization`
- `Connection`
- `Content-Length`
- `Content-MD5`
- `Content-Type`
- `Date`
- `Expect`
- `Host`

- x-amz-date

Related information

- [S3 REST API implementation details](#)
- [Amazon Simple Storage Service API Reference: Common Request Headers](#)

AbortMultipartUpload

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus this additional URI query parameter:

- uploadId

Request body

None

StorageGRID documentation

[Operations for multipart uploads](#)

CompleteMultipartUpload

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus this additional URI query parameter:

- uploadId
- x-amz-checksum-sha256

Request body XML tags

StorageGRID supports these request body XML tags:

- ChecksumSHA256
- CompleteMultipartUpload
- ETag
- Part
- PartNumber

StorageGRID documentation

[CompleteMultipartUpload](#)

CopyObject

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional headers:

- x-amz-copy-source
- x-amz-copy-source-if-match

- x-amz-copy-source-if-modified-since
- x-amz-copy-source-if-none-match
- x-amz-copy-source-if-unmodified-since
- x-amz-copy-source-server-side-encryption-customer-algorithm
- x-amz-copy-source-server-side-encryption-customer-key
- x-amz-copy-source-server-side-encryption-customer-key-MD5
- x-amz-metadata-directive
- x-amz-object-lock-legal-hold
- x-amz-object-lock-mode
- x-amz-object-lock-retain-until-date
- x-amz-server-side-encryption
- x-amz-server-side-encryption-customer-algorithm
- x-amz-server-side-encryption-customer-key
- x-amz-server-side-encryption-customer-key-MD5
- x-amz-storage-class
- x-amz-tagging
- x-amz-tagging-directive
- x-amz-meta-`<metadata-name>`

Request body

None

StorageGRID documentation

[CopyObject](#)

CreateBucket

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional headers:

- x-amz-bucket-object-lock-enabled

Request body

StorageGRID supports all request body parameters defined by the Amazon S3 REST API at the time of implementation.

StorageGRID documentation

[Operations on buckets](#)

CreateMultipartUpload

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional headers:

- Cache-Control
- Content-Disposition
- Content-Encoding
- Content-Language
- Expires
- x-amz-checksum-algorithm
- x-amz-server-side-encryption
- x-amz-storage-class
- x-amz-server-side-encryption-customer-algorithm
- x-amz-server-side-encryption-customer-key
- x-amz-server-side-encryption-customer-key-MD5
- x-amz-tagging
- x-amz-object-lock-mode
- x-amz-object-lock-retain-until-date
- x-amz-object-lock-legal-hold
- x-amz-meta-`<metadata-name>`

Request body

None

StorageGRID documentation

[CreateMultipartUpload](#)

DeleteBucket

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

StorageGRID documentation

[Operations on buckets](#)

DeleteBucketCors

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

DeleteBucketEncryption

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

DeleteBucketLifecycle

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

- [Operations on buckets](#)
- [Create S3 lifecycle configuration](#)

DeleteBucketPolicy

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

DeleteBucketReplication

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

DeleteBucketTagging

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

DeleteObject

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus this additional request header:

- `x-amz-bypass-governance-retention`

Request body

None

StorageGRID documentation

[Operations on objects](#)

DeleteObjects

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus this additional request header:

- `x-amz-bypass-governance-retention`

Request body

StorageGRID supports all request body parameters defined by the Amazon S3 REST API at the time of implementation.

StorageGRID documentation

[Operations on objects](#)

DeleteObjectTagging

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on objects](#)

GetBucketAcl

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

GetBucketCors

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

GetBucketEncryption

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

GetBucketLifecycleConfiguration

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

- [Operations on buckets](#)
- [Create S3 lifecycle configuration](#)

GetBucketLocation

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

GetBucketNotificationConfiguration

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

GetBucketPolicy

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

GetBucketReplication

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

GetBucketTagging

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

GetBucketVersioning

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

GetObject

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional URI query parameters:

- x-amz-checksum-mode
- partNumber
- response-cache-control
- response-content-disposition
- response-content-encoding
- response-content-language
- response-content-type
- response-expires

And these additional request headers:

- Range
- x-amz-server-side-encryption-customer-algorithm
- x-amz-server-side-encryption-customer-key
- x-amz-server-side-encryption-customer-key-MD5
- If-Match
- If-Modified-Since
- If-None-Match
- If-Unmodified-Since

Request body

None

StorageGRID documentation

[GetObject](#)

GetObjectAcl

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on objects](#)

GetObjectLegalHold

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Use S3 REST API to configure S3 Object Lock](#)

GetObjectLockConfiguration

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Use S3 REST API to configure S3 Object Lock](#)

GetObjectRetention

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Use S3 REST API to configure S3 Object Lock](#)

GetObjectTagging

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on objects](#)

HeadBucket

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on buckets](#)

HeadObject

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional headers:

- x-amz-checksum-mode
- x-amz-server-side-encryption-customer-algorithm
- x-amz-server-side-encryption-customer-key
- x-amz-server-side-encryption-customer-key-MD5
- If-Match
- If-Modified-Since
- If-None-Match
- If-Unmodified-Since
- Range

Request body

None

StorageGRID documentation

[HeadObject](#)

ListBuckets

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

None

StorageGRID documentation

[Operations on the service > ListBuckets](#)

ListMultipartUploads

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional parameters:

- encoding-type
- key-marker
- max-uploads
- prefix
- upload-id-marker

Request body

None

StorageGRID documentation

[ListMultipartUploads](#)

ListObjects

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional parameters:

- delimiter
- encoding-type
- marker
- max-keys
- prefix

Request body

None

StorageGRID documentation

[Operations on buckets](#)

ListObjectsV2

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional parameters:

- continuation-token
- delimiter
- encoding-type
- fetch-owner
- max-keys
- prefix

- start-after

Request body

None

StorageGRID documentation

[Operations on buckets](#)

ListObjectVersions**URI query parameters and request headers**

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional parameters:

- delimiter
- encoding-type
- key-marker
- max-keys
- prefix
- version-id-marker

Request body

None

StorageGRID documentation

[Operations on buckets](#)

ListParts**URI query parameters and request headers**

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional parameters:

- max-parts
- part-number-marker
- uploadId

Request body

None

StorageGRID documentation

[ListMultipartUploads](#)

PutBucketCors**URI query parameters and request headers**

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

StorageGRID supports all request body parameters defined by the Amazon S3 REST API at the time of

implementation.

StorageGRID documentation

[Operations on buckets](#)

PutBucketEncryption

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body XML tags

StorageGRID supports these request body XML tags:

- ApplyServerSideEncryptionByDefault
- Rule
- ServerSideEncryptionConfiguration
- SSEAlgorithm

StorageGRID documentation

[Operations on buckets](#)

PutBucketLifecycleConfiguration

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body XML tags

StorageGRID supports these request body XML tags:

- And
- Days
- Expiration
- ExpiredObjectDeleteMarker
- Filter
- ID
- Key
- LifecycleConfiguration
- NewerNoncurrentVersions
- NoncurrentDays
- NoncurrentVersionExpiration
- Prefix
- Rule
- Status

- Tag
- Value

StorageGRID documentation

- [Operations on buckets](#)
- [Create S3 lifecycle configuration](#)

PutBucketNotificationConfiguration

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body XML tags

StorageGRID supports these request body XML tags:

- Event
- Filter
- FilterRule
- Id
- Name
- NotificationConfiguration
- Prefix
- S3Key
- Suffix
- Topic
- TopicConfiguration
- Value

StorageGRID documentation

[Operations on buckets](#)

PutBucketPolicy

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

For details about the supported JSON body fields, see [Use bucket and group access policies](#).

PutBucketReplication

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body XML tags

- Bucket
- Destination
- Prefix
- ReplicationConfiguration
- Rule
- Status
- StorageClass

StorageGRID documentation

[Operations on buckets](#)

PutBucketTagging

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

StorageGRID supports all request body parameters defined by the Amazon S3 REST API at the time of implementation.

StorageGRID documentation

[Operations on buckets](#)

PutBucketVersioning

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body parameters

StorageGRID supports these request body parameters:

- VersioningConfiguration
- Status

StorageGRID documentation

[Operations on buckets](#)

PutObject

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional headers:

- Cache-Control
- Content-Disposition
- Content-Encoding

- Content-Language
- Expires
- x-amz-checksum-sha256
- x-amz-server-side-encryption
- x-amz-storage-class
- x-amz-server-side-encryption-customer-algorithm
- x-amz-server-side-encryption-customer-key
- x-amz-server-side-encryption-customer-key-MD5
- x-amz-tagging
- x-amz-object-lock-mode
- x-amz-object-lock-retain-until-date
- x-amz-object-lock-legal-hold
- x-amz-meta-`<metadata-name>`

Request body

- Binary data of the object

StorageGRID documentation

[PutObject](#)

PutObjectLegalHold

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

StorageGRID supports all request body parameters defined by the Amazon S3 REST API at the time of implementation.

StorageGRID documentation

[Use S3 REST API to configure S3 Object Lock](#)

PutObjectLockConfiguration

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

StorageGRID supports all request body parameters defined by the Amazon S3 REST API at the time of implementation.

StorageGRID documentation

[Use S3 REST API to configure S3 Object Lock](#)

PutObjectRetention

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus this additional header:

- `x-amz-bypass-governance-retention`

Request body

StorageGRID supports all request body parameters defined by the Amazon S3 REST API at the time of implementation.

StorageGRID documentation

[Use S3 REST API to configure S3 Object Lock](#)

PutObjectTagging

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

StorageGRID supports all request body parameters defined by the Amazon S3 REST API at the time of implementation.

StorageGRID documentation

[Operations on objects](#)

RestoreObject

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

For details about the supported body fields, see [RestoreObject](#).

SelectObjectContent

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request.

Request body

For details about the supported body fields, see the following:

- [Use S3 Select](#)
- [SelectObjectContent](#)

UploadPart

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional URI query parameters:

- partNumber
- uploadId

And these additional request headers:

- x-amz-checksum-sha256
- x-amz-server-side-encryption-customer-algorithm
- x-amz-server-side-encryption-customer-key
- x-amz-server-side-encryption-customer-key-MD5

Request body

- Binary data of the part

StorageGRID documentation

[UploadPart](#)

[UploadPartCopy](#)

URI query parameters and request headers

StorageGRID supports all [common parameters and headers](#) for this request, plus these additional URI query parameters:

- partNumber
- uploadId

And these additional request headers:

- x-amz-copy-source
- x-amz-copy-source-if-match
- x-amz-copy-source-if-modified-since
- x-amz-copy-source-if-none-match
- x-amz-copy-source-if-unmodified-since
- x-amz-copy-source-range
- x-amz-server-side-encryption-customer-algorithm
- x-amz-server-side-encryption-customer-key
- x-amz-server-side-encryption-customer-key-MD5
- x-amz-copy-source-server-side-encryption-customer-algorithm
- x-amz-copy-source-server-side-encryption-customer-key
- x-amz-copy-source-server-side-encryption-customer-key-MD5

Request body

None

Test S3 REST API configuration

You can use the Amazon Web Services Command Line Interface (AWS CLI) to test your connection to the system and to verify that you can read and write objects.

Before you begin

- You have downloaded and installed the AWS CLI from aws.amazon.com/cli.
- Optionally, you have [created a load balancer endpoint](#). Otherwise, you know the IP address of the Storage Node you want to connect to and the port number to use. See [IP addresses and ports for client connections](#).
- You have [created an S3 tenant account](#).
- You have signed in to the tenant and [created an access key](#).

For details on these steps, see [Configure client connections](#).

Steps

1. Configure the AWS CLI settings to use the account you created in the StorageGRID system:
 - a. Enter configuration mode: `aws configure`
 - b. Enter the access key ID for the account you created.
 - c. Enter the secret access key for the account you created.
 - d. Enter the default region to use. For example, `us-east-1`.
 - e. Enter the default output format to use, or press **Enter** to select JSON.
2. Create a bucket.

This example assumes you configured a load balancer endpoint to use IP address 10.96.101.17 and port 10443.

```
aws s3api --endpoint-url https://10.96.101.17:10443
--no-verify-ssl create-bucket --bucket testbucket
```

If the bucket is created successfully, the location of the bucket is returned, as seen in the following example:

```
"Location": "/testbucket"
```

3. Upload an object.

```
aws s3api --endpoint-url https://10.96.101.17:10443 --no-verify-ssl
put-object --bucket testbucket --key s3.pdf --body C:\s3-
test\upload\s3.pdf
```

If the object is uploaded successfully, an Etag is returned which is a hash of the object data.

4. List the contents of the bucket to verify that the object was uploaded.

```
aws s3api --endpoint-url https://10.96.101.17:10443 --no-verify-ssl  
list-objects --bucket testbucket
```

5. Delete the object.

```
aws s3api --endpoint-url https://10.96.101.17:10443 --no-verify-ssl  
delete-object --bucket testbucket --key s3.pdf
```

6. Delete the bucket.

```
aws s3api --endpoint-url https://10.96.101.17:10443 --no-verify-ssl  
delete-bucket --bucket testbucket
```

How StorageGRID implements S3 REST API

Conflicting client requests

Conflicting client requests, such as two clients writing to the same key, are resolved on a "latest-wins" basis.

The timing for the "latest-wins" evaluation is based on when the StorageGRID system completes a given request, and not on when S3 clients begin an operation.

Consistency values

Consistency provides a balance between the availability of the objects and the consistency of those objects across different Storage Nodes and sites. You can change the consistency as required by your application.

By default, StorageGRID guarantees read-after-write consistency for newly created objects. Any GET following a successfully completed PUT will be able to read the newly written data. Overwrites of existing objects, metadata updates, and deletes are eventually consistent. Overwrites generally take seconds or minutes to propagate, but can take up to 15 days.

If you want to perform object operations at a different consistency, you can:

- Specify a consistency for [each bucket](#).
- Specify a consistency for [each API operation](#).
- Change the default grid-wide consistency by performing one of the following tasks:
 - In the Grid Manager, go to **CONFIGURATION > System > Storage settings > Default consistency**.
 - [Use the grid-config endpoint of the Grid Management private API](#).



A change to the grid-wide consistency applies only to buckets created after the setting was changed. To determine the details of a change, see the audit log located at `/var/local/log` (search for **consistencyLevel**).

Consistency values

The consistency affects how the metadata that StorageGRID uses to track objects is distributed between nodes, and therefore the availability of objects for client requests.

You can set the consistency for a bucket or an API operation to one of the following values:

- **All**: All nodes receive the data immediately, or the request will fail.
- **Strong-global**: Guarantees read-after-write consistency for all client requests across all sites.
- **Strong-site**: Guarantees read-after-write consistency for all client requests within a site.
- **Read-after-new-write**: (Default) Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.
- **Available**: Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that don't exist). Not supported for S3 FabricPool buckets.

Use "Read-after-new-write" and "Available" consistency

When a HEAD or GET operation uses the "Read-after-new-write" consistency, StorageGRID performs the lookup in multiple steps, as follows:

- It first looks up the object using a low consistency.
- If that lookup fails, it repeats the lookup at the next consistency value until it reaches a consistency equivalent to the behavior for strong-global.

If a HEAD or GET operation uses the "Read-after-new-write" consistency but the object does not exist, the object lookup will always reach a consistency equivalent to the behavior for strong-global. Because this consistency requires multiple copies of the object metadata to be available at each site, you can receive a high number of 500 Internal Server errors if two or more Storage Nodes at the same site are unavailable.

Unless you require consistency guarantees similar to Amazon S3, you can prevent these errors for HEAD and GET operations by setting the consistency to "Available." When a HEAD or GET operation uses the "Available" consistency, StorageGRID provides eventual consistency only. It does not retry a failed operation at increasing consistency, so it does not require that multiple copies of the object metadata be available.

Specify consistency for API operation

To set the consistency for an individual API operation, the consistency values must be supported for the operation, and you must specify the consistency in the request header. This example sets the consistency to "Strong-site" for a GetObject operation.


```
GET /bucket/object HTTP/1.1
Date: date
Authorization: authorization name
Host: host
Consistency-Control: strong-site
```



You must use the same consistency for both the PutObject and GetObject operations.

Specify consistency for bucket

To set the consistency for bucket, you can use the StorageGRID [PUT Bucket consistency](#) request. Or you can [change a bucket's consistency](#) from the Tenant Manager.

When setting the consistency for a bucket, be aware of the following:

- Setting the consistency for a bucket determines which consistency is used for S3 operations performed on the objects in the bucket or on the bucket configuration. It does not affect operations on the bucket itself.
- The consistency for an individual API operation overrides the consistency for the bucket.
- In general, buckets should use the default consistency, "Read-after-new-write." If requests aren't working correctly, change the application client behavior if possible. Or, configure the client to specify the consistency for each API request. Set the consistency at the bucket level only as a last resort.

How consistency and ILM rules interact to affect data protection

Both your choice of consistency and your ILM rule affect how objects are protected. These settings can interact.

For example, the consistency used when an object is stored affects the initial placement of object metadata, while the ingest behavior selected for the ILM rule affects the initial placement of object copies. Because StorageGRID requires access to both an object's metadata and its data to fulfill client requests, selecting matching levels of protection for the consistency and ingest behavior can provide better initial data protection and more predictable system responses.

The following [ingest options](#) are available for ILM rules:

Dual commit

StorageGRID immediately makes interim copies of the object and returns success to the client. Copies specified in the ILM rule are made when possible.

Strict

All copies specified in the ILM rule must be made before success is returned to the client.

Balanced

StorageGRID attempts to make all copies specified in the ILM rule at ingest; if this is not possible, interim copies are made and success is returned to the client. The copies specified in the ILM rule are made when possible.

Example of how the consistency and ILM rule can interact

Suppose you have a two-site grid with the following ILM rule and the following consistency:

- **ILM rule:** Create two object copies, one at the local site and one at a remote site. Use Strict ingest behavior.
- **consistency:** Strong-global (object metadata is immediately distributed to all sites).

When a client stores an object to the grid, StorageGRID makes both object copies and distributes metadata to both sites before returning success to the client.

The object is fully protected against loss at the time of the ingest successful message. For example, if the local site is lost shortly after ingest, copies of both the object data and the object metadata still exist at the remote site. The object is fully retrievable.

If you instead used the same ILM rule and the strong-site consistency, the client might receive a success message after object data is replicated to the remote site but before object metadata is distributed there. In this case, the level of protection of object metadata does not match the level of protection for object data. If the local site is lost shortly after ingest, object metadata is lost. The object can't be retrieved.

The inter-relationship between consistency and ILM rules can be complex. Contact NetApp if you need assistance.

Object versioning

You can set the versioning state of a bucket if you want to retain multiple versions of each object. Enabling versioning for a bucket can help protect against accidental deletion of objects and enables you to retrieve and restore earlier versions of an object.

The StorageGRID system implements versioning with support for most features, and with some limitations. StorageGRID supports up to 10,000 versions of each object.

Object versioning can be combined with StorageGRID information lifecycle management (ILM) or with S3 bucket lifecycle configuration. You must explicitly enable versioning for each bucket. When versioning is enabled for a bucket, each object added to the bucket is assigned a version ID, which is generated by the StorageGRID system.

Using MFA (multi-factor authentication) Delete is not supported.



Versioning can be enabled only on buckets created with StorageGRID version 10.3 or later.

ILM and versioning

ILM policies are applied to each version of an object. An ILM scanning process continuously scans all objects and re-evaluates them against the current ILM policy. Any changes you make to ILM policies are applied to all previously ingested objects. This includes previously ingested versions if versioning is enabled. ILM scanning applies new ILM changes to previously ingested objects.

For S3 objects in versioning-enabled buckets, versioning support allows you to create ILM rules that use "Noncurrent time" as the Reference time (select **Yes** for the question, "Apply this rule to older object versions only?" in [Step 1 of the Create an ILM rule wizard](#)). When an object is updated, its previous versions become noncurrent. Using a "Noncurrent time" filter allows you to create policies that reduce the storage impact of previous versions of objects.



When you upload a new version of an object using a multipart upload operation, the noncurrent time for the original version of the object reflects when the multipart upload was created for the new version, not when the multipart upload was completed. In limited cases, the noncurrent time for the original version might be hours or days earlier than the time for the current version.

Related information

- [How S3 versioned objects are deleted](#)
- [ILM rules and policies for S3 versioned objects \(Example 4\)](#).

Use S3 REST API to configure S3 Object Lock

If the global S3 Object Lock setting is enabled for your StorageGRID system, you can create buckets with S3 Object Lock enabled. You can specify default retention for each bucket or retention settings for each object version.

How to enable S3 Object Lock for a bucket

If the global S3 Object Lock setting is enabled for your StorageGRID system, you can optionally enable S3 Object Lock when you create each bucket.

S3 Object Lock is a permanent setting that can only be enabled when you create a bucket. You can't add or disable S3 Object Lock after a bucket is created.

To enable S3 Object Lock for a bucket, use either of these methods:

- Create the bucket using the Tenant Manager. See [Create S3 bucket](#).
- Create the bucket using a CreateBucket request with the `x-amz-bucket-object-lock-enabled` request header. See [Operations on buckets](#).

S3 Object Lock requires bucket versioning, which is enabled automatically when the bucket is created. You can't suspend versioning for the bucket. See [Object versioning](#).

Default retention settings for a bucket

When S3 Object Lock is enabled for a bucket, you can optionally enable default retention for the bucket and specify a default retention mode and default retention period.

Default retention mode

- In COMPLIANCE mode:
 - The object can't be deleted until its retain-until-date is reached.
 - The object's retain-until-date can be increased, but it can't be decreased.
 - The object's retain-until-date can't be removed until that date is reached.
- In GOVERNANCE mode:
 - Users with the `s3:BypassGovernanceRetention` permission can use the `x-amz-bypass-governance-retention: true` request header to bypass retention settings.
 - These users can delete an object version before its retain-until-date is reached.
 - These users can increase, decrease, or remove an object's retain-until-date.

Default retention period

Each bucket can have a default retention period specified in years or days.

How to set default retention for a bucket

To set the default retention for a bucket, use either of these methods:

- Manage bucket settings from the Tenant Manager. See [Create an S3 bucket](#) and [Update S3 Object Lock default retention](#).
- Issue a `PutObjectLockConfiguration` request for the bucket to specify the default mode and default number of days or years.

PutObjectLockConfiguration

The `PutObjectLockConfiguration` request allows you to set and modify the default retention mode and default retention period for a bucket that has S3 Object Lock enabled. You can also remove previously configured default retention settings.

When new object versions are ingested to the bucket, the default retention mode is applied if `x-amz-object-lock-mode` and `x-amz-object-lock-retain-until-date` aren't specified. The default retention period is used to calculate the `retain-until-date` if `x-amz-object-lock-retain-until-date` is not specified.

If the default retention period is modified after ingest of an object version, the `retain-until-date` of the object version remains the same and is not recalculated using the new default retention period.

You must have the `s3:PutBucketObjectLockConfiguration` permission, or be account root, to complete this operation.

The `Content-MD5` request header must be specified in the PUT request.

Request example

This example enables S3 Object Lock for a bucket and sets the default retention mode to COMPLIANCE and the default retention period to 6 years.

```
PUT /bucket?object-lock HTTP/1.1
Accept-Encoding: identity
Content-Length: 308
Host: host
Content-MD5: request header
User-Agent: s3sign/1.0.0 requests/2.24.0 python/3.8.2
X-Amz-Date: date
X-Amz-Content-SHA256: authorization-string
Authorization: authorization-string

<ObjectLockConfiguration>
  <ObjectLockEnabled>Enabled</ObjectLockEnabled>
  <Rule>
    <DefaultRetention>
      <Mode>COMPLIANCE</Mode>
      <Years>6</Years>
    </DefaultRetention>
  </Rule>
</ObjectLockConfiguration>
```

How to determine the default retention for a bucket

To determine if S3 Object Lock is enabled for a bucket and to see the default retention mode and retention period, use either of these methods:

- View the bucket in the Tenant Manager. See [View S3 buckets](#).
- Issue a `GetObjectLockConfiguration` request.

GetObjectLockConfiguration

The `GetObjectLockConfiguration` request allows you to determine if S3 Object Lock is enabled for a bucket and, if it is enabled, see if there is a default retention mode and retention period configured for the bucket.

When new object versions are ingested to the bucket, the default retention mode is applied if `x-amz-object-lock-mode` is not specified. The default retention period is used to calculate the `retain-until-date` if `x-amz-object-lock-retain-until-date` is not specified.

You must have the `s3:GetBucketObjectLockConfiguration` permission, or be account root, to complete this operation.

Request example

```
GET /bucket?object-lock HTTP/1.1
Host: host
Accept-Encoding: identity
User-Agent: aws-cli/1.18.106 Python/3.8.2 Linux/4.4.0-18362-Microsoft
botocore/1.17.29
x-amz-date: date
x-amz-content-sha256: authorization-string
Authorization: authorization-string
```

Response example

```
HTTP/1.1 200 OK
x-amz-id-2:
iVmcB7OXXJRkRH1FiVq1151/T24gRfpwpuZrEG11Bb9ImOMAAe98oxSpX1knabA0LTvBYJpSIX
k=
x-amz-request-id: B34E94CACB2CEF6D
Date: Fri, 04 Sep 2020 22:47:09 GMT
Transfer-Encoding: chunked
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<ObjectLockConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <ObjectLockEnabled>Enabled</ObjectLockEnabled>
  <Rule>
    <DefaultRetention>
      <Mode>COMPLIANCE</Mode>
      <Years>6</Years>
    </DefaultRetention>
  </Rule>
</ObjectLockConfiguration>
```

How to specify retention settings for an object

A bucket with S3 Object Lock enabled can contain a combination of objects with and without S3 Object Lock retention settings.

Object-level retention settings are specified using the S3 REST API. The retention settings for an object override any default retention settings for the bucket.

You can specify the following settings for each object:

- **Retention mode:** Either COMPLIANCE or GOVERNANCE.
- **Retain-until-date:** A date specifying how long the object version must be retained by StorageGRID.
 - In COMPLIANCE mode, if the retain-until-date is in the future, the object can be retrieved, but it can't be modified or deleted. The retain-until-date can be increased, but this date can't be decreased or removed.

- In GOVERNANCE mode, users with special permission can bypass the retain-until-date setting. They can delete an object version before its retention period has elapsed. They can also increase, decrease, or even remove the retain-until-date.
- **Legal hold:** Applying a legal hold to an object version immediately locks that object. For example, you might need to put a legal hold on an object that is related to an investigation or legal dispute. A legal hold has no expiration date, but remains in place until it is explicitly removed.

The legal hold setting for an object is independent of the retention mode and the retain-until-date. If an object version is under a legal hold, no one can delete that version.

To specify S3 Object Lock settings when adding an object version to a bucket, issue a [PutObject](#), [CopyObject](#), or [CreateMultipartUpload](#) request.

You can use the following:

- `x-amz-object-lock-mode`, which can be COMPLIANCE or GOVERNANCE (case sensitive).



If you specify `x-amz-object-lock-mode`, you must also specify `x-amz-object-lock-retain-until-date`.

- `x-amz-object-lock-retain-until-date`
 - The retain-until-date value must be in the format `2020-08-10T21:46:00Z`. Fractional seconds are allowed, but only 3 decimal digits are preserved (milliseconds precision). Other ISO 8601 formats aren't allowed.
 - The retain-until-date must be in the future.
- `x-amz-object-lock-legal-hold`

If legal hold is ON (case-sensitive), the object is placed under a legal hold. If legal hold is OFF, no legal hold is placed. Any other value results in a 400 Bad Request (InvalidArgument) error.

If you use any of these request headers, be aware of these restrictions:

- The `Content-MD5` request header is required if any `x-amz-object-lock-*` request header is present in the `PutObject` request. `Content-MD5` is not required for `CopyObject` or `CreateMultipartUpload`.
- If the bucket does not have S3 Object Lock enabled and a `x-amz-object-lock-*` request header is present, a 400 Bad Request (InvalidRequest) error is returned.
- The `PutObject` request supports the use of `x-amz-storage-class: REDUCED_REDUNDANCY` to match AWS behavior. However, when an object is ingested into a bucket with S3 Object Lock enabled, StorageGRID will always perform a dual-commit ingest.
- A subsequent GET or `HeadObject` version response will include the headers `x-amz-object-lock-mode`, `x-amz-object-lock-retain-until-date`, and `x-amz-object-lock-legal-hold`, if configured and if the request sender has the correct `s3:Get*` permissions.

You can use the `s3:object-lock-remaining-retention-days` policy condition key to limit the minimum and maximum allowable retention periods for your objects.

How to update retention settings for an object

If you need to update the legal hold or retention settings for an existing object version, you can perform the following object subresource operations:

- `PutObjectLegalHold`

If the new legal-hold value is ON, the object is placed under a legal hold. If the legal-hold value is OFF, the legal hold is lifted.

- `PutObjectRetention`
 - The mode value can be COMPLIANCE or GOVERNANCE (case sensitive).
 - The retain-until-date value must be in the format `2020-08-10T21:46:00Z`. Fractional seconds are allowed, but only 3 decimal digits are preserved (milliseconds precision). Other ISO 8601 formats aren't allowed.
 - If an object version has an existing retain-until-date, you can only increase it. The new value must be in the future.

How to use GOVERNANCE mode

Users who have the `s3:BypassGovernanceRetention` permission can bypass the active retention settings of an object that uses GOVERNANCE mode. Any DELETE or `PutObjectRetention` operations must include the `x-amz-bypass-governance-retention:true` request header. These users can perform these additional operations:

- Perform `DeleteObject` or `DeleteObjects` operations to delete an object version before its retention period has elapsed.

Objects that are under a legal hold can't be deleted. Legal hold must be OFF.

- Perform `PutObjectRetention` operations that change an object version's mode from GOVERNANCE to COMPLIANCE before the object's retention period has elapsed.

Changing the mode from COMPLIANCE to GOVERNANCE is never allowed.

- Perform `PutObjectRetention` operations to increase, decrease, or remove an object version's retention period.

Related information

- [Manage objects with S3 Object Lock](#)
- [Use S3 Object Lock to retain objects](#)
- [Amazon Simple Storage Service User Guide: Locking Objects](#)

Create S3 lifecycle configuration

You can create an S3 lifecycle configuration to control when specific objects are deleted from the StorageGRID system.

The simple example in this section illustrates how an S3 lifecycle configuration can control when certain objects are deleted (expired) from specific S3 buckets. The example in this section is for illustration purposes only. For complete details on creating S3 lifecycle configurations, see [Amazon Simple Storage Service User Guide: Object lifecycle management](#). Note that StorageGRID only supports Expiration actions; it does not

support Transition actions.

What lifecycle configuration is

A lifecycle configuration is a set of rules that are applied to the objects in specific S3 buckets. Each rule specifies which objects are affected and when those objects will expire (on a specific date or after some number of days).

StorageGRID supports up to 1,000 lifecycle rules in a lifecycle configuration. Each rule can include the following XML elements:

- **Expiration:** Delete an object when a specified date is reached or when a specified number of days is reached, starting from when the object was ingested.
- **NoncurrentVersionExpiration:** Delete an object when a specified number of days is reached, starting from when the object became noncurrent.
- **Filter (Prefix, Tag)**
- **Status**
- **ID**

Each object follows the retention settings of either an S3 bucket lifecycle or an ILM policy. When an S3 bucket lifecycle is configured, the lifecycle expiration actions override the ILM policy for objects matching the bucket lifecycle filter. Objects that do not match the bucket lifecycle filter use the retention settings of the ILM policy. If an object matches a bucket lifecycle filter and no expiration actions are explicitly specified, the retention settings of the ILM policy are not used and it is implied that object versions are retained forever. See [Example priorities for S3 bucket lifecycle and ILM policy](#).

As a result, an object might be removed from the grid even though the placement instructions in an ILM rule still apply to the object. Or, an object might be retained on the grid even after any ILM placement instructions for the object have lapsed. For details, see [How ILM operates throughout an object's life](#).



Bucket lifecycle configuration can be used with buckets that have S3 Object Lock enabled, but bucket lifecycle configuration is not supported for legacy Compliant buckets.

StorageGRID supports the use of the following bucket operations to manage lifecycle configurations:

- `DeleteBucketLifecycle`
- `GetBucketLifecycleConfiguration`
- `PutBucketLifecycleConfiguration`

Create lifecycle configuration

As the first step in creating a lifecycle configuration, you create a JSON file that includes one or more rules. For example, this JSON file includes three rules, as follows:

1. Rule 1 applies only to objects that match the prefix `category1/` and that have a `key2` value of `tag2`. The `Expiration` parameter specifies that objects matching the filter will expire at midnight on 22 August 2020.
2. Rule 2 applies only to objects that match the prefix `category2/`. The `Expiration` parameter specifies that objects matching the filter will expire 100 days after they are ingested.



Rules that specify a number of days are relative to when the object was ingested. If the current date exceeds the ingest date plus the number of days, some objects might be removed from the bucket as soon as the lifecycle configuration is applied.

3. Rule 3 applies only to objects that match the prefix `category3/`. The `Expiration` parameter specifies that any noncurrent versions of matching objects will expire 50 days after they become noncurrent.

```

{
  "Rules": [
    {
      "ID": "rule1",
      "Filter": {
        "And": {
          "Prefix": "category1/",
          "Tags": [
            {
              "Key": "key2",
              "Value": "tag2"
            }
          ]
        }
      },
      "Expiration": {
        "Date": "2020-08-22T00:00:00Z"
      },
      "Status": "Enabled"
    },
    {
      "ID": "rule2",
      "Filter": {
        "Prefix": "category2/"
      },
      "Expiration": {
        "Days": 100
      },
      "Status": "Enabled"
    },
    {
      "ID": "rule3",
      "Filter": {
        "Prefix": "category3/"
      },
      "NoncurrentVersionExpiration": {
        "NoncurrentDays": 50
      },
      "Status": "Enabled"
    }
  ]
}

```

Apply lifecycle configuration to bucket

After you have created the lifecycle configuration file, you apply it to a bucket by issuing a `PutBucketLifecycleConfiguration` request.

This request applies the lifecycle configuration in the example file to objects in a bucket named `testbucket`.

```
aws s3api --endpoint-url <StorageGRID endpoint> put-bucket-lifecycle-configuration
--bucket testbucket --lifecycle-configuration file://bktjson.json
```

To validate that a lifecycle configuration was successfully applied to the bucket, issue a `GetBucketLifecycleConfiguration` request. For example:

```
aws s3api --endpoint-url <StorageGRID endpoint> get-bucket-lifecycle-configuration
--bucket testbucket
```

A successful response lists the lifecycle configuration you just applied.

Validate that bucket lifecycle expiration applies to object

You can determine if an expiration rule in the lifecycle configuration applies to a specific object when issuing a `PutObject`, `HeadObject`, or `GetObject` request. If a rule applies, the response includes an `Expiration` parameter that indicates when the object expires and which expiration rule was matched.



Because bucket lifecycle overrides ILM, the `expiry-date` shown is the actual date the object will be deleted. For details, see [How object retention is determined](#).

For example, this `PutObject` request was issued on 22 Jun 2020 and places an object in the `testbucket` bucket.

```
aws s3api --endpoint-url <StorageGRID endpoint> put-object
--bucket testbucket --key obj2test2 --body bktjson.json
```

The success response indicates that the object will expire in 100 days (01 Oct 2020) and that it matched Rule 2 of the lifecycle configuration.

```
{
  *Expiration: "expiry-date=\"Thu, 01 Oct 2020 09:07:49 GMT\", rule-id=\"rule2\"",
  ETag: "\"9762f8a803bc34f5340579d4446076f7\""
}
```

For example, this `HeadObject` request was used to get metadata for the same object in the `testbucket` bucket.

```
aws s3api --endpoint-url <StorageGRID endpoint> head-object
--bucket testbucket --key obj2test2
```

The success response includes the object's metadata and indicates that the object will expire in 100 days and that it matched Rule 2.

```
{
  "AcceptRanges": "bytes",
  *Expiration: "expiry-date=\"Thu, 01 Oct 2020 09:07:48 GMT\", rule-
id=\"rule2\"",
  "LastModified": "2020-06-23T09:07:48+00:00",
  "ContentLength": 921,
  "ETag": "\"9762f8a803bc34f5340579d4446076f7\""
  "ContentType": "binary/octet-stream",
  "Metadata": {}
}
```



For versioning-enabled buckets, the `x-amz-expiration` response header applies only to current versions of objects.

Recommendations for implementing S3 REST API

You should follow these recommendations when implementing the S3 REST API for use with StorageGRID.

Recommendations for HEADs to non-existent objects

If your application routinely checks to see if an object exists at a path where you don't expect the object to actually exist, you should use the "Available" [consistency](#). For example, you should use the "Available" consistency if your application HEADs a location before PUT-ing to it.

Otherwise, if the HEAD operation does not find the object, you might receive a high number of 500 Internal Server errors if two or more Storage Nodes at the same site are unavailable or a remote site is unreachable.

You can set the "Available" consistency for each bucket using the [PUT Bucket consistency](#) request, or you can specify the consistency in the request header for an individual API operation.

Recommendations for object keys

Follow these recommendations for object key names, based on when the bucket was first created.

Buckets created in StorageGRID 11.4 or earlier

- Don't use random values as the first four characters of object keys. This is in contrast to the former AWS recommendation for key prefixes. Instead, use non-random, non-unique prefixes, such as `image`.
- If you do follow the former AWS recommendation to use random and unique characters in key prefixes, prefix the object keys with a directory name. That is, use this format:

```
mybucket/mydir/f8e3-image3132.jpg
```

Instead of this format:

```
mybucket/f8e3-image3132.jpg
```

Buckets created in StorageGRID 11.4 or later

Restricting object key names to meet performance best practices is not required. In most cases, you can use random values for the first four characters of object key names.



An exception to this is an S3 workload that continuously removes all objects after a short period of time. To minimize the performance impact for this use case, vary a leading portion of the key name every several thousand objects with something like the date. For example, suppose an S3 client typically writes 2,000 objects/second and the ILM or bucket lifecycle policy removes all objects after three days. To minimize the performance impact, you might name keys using a pattern like this: `/mybucket/mydir/yyyymmddhhmmss-random_UUID.jpg`

Recommendations for "range reads"

If the [global option to compress stored objects](#) is enabled, S3 client applications should avoid performing GetObject 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. GetObject 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.

Support for Amazon S3 REST API

S3 REST API implementation details

The StorageGRID system implements the Simple Storage Service API (API Version 2006-03-01) with support for most operations, and with some limitations. You need to understand the implementation details when you are integrating S3 REST API client applications.

The StorageGRID system supports both virtual hosted-style requests and path-style requests.

Date handling

The StorageGRID implementation of the S3 REST API only supports valid HTTP date formats.

The StorageGRID system only supports valid HTTP date formats for any headers that accept date values. The time portion of the date can be specified in Greenwich Mean Time (GMT) format, or in Universal Coordinated Time (UTC) format with no time zone offset (+0000 must be specified). If you include the `x-amz-date` header in your request, it overrides any value specified in the Date request header. When using AWS Signature Version 4, the `x-amz-date` header must be present in the signed request because the date header is not supported.

Common request headers

The StorageGRID system supports the common request headers defined by [Amazon Simple Storage Service API Reference: Common Request Headers](#), with one exception.

Request header	Implementation
Authorization	Full support for AWS Signature Version 2 Support for AWS Signature Version 4, with the following exceptions: • When you provide the actual payload checksum value in <code>x-amz-content-sha256</code>, the value is accepted without validation, as if the value <code>UNSIGNED-PAYLOAD</code> had been provided for the header. When you provide an <code>x-amz-content-sha256</code> header value that implies <code>aws-chunked</code> streaming (for example, <code>STREAMING-AWS4-HMAC-SHA256-PAYLOAD</code>), the chunk signatures aren't verified against the chunk data.
x-amz-security-token	Not implemented. Returns <code>XNotImplemented</code> .

Common response headers

The StorageGRID system supports all of the common response headers defined by the *Simple Storage Service API Reference*, with one exception.

Response header	Implementation
x-amz-id-2	Not used

Authenticate requests

The StorageGRID system supports both authenticated and anonymous access to objects using the S3 API.

The S3 API supports Signature version 2 and Signature version 4 for authenticating S3 API requests.

Authenticated requests must be signed using your access key ID and secret access key.

The StorageGRID system supports two authentication methods: the HTTP `Authorization` header and using query parameters.

Use the HTTP Authorization header

The HTTP `Authorization` header is used by all S3 API operations except Anonymous requests where permitted by the bucket policy. The `Authorization` header contains all of the required signing information to authenticate a request.

Use query parameters

You can use query parameters to add authentication information to a URL. This is known as presigning the URL, which can be used to grant temporary access to specific resources. Users with the presigned URL don't

need to know the secret access key to access the resource, which enables you to provide third-party restricted access to a resource.

Operations on the service

The StorageGRID system supports the following operations on the service.

Operation	Implementation
ListBuckets (previously named GET Service)	Implemented with all Amazon S3 REST API behavior. Subject to change without notice.
GET Storage Usage	The StorageGRID GET Storage Usage request tells you the total amount of storage in use by an account, and for each bucket associated with the account. This is an operation on the service with a path of / and a custom query parameter (?x-ntap-sg-usage) added.
OPTIONS /	Client applications can issue OPTIONS / requests to the S3 port on a Storage Node, without providing S3 authentication credentials, to determine whether the Storage Node is available. You can use this request for monitoring, or to allow external load balancers to identify when a Storage Node is down.

Operations on buckets

The StorageGRID system supports a maximum of 5,000 buckets for each S3 tenant account.

Each grid can have a maximum of 100,000 buckets.

To support 5,000 buckets, each Storage Node in the grid must have a minimum of 64 GB of RAM.

Bucket name restrictions follow the AWS US Standard region restrictions, but you should further restrict them to DNS naming conventions to support S3 virtual hosted-style requests.

See the following for more information:

- [Amazon Simple Storage Service User Guide: Bucket quotas, restrictions, and limitations](#)
- [Configure S3 endpoint domain names](#)

The ListObjects (GET Bucket) and ListObjectVersions (GET Bucket object versions) operations support StorageGRID [consistency values](#).

You can check whether updates to last access time are enabled or disabled for individual buckets. See [GET Bucket last access time](#).

The following table describes how StorageGRID implements S3 REST API bucket operations. To perform any of these operations, the necessary access credentials must be provided for the account.

Operation	Implementation
CreateBucket	Creates a new bucket. By creating the bucket, you become the bucket owner. • Bucket names must comply with the following rules: ◦ Must be unique across each StorageGRID system (not just unique within the tenant account). ◦ Must be DNS compliant. ◦ Must contain at least 3 and no more than 63 characters. ◦ Can be a series of one or more labels, with adjacent labels separated by a period. Each label must start and end with a lowercase letter or a number and can only use lowercase letters, numbers, and hyphens. ◦ Must not look like a text-formatted IP address. ◦ Should not use periods in virtual hosted style requests. Periods will cause problems with server wildcard certificate verification. • By default, buckets are created in the <code>us-east-1</code> region; however, you can use the <code>LocationConstraint</code> request element in the request body to specify a different region. When using the <code>LocationConstraint</code> element, you must specify the exact name of a region that has been defined using the Grid Manager or the Grid Management API. Contact your system administrator if you don't know the region name you should use. Note: An error will occur if your CreateBucket request uses a region that has not been defined in StorageGRID. • You can include the <code>x-amz-bucket-object-lock-enabled</code> request header to create a bucket with S3 Object Lock enabled. See Use S3 REST API to configure S3 Object Lock. You must enable S3 Object Lock when you create the bucket. You can't add or disable S3 Object Lock after a bucket is created. S3 Object Lock requires bucket versioning, which is enabled automatically when you create the bucket.
DeleteBucket	Deletes the bucket.
DeleteBucketCors	Deletes the CORS configuration for the bucket.
DeleteBucketEncryption	Deletes the default encryption from the bucket. Existing encrypted objects remain encrypted, but any new objects added to the bucket aren't encrypted.
DeleteBucketLifecycle	Deletes the lifecycle configuration from the bucket. See Create S3 lifecycle configuration .
DeleteBucketPolicy	Deletes the policy attached to the bucket.
DeleteBucketReplication	Deletes the replication configuration attached to the bucket.

Operation	Implementation
DeleteBucketTagging	Uses the <code>tagging</code> subresource to remove all tags from a bucket. Caution: If a non-default ILM policy tag is set for this bucket, there will be a <code>NTAP-SG-ILM-BUCKET-TAG</code> bucket tag with a value assigned to it. Do not issue a <code>DeleteBucketTagging</code> request if there is a <code>NTAP-SG-ILM-BUCKET-TAG</code> bucket tag. Instead, issue a <code>PutBucketTagging</code> request with only the <code>NTAP-SG-ILM-BUCKET-TAG</code> tag and its assigned value to remove all other tags from the bucket. Do not modify or remove the <code>NTAP-SG-ILM-BUCKET-TAG</code> bucket tag.
GetBucketAcl	Returns a positive response and the ID, DisplayName, and Permission of the bucket owner, indicating that the owner has full access to the bucket.
GetBucketCors	Returns the <code>cors</code> configuration for the bucket.
GetBucketEncryption	Returns the default encryption configuration for the bucket.
GetBucketLifecycleConfiguration (previously named GET Bucket lifecycle)	Returns the lifecycle configuration for the bucket. See Create S3 lifecycle configuration .
GetBucketLocation	Returns the region that was set using the <code>LocationConstraint</code> element in the <code>CreateBucket</code> request. If the bucket's region is <code>us-east-1</code> , an empty string is returned for the region.
GetBucketNotificationConfiguration (previously named GET Bucket notification)	Returns the notification configuration attached to the bucket.
GetBucketPolicy	Returns the policy attached to the bucket.
GetBucketReplication	Returns the replication configuration attached to the bucket.
GetBucketTagging	Uses the <code>tagging</code> subresource to return all tags for a bucket. Caution: If a non-default ILM policy tag is set for this bucket, there will be a <code>NTAP-SG-ILM-BUCKET-TAG</code> bucket tag with a value assigned to it. Do not modify or remove this tag.

Operation	Implementation
GetBucketVersioning	This implementation uses the <code>versioning</code> subresource to return the versioning state of a bucket. • <i>blank</i>: Versioning has never been enabled (bucket is "Unversioned") • Enabled: Versioning is enabled • Suspended: Versioning was previously enabled and is suspended
GetObjectLockConfiguration	Returns the bucket default retention mode and default retention period, if configured. See Use S3 REST API to configure S3 Object Lock.
HeadBucket	Determines if a bucket exists and you have permission to access it. This operation returns: • <code>x-ntap-sg-bucket-id</code>: The UUID of the bucket in UUID format. • <code>x-ntap-sg-trace-id</code>: The unique trace ID of the associated request.
ListObjects and ListObjectsV2 (previously named GET Bucket)	Returns some or all (up to 1,000) of the objects in a bucket. The Storage Class for objects can have either of two values, even if the object was ingested with the <code>REDUCED_REDUNDANCY</code> storage class option: • <code>STANDARD</code>, which indicates the object is stored in a storage pool consisting of Storage Nodes. • <code>GLACIER</code>, which indicates that the object has been moved to the external bucket specified by the Cloud Storage Pool. If the bucket contains large numbers of deleted keys that have the same prefix, the response might include some <code>CommonPrefixes</code> that don't contain keys.
ListObjectVersions (previously named GET Bucket Object versions)	With <code>READ</code> access on a bucket, using this operation with the <code>versions</code> subresource lists metadata of all of the versions of objects in the bucket.
PutBucketCors	Sets the CORS configuration for a bucket so that the bucket can service cross-origin requests. Cross-origin resource sharing (CORS) is a security mechanism that allows client web applications in one domain to access resources in a different domain. For example, suppose you use an S3 bucket named <code>images</code> to store graphics. By setting the CORS configuration for the <code>images</code> bucket, you can allow the images in that bucket to be displayed on the website <code>http://www.example.com</code>.

Operation	Implementation
PutBucketEncryption	Sets the default encryption state of an existing bucket. When bucket-level encryption is enabled, any new objects added to the bucket are encrypted. StorageGRID supports server-side encryption with StorageGRID-managed keys. When specifying the server-side encryption configuration rule, set the <code>SSEAlgorithm</code> parameter to <code>AES256</code>, and don't use the <code>KMSMasterKeyID</code> parameter. Bucket default encryption configuration is ignored if the object upload request already specifies encryption (that is, if the request includes the <code>x-amz-server-side-encryption-*</code> request header).
PutBucketLifecycleConfiguration <pre>(previously named PUT Bucket lifecycle)</pre>	Creates a new lifecycle configuration for the bucket or replaces an existing lifecycle configuration. StorageGRID supports up to 1,000 lifecycle rules in a lifecycle configuration. Each rule can include the following XML elements: • Expiration (Days, Date, ExpiredObjectDeleteMarker) • NoncurrentVersionExpiration (NewerNoncurrentVersions, NoncurrentDays) • Filter (Prefix, Tag) • Status • ID StorageGRID does not support these actions: • AbortIncompleteMultipartUpload • Transition See Create S3 lifecycle configuration. To understand how the Expiration action in a bucket lifecycle interacts with ILM placement instructions, see How ILM operates throughout an object's life. Note: Bucket lifecycle configuration can be used with buckets that have S3 Object Lock enabled, but bucket lifecycle configuration is not supported for legacy Compliant buckets.

Operation	Implementation
PutBucketNotificationConfiguration (previously named PUT Bucket notification)	Configures notifications for the bucket using the notification configuration XML included in the request body. You should be aware of the following implementation details: - StorageGRID supports Amazon Simple Notification Service (Amazon SNS) or Kafka topics as destinations. Simple Queue Service (SQS) or Amazon Lambda endpoints aren't supported. - The destination for notifications must be specified as the URN of an StorageGRID endpoint. Endpoints can be created using the Tenant Manager or the Tenant Management API. The endpoint must exist for notification configuration to succeed. If the endpoint does not exist, a 400 Bad Request error is returned with the code <code>InvalidArgument</code>. - You can't configure a notification for the following event types. These event types are not supported. <code>s3:ReducedRedundancyLostObject</code> <code>s3:ObjectRestore:Completed</code> - Event notifications sent from StorageGRID use the standard JSON format except that they don't include some keys and use specific values for others, as shown in the following list: eventSource <code>sgws:s3</code> awsRegion not included x-amz-id-2 not included arn <code>urn:sgws:s3:::bucket_name</code>
PutBucketPolicy	Sets the policy attached to the bucket. See Use bucket and group access policies .

Operation	Implementation
PutBucketReplication	Configures StorageGRID CloudMirror replication for the bucket using the replication configuration XML provided in the request body. For CloudMirror replication, you should be aware of the following implementation details: • StorageGRID only supports V1 of the replication configuration. This means that StorageGRID does not support the use of the <code>Filter</code> element for rules, and follows V1 conventions for deletion of object versions. For details, see Amazon Simple Storage Service User Guide: Replication configuration. • Bucket replication can be configured on versioned or unversioned buckets. • You can specify a different destination bucket in each rule of the replication configuration XML. A source bucket can replicate to more than one destination bucket. • Destination buckets must be specified as the URN of StorageGRID endpoints as specified in the Tenant Manager or the Tenant Management API. See Configure CloudMirror replication. The endpoint must exist for replication configuration to succeed. If the endpoint does not exist, the request fails as a 400 Bad Request. The error message states: <code>Unable to save the replication policy. The specified endpoint URN does not exist: URN.</code> • You don't need to specify a <code>Role</code> in the configuration XML. This value is not used by StorageGRID and will be ignored if submitted. • If you omit the storage class from the configuration XML, StorageGRID uses the <code>STANDARD</code> storage class by default. • If you delete an object from the source bucket or you delete the source bucket itself, the cross-region replication behavior is as follows: ◦ If you delete the object or bucket before it has been replicated, the object/bucket is not replicated and you aren't notified. ◦ If you delete the object or bucket after it has been replicated, StorageGRID follows standard Amazon S3 delete behavior for V1 of cross-region replication.

Operation	Implementation
PutBucketTagging	Uses the <code>tagging</code> subresource to add or update a set of tags for a bucket. When adding bucket tags, be aware of the following limitations: • Both StorageGRID and Amazon S3 support up to 50 tags for each bucket. • Tags associated with a bucket must have unique tag keys. A tag key can be up to 128 Unicode characters in length. • Tag values can be up to 256 Unicode characters in length. • Key and values are case sensitive. Caution: If a non-default ILM policy tag is set for this bucket, there will be a <code>NTAP-SG-ILM-BUCKET-TAG</code> bucket tag with a value assigned to it. Make sure that the <code>NTAP-SG-ILM-BUCKET-TAG</code> bucket tag is included with the assigned value in all PutBucketTagging requests. Do not modify or remove this tag. Note: This operation will overwrite any current tags the bucket already has. If any existing tags are omitted from the set, those tags will be removed for the bucket.
PutBucketVersioning	Uses the <code>versioning</code> subresource to set the versioning state of an existing bucket. You can set the versioning state with one of the following values: • Enabled: Enables versioning for the objects in the bucket. All objects added to the bucket receive a unique version ID. • Suspended: Disables versioning for the objects in the bucket. All objects added to the bucket receive the version ID <code>null</code>.
PutObjectLockConfiguration	Configures or removes the bucket default retention mode and default retention period. If the default retention period is modified, the retain-until-date of existing object versions remains the same and is not recalculated using the new default retention period. See Use S3 REST API to configure S3 Object Lock for detailed information.

Operations on objects

Operations on objects

This section describes how the StorageGRID system implements S3 REST API operations for objects.

The following conditions apply to all object operations:

- StorageGRID [consistency values](#) are supported by all operations on objects, with the exception of the following:
 - `GetObjectAcl`
 - `OPTIONS /`

- PutObjectLegalHold
- PutObjectRetention
- SelectObjectContent
- Conflicting client requests, such as two clients writing to the same key, are resolved on a "latest-wins" basis. The timing for the "latest-wins" evaluation is based on when the StorageGRID system completes a given request, and not on when S3 clients begin an operation.
- All objects in a StorageGRID bucket are owned by the bucket owner, including objects created by an anonymous user, or by another account.
- Data objects ingested to the StorageGRID system through Swift can't be accessed through S3.

The following table describes how StorageGRID implements S3 REST API object operations.

Operation	Implementation
DeleteObject	Multi-Factor Authentication (MFA) and the response header <code>x-amz-mfa</code> aren't supported. When processing a DeleteObject request, StorageGRID attempts to immediately remove all copies of the object from all stored locations. If successful, StorageGRID returns a response to the client immediately. If all copies can't be removed within 30 seconds (for example, because a location is temporarily unavailable), StorageGRID queues the copies for removal and then indicates success to the client. Versioning To remove a specific version, the requestor must be the bucket owner and use the <code>versionId</code> subresource. Using this subresource permanently deletes the version. If the <code>versionId</code> corresponds to a delete marker, the response header <code>x-amz-delete-marker</code> is returned set to <code>true</code>. • If an object is deleted without the <code>versionId</code> subresource on a bucket with versioning enabled, it results in the generation of a delete marker. The <code>versionId</code> for the delete marker is returned using the <code>x-amz-version-id</code> response header, and the <code>x-amz-delete-marker</code> response header is returned set to <code>true</code>. • If an object is deleted without the <code>versionId</code> subresource on a bucket with versioning suspended, it results in a permanent deletion of an already existing 'null' version or a 'null' delete marker, and the generation of a new 'null' delete marker. The <code>x-amz-delete-marker</code> response header is returned set to <code>true</code>. Note: In certain cases, multiple delete markers might exist for an object. See Use S3 REST API to configure S3 Object Lock to learn how to delete object versions in GOVERNANCE mode.

Operation	Implementation
DeleteObjects (previously named DELETE Multiple Objects)	Multi-Factor Authentication (MFA) and the response header <code>x-amz-mfa</code> aren't supported. Multiple objects can be deleted in the same request message. See Use S3 REST API to configure S3 Object Lock to learn how to delete object versions in GOVERNANCE mode.
DeleteObjectTagging	Uses the <code>tagging</code> subresource to remove all tags from an object. Versioning If the <code>versionId</code> query parameter is not specified in the request, the operation deletes all tags from the most recent version of the object in a versioned bucket. If the current version of the object is a delete marker, a "MethodNotAllowed" status is returned with the <code>x-amz-delete-marker</code> response header set to <code>true</code>.
GetObject	GetObject
GetObjectAcl	If the necessary access credentials are provided for the account, the operation returns a positive response and the ID, DisplayName, and Permission of the object owner, indicating that the owner has full access to the object.
GetObjectLegalHold	Use S3 REST API to configure S3 Object Lock
GetObjectRetention	Use S3 REST API to configure S3 Object Lock
GetObjectTagging	Uses the <code>tagging</code> subresource to return all tags for an object. Versioning If the <code>versionId</code> query parameter is not specified in the request, the operation returns all tags from the most recent version of the object in a versioned bucket. If the current version of the object is a delete marker, a "MethodNotAllowed" status is returned with the <code>x-amz-delete-marker</code> response header set to <code>true</code>.
HeadObject	HeadObject
RestoreObject	RestoreObject
PutObject	PutObject

Operation	Implementation
CopyObject (previously named PUT Object - Copy)	CopyObject
PutObjectLegalHold	Use S3 REST API to configure S3 Object Lock
PutObjectRetention	Use S3 REST API to configure S3 Object Lock
PutObjectTagging	Uses the <code>tagging</code> subresource to add a set of tags to an existing object. Object tag limits You can add tags to new objects when you upload them, or you can add them to existing objects. Both StorageGRID and Amazon S3 support up to 10 tags for each object. Tags associated with an object must have unique tag keys. A tag key can be up to 128 Unicode characters in length and tag values can be up to 256 Unicode characters in length. Key and values are case sensitive. Tag updates and ingest behavior When you use PutObjectTagging to update an object's tags, StorageGRID does not re-ingest the object. This means that the option for Ingest Behavior specified in the matching ILM rule is not used. Any changes to object placement that are triggered by the update are made when ILM is re-evaluated by normal background ILM processes. This means that if the ILM rule uses the Strict option for ingest behavior, no action is taken if the required object placements can't be made (for example, because a newly required location is unavailable). The updated object retains its current placement until the required placement is possible. Resolving conflicts Conflicting client requests, such as two clients writing to the same key, are resolved on a "latest-wins" basis. The timing for the "latest-wins" evaluation is based on when the StorageGRID system completes a given request, and not on when S3 clients begin an operation. Versioning If the <code>versionId</code> query parameter is not specified in the request, the operation add tags to the most recent version of the object in a versioned bucket. If the current version of the object is a delete marker, a "MethodNotAllowed" status is returned with the <code>x-amz-delete-marker</code> response header set to <code>true</code>.
SelectObjectContent	SelectObjectContent

Use S3 Select

StorageGRID supports the following Amazon S3 Select clauses, data types, and operators for the [SelectObjectContent](#) command.



Any items not listed aren't supported.

For syntax, see [SelectObjectContent](#). For more information about S3 Select, see the [AWS documentation for S3 Select](#).

Only tenant accounts that have S3 Select enabled can issue SelectObjectContent queries. See the [considerations and requirements for using S3 Select](#).

Clauses

- SELECT list
- FROM clause
- WHERE clause
- LIMIT clause

Data types

- bool
- integer
- string
- float
- decimal, numeric
- timestamp

Operators

Logical operators

- AND
- NOT
- OR

Comparison operators

- <
- >
- <=
- >=
- =
- =
- <>

- !=
- BETWEEN
- IN

Pattern matching operators

- LIKE
- _
- %

Unitary operators

- IS NULL
- IS NOT NULL

Math operators

- +
- -
- *
- /
- %

StorageGRID follows the Amazon S3 Select operator precedence.

Aggregate functions

- AVG()
- COUNT(*)
- MAX()
- MIN()
- SUM()

Conditional functions

- CASE
- COALESCE
- NULLIF

Conversion functions

- CAST (for supported datatype)

Date functions

- DATE_ADD
- DATE_DIFF

- EXTRACT
- TO_STRING
- TO_TIMESTAMP
- UTCNOW

String functions

- CHAR_LENGTH, CHARACTER_LENGTH
- LOWER
- SUBSTRING
- TRIM
- UPPER

Use server-side encryption

Server-side encryption allows you to protect your object data at rest. StorageGRID encrypts the data as it writes the object and decrypts the data when you access the object.

If you want to use server-side encryption, you can choose either of two mutually exclusive options, based on how the encryption keys are managed:

- **SSE (server-side encryption with StorageGRID-managed keys)**: When you issue an S3 request to store an object, StorageGRID encrypts the object with a unique key. When you issue an S3 request to retrieve the object, StorageGRID uses the stored key to decrypt the object.
- **SSE-C (server-side encryption with customer-provided keys)**: When you issue an S3 request to store an object, you provide your own encryption key. When you retrieve an object, you provide the same encryption key as part of your request. If the two encryption keys match, the object is decrypted and your object data is returned.

While StorageGRID manages all object encryption and decryption operations, you must manage the encryption keys you provide.



The encryption keys you provide are never stored. If you lose an encryption key, you lose the corresponding object.



If an object is encrypted with SSE or SSE-C, any bucket-level or grid-level encryption settings are ignored.

Use SSE

To encrypt an object with a unique key managed by StorageGRID, you use the following request header:

```
x-amz-server-side-encryption
```

The SSE request header is supported by the following object operations:

- [PutObject](#)

- [CopyObject](#)
- [CreateMultipartUpload](#)

Use SSE-C

To encrypt an object with a unique key that you manage, you use three request headers:

Request header	Description
x-amz-server-side-encryption-customer-algorithm	Specify the encryption algorithm. The header value must be AES256.
x-amz-server-side-encryption-customer-key	Specify the encryption key that will be used to encrypt or decrypt the object. The value for the key must be 256-bit, base64-encoded.
x-amz-server-side-encryption-customer-key-MD5	Specify the MD5 digest of the encryption key according to RFC 1321, which is used to ensure the encryption key was transmitted without error. The value for the MD5 digest must be base64-encoded 128-bit.

The SSE-C request headers are supported by the following object operations:

- [GetObject](#)
- [HeadObject](#)
- [PutObject](#)
- [CopyObject](#)
- [CreateMultipartUpload](#)
- [UploadPart](#)
- [UploadPartCopy](#)

Considerations for using server-side encryption with customer-provided keys (SSE-C)

Before using SSE-C, be aware of the following considerations:

- You must use https.



StorageGRID rejects any requests made over http when using SSE-C. For security considerations, you should consider any key you send accidentally using http to be compromised. Discard the key, and rotate as appropriate.

- The ETag in the response is not the MD5 of the object data.
- You must manage the mapping of encryption keys to objects. StorageGRID does not store encryption keys. You are responsible for tracking the encryption key you provide for each object.
- If your bucket is versioning-enabled, each object version should have its own encryption key. You are responsible for tracking the encryption key used for each object version.
- Because you manage encryption keys on the client side, you must also manage any additional safeguards, such as key rotation, on the client side.



The encryption keys you provide are never stored. If you lose an encryption key, you lose the corresponding object.

- If cross-grid replication or CloudMirror replication is configured for the bucket, you can't ingest SSE-C objects. The ingest operation will fail.

Related information

[Amazon S3 User Guide: Using server-side encryption with customer-provided keys \(SSE-C\)](#)

CopyObject

You can use the S3 CopyObject request to create a copy of an object that is already stored in S3. A CopyObject operation is the same as performing GetObject followed by PutObject.

Resolve conflicts

Conflicting client requests, such as two clients writing to the same key, are resolved on a "latest-wins" basis. The timing for the "latest-wins" evaluation is based on when the StorageGRID system completes a given request, and not on when S3 clients begin an operation.

Object size

The maximum *recommended* size for a single PutObject operation is 5 GiB (5,368,709,120 bytes). If you have objects that are larger than 5 GiB, use [multipart upload](#) instead.

The maximum *supported* size for a single PutObject operation is 5 TiB (5,497,558,138,880 bytes).



If you upgraded from StorageGRID 11.6 or earlier, the S3 PUT Object size too large alert will be triggered if you attempt to upload an object that exceeds 5 GiB. If you have a new installation of StorageGRID 11.7 or 11.8, the alert won't be triggered in this case. However, to align with the AWS S3 standard, future releases of StorageGRID won't support uploads of objects larger than 5 GiB.

UTF-8 characters in user metadata

If a request includes (unescaped) UTF-8 values in the key name or value of user-defined metadata, StorageGRID behavior is undefined.

StorageGRID does not parse or interpret escaped UTF-8 characters included in the key name or value of user-defined metadata. Escaped UTF-8 characters are treated as ASCII characters:

- Requests succeed if user-defined metadata includes escaped UTF-8 characters.
- StorageGRID does not return the `x-amz-missing-meta` header if the interpreted value of the key name or value includes unprintable characters.

Supported request headers

The following request headers are supported:

- `Content-Type`

- `x-amz-copy-source`
- `x-amz-copy-source-if-match`
- `x-amz-copy-source-if-none-match`
- `x-amz-copy-source-if-unmodified-since`
- `x-amz-copy-source-if-modified-since`
- `x-amz-meta-`, followed by a name-value pair containing user-defined metadata
- `x-amz-metadata-directive`: The default value is `COPY`, which enables you to copy the object and associated metadata.

You can specify `REPLACE` to overwrite the existing metadata when copying the object, or to update the object metadata.

- `x-amz-storage-class`
- `x-amz-tagging-directive`: The default value is `COPY`, which enables you to copy the object and all tags.

You can specify `REPLACE` to overwrite the existing tags when copying the object, or to update the tags.

- S3 Object Lock request headers:
 - `x-amz-object-lock-mode`
 - `x-amz-object-lock-retain-until-date`
 - `x-amz-object-lock-legal-hold`

If a request is made without these headers, the bucket default retention settings are used to calculate the object version mode and retain-until-date. See [Use S3 REST API to configure S3 Object Lock](#).

- SSE request headers:
 - `x-amz-copy-source-server-side-encryption-customer-algorithm`
 - `x-amz-copy-source-server-side-encryption-customer-key`
 - `x-amz-copy-source-server-side-encryption-customer-key-MD5`
 - `x-amz-server-side-encryption`
 - `x-amz-server-side-encryption-customer-key-MD5`
 - `x-amz-server-side-encryption-customer-key`
 - `x-amz-server-side-encryption-customer-algorithm`

See [Request headers for server-side encryption](#)

Unsupported request headers

The following request headers aren't supported:

- `Cache-Control`
- `Content-Disposition`

- Content-Encoding
- Content-Language
- Expires
- x-amz-checksum-algorithm

When you copy an object, if the source object has a checksum, StorageGRID doesn't copy that checksum value to the new object. This behavior applies whether or not you try to use `x-amz-checksum-algorithm` in the object request.

- x-amz-website-redirect-location

Storage class options

The `x-amz-storage-class` request header is supported, and affects how many object copies StorageGRID creates if the matching ILM rule uses the Dual commit or Balanced [ingest option](#).

- STANDARD

(Default) Specifies a dual-commit ingest operation when the ILM rule uses the Dual commit option, or when the Balanced option falls back to creating interim copies.

- REDUCED_REDUNDANCY

Specifies a single-commit ingest operation when the ILM rule uses the Dual commit option, or when the Balanced option falls back to creating interim copies.



If you are ingesting an object into a bucket with S3 Object Lock enabled, the `REDUCED_REDUNDANCY` option is ignored. If you are ingesting an object into a legacy Compliant bucket, the `REDUCED_REDUNDANCY` option returns an error. StorageGRID will always perform a dual-commit ingest to ensure that compliance requirements are satisfied.

Using x-amz-copy-source in CopyObject

If the source bucket and key, specified in the `x-amz-copy-source` header, are different from the destination bucket and key, a copy of the source object data is written to the destination.

If the source and destination match, and the `x-amz-metadata-directive` header is specified as `REPLACE`, the object's metadata is updated with the metadata values supplied in the request. In this case, StorageGRID does not re-ingest the object. This has two important consequences:

- You can't use CopyObject to encrypt an existing object in place, or to change the encryption of an existing object in place. If you supply the `x-amz-server-side-encryption` header or the `x-amz-server-side-encryption-customer-algorithm` header, StorageGRID rejects the request and returns `XNotImplemented`.
- The option for Ingest Behavior specified in the matching ILM rule is not used. Any changes to object placement that are triggered by the update are made when ILM is re-evaluated by normal background ILM processes.

This means that if the ILM rule uses the Strict option for ingest behavior, no action is taken if the required object placements can't be made (for example, because a newly required location is unavailable). The

updated object retains its current placement until the required placement is possible.

Request headers for server-side encryption

If you [use server-side encryption](#), the request headers you provide depend on whether the source object is encrypted and on whether you plan to encrypt the target object.

- If the source object is encrypted using a customer-provided key (SSE-C), you must include the following three headers in the CopyObject request, so the object can be decrypted and then copied:
 - `x-amz-copy-source-server-side-encryption-customer-algorithm`: Specify AES256.
 - `x-amz-copy-source-server-side-encryption-customer-key`: Specify the encryption key you provided when you created the source object.
 - `x-amz-copy-source-server-side-encryption-customer-key-MD5`: Specify the MD5 digest you provided when you created the source object.
- If you want to encrypt the target object (the copy) with a unique key that you provide and manage, include the following three headers:
 - `x-amz-server-side-encryption-customer-algorithm`: Specify AES256.
 - `x-amz-server-side-encryption-customer-key`: Specify a new encryption key for the target object.
 - `x-amz-server-side-encryption-customer-key-MD5`: Specify the MD5 digest of the new encryption key.



The encryption keys you provide are never stored. If you lose an encryption key, you lose the corresponding object. Before using customer-provided keys to secure object data, review the considerations for [using server-side encryption](#).

- If you want to encrypt the target object (the copy) with a unique key managed by StorageGRID (SSE), include this header in the CopyObject request:

- `x-amz-server-side-encryption`



The `server-side-encryption` value of the object can't be updated. Instead, make a copy with a new `server-side-encryption` value using `x-amz-metadata-directive: REPLACE`.

Versioning

If the source bucket is versioned, you can use the `x-amz-copy-source` header to copy the latest version of an object. To copy a specific version of an object, you must explicitly specify the version to copy using the `versionId` subresource. If the destination bucket is versioned, the generated version is returned in the `x-amz-version-id` response header. If versioning is suspended for the target bucket, then `x-amz-version-id` returns a "null" value.

GetObject

You can use the S3 GetObject request to retrieve an object from an S3 bucket.

GetObject and multipart objects

You can use the `partNumber` request parameter to retrieve a specific part of a multipart or segmented object. The `x-amz-mp-parts-count` response element indicates how many parts the object has.

You can set `partNumber` to 1 for both segmented/multipart objects and non-segmented/non-multipart objects; however, the `x-amz-mp-parts-count` response element is only returned for segmented or multipart objects.

UTF-8 characters in user metadata

StorageGRID does not parse or interpret escaped UTF-8 characters in user-defined metadata. GET requests for an object with escaped UTF-8 characters in user-defined metadata don't return the `x-amz-missing-meta` header if the key name or value includes unprintable characters.

Supported request header

The following request header is supported:

- `x-amz-checksum-mode`: Specify `ENABLED`

The `Range` header isn't supported with `x-amz-checksum-mode` for `GetObject`. When you include `Range` in the request with `x-amz-checksum-mode` enabled, StorageGRID doesn't return a checksum value in the response.

Unsupported request header

The following request header is not supported and returns `XNotImplemented`:

- `x-amz-website-redirect-location`

Versioning

If a `versionId` subresource is not specified, the operation fetches the most recent version of the object in a versioned bucket. If the current version of the object is a delete marker, a "Not Found" status is returned with the `x-amz-delete-marker` response header set to `true`.

Request headers for server-side encryption with customer-provided encryption keys (SSE-C)

Use all three of the headers if the object is encrypted with a unique key that you provided.

- `x-amz-server-side-encryption-customer-algorithm`: Specify `AES256`.
- `x-amz-server-side-encryption-customer-key`: Specify your encryption key for the object.
- `x-amz-server-side-encryption-customer-key-MD5`: Specify the MD5 digest of the object's encryption key.



The encryption keys you provide are never stored. If you lose an encryption key, you lose the corresponding object. Before using customer-provided keys to secure object data, review the considerations in [Use server-side encryption](#).

Behavior of GetObject for Cloud Storage Pool objects

If an object has been stored in a [Cloud Storage Pool](#), the behavior of a GetObject request depends on the state of the object. See [HeadObject](#) for more details.



If an object is stored in a Cloud Storage Pool and one or more copies of the object also exist on the grid, GetObject requests will attempt to retrieve data from the grid, before retrieving it from the Cloud Storage Pool.

State of object	Behavior of GetObject
Object ingested into StorageGRID but not yet evaluated by ILM, or object stored in a traditional storage pool or using erasure coding	200 OK A copy of the object is retrieved.
Object in Cloud Storage Pool but not yet transitioned to a non-retrievable state	200 OK A copy of the object is retrieved.
Object transitioned to a non-retrievable state	403 Forbidden, InvalidObjectState Use a RestoreObject request to restore the object to a retrievable state.
Object in process of being restored from a non-retrievable state	403 Forbidden, InvalidObjectState Wait for the RestoreObject request to complete.
Object fully restored to the Cloud Storage Pool	200 OK A copy of the object is retrieved.

Multipart or segmented objects in a Cloud Storage Pool

If you uploaded a multipart object or if StorageGRID split a large object into segments, StorageGRID determines whether the object is available in the Cloud Storage Pool by sampling a subset of the object's parts or segments. In some cases, a GetObject request might incorrectly return 200 OK when some parts of the object have already been transitioned to a non-retrievable state or when some parts of the object have not yet been restored.

In these cases:

- The GetObject request might return some data but stop midway through the transfer.
- A subsequent GetObject request might return 403 Forbidden.

GetObject and cross-grid replication

If you are using [grid federation](#) and [cross-grid replication](#) is enabled for a bucket, the S3 client can verify an object's replication status by issuing a GetObject request. The response includes the StorageGRID-specific `x-ntap-sg-cgr-replication-status` response header, which will have one of the following values:

Grid	Replication status
Source	• COMPLETED: The replication was successful. • PENDING: The object hasn't been replicated yet. • FAILURE: The replication failed with a permanent failure. A user must resolve the error.
Destination	REPLICA: The object was replicated from the source grid.



StorageGRID does not support the `x-amz-replication-status` header.

HeadObject

You can use the S3 HeadObject request to retrieve metadata from an object without returning the object itself. If the object is stored in a Cloud Storage Pool, you can use HeadObject to determine the object's transition state.

HeadObject and multipart objects

You can use the `partNumber` request parameter to retrieve metadata for a specific part of a multipart or segmented object. The `x-amz-mp-parts-count` response element indicates how many parts the object has.

You can set `partNumber` to 1 for both segmented/multipart objects and non-segmented/non-multipart objects; however, the `x-amz-mp-parts-count` response element is only returned for segmented or multipart objects.

UTF-8 characters in user metadata

StorageGRID does not parse or interpret escaped UTF-8 characters in user-defined metadata. HEAD requests for an object with escaped UTF-8 characters in user-defined metadata don't return the `x-amz-missing-meta` header if the key name or value includes unprintable characters.

Supported request header

The following request header is supported:

- `x-amz-checksum-mode`

The `partNumber` parameter and Range header aren't supported with `x-amz-checksum-mode` for HeadObject. When you include them in the request with `x-amz-checksum-mode` enabled, StorageGRID doesn't return a checksum value in the response.

Unsupported request header

The following request header isn't supported and returns `XNotImplemented`:

- `x-amz-website-redirect-location`

Versioning

If a `versionId` subresource is not specified, the operation fetches the most recent version of the object in a versioned bucket. If the current version of the object is a delete marker, a "Not Found" status is returned with the `x-amz-delete-marker` response header set to `true`.

Request headers for server-side encryption with customer-provided encryption keys (SSE-C)

Use all three of these headers if the object is encrypted with a unique key that you provided.

- `x-amz-server-side-encryption-customer-algorithm`: Specify AES256.
- `x-amz-server-side-encryption-customer-key`: Specify your encryption key for the object.
- `x-amz-server-side-encryption-customer-key-MD5`: Specify the MD5 digest of the object's encryption key.



The encryption keys you provide are never stored. If you lose an encryption key, you lose the corresponding object. Before using customer-provided keys to secure object data, review the considerations in [Use server-side encryption](#).

HeadObject responses for Cloud Storage Pool objects

If the object is stored in a [Cloud Storage Pool](#), the following response headers are returned:

- `x-amz-storage-class`: GLACIER
- `x-amz-restore`

The response headers provide information about the state of an object as it is moved to a Cloud Storage Pool, optionally transitioned to a non-retrievable state, and restored.

State of object	Response to HeadObject
Object ingested into StorageGRID but not yet evaluated by ILM, or object stored in a traditional storage pool or using erasure coding	200 OK (No special response header is returned.)
Object in Cloud Storage Pool but not yet transitioned to a non-retrievable state	200 OK <code>x-amz-storage-class: GLACIER</code> <code>x-amz-restore: ongoing-request="false", expiry-date="Sat, 23 July 20 2030 00:00:00 GMT"</code> Until the object is transitioned to a non-retrievable state, the value for <code>expiry-date</code> is set to some distant time in the future. The exact time of transition is not controlled by the StorageGRID system.

State of object	Response to HeadObject
Object has transitioned to non-retrievable state, but at least one copy also exists on the grid	200 OK x-amz-storage-class: GLACIER x-amz-restore: ongoing-request="false", expiry-date="Sat, 23 July 20 2030 00:00:00 GMT" The value for expiry-date is set to some distant time in the future. Note: If the copy on the grid is not available (for example, a Storage Node is down), you must issue a RestoreObject request to restore the copy from the Cloud Storage Pool before you can successfully retrieve the object.
Object transitioned to a non-retrievable state, and no copy exists on the grid	200 OK x-amz-storage-class: GLACIER
Object in process of being restored from a non-retrievable state	200 OK x-amz-storage-class: GLACIER x-amz-restore: ongoing-request="true"
Object fully restored to the Cloud Storage Pool	200 OK x-amz-storage-class: GLACIER x-amz-restore: ongoing-request="false", expiry-date="Sat, 23 July 20 2018 00:00:00 GMT" The expiry-date indicates when the object in the Cloud Storage Pool will be returned to a non-retrievable state.

Multipart or segmented objects in Cloud Storage Pool

If you uploaded a multipart object or if StorageGRID split a large object into segments, StorageGRID determines whether the object is available in the Cloud Storage Pool by sampling a subset of the object's parts or segments. In some cases, a HeadObject request might incorrectly return `x-amz-restore: ongoing-request="false"` when some parts of the object have already been transitioned to a non-retrievable state or when some parts of the object have not yet been restored.

HeadObject and cross-grid replication

If you are using [grid federation](#) and [cross-grid replication](#) is enabled for a bucket, the S3 client can verify an object's replication status by issuing a HeadObject request. The response includes the StorageGRID-specific `x-ntap-sg-cgr-replication-status` response header, which will have one of the following values:

Grid	Replication status
Source	• COMPLETED: The replication was successful. • PENDING: The object hasn't been replicated yet. • FAILURE: The replication failed with a permanent failure. A user must resolve the error.
Destination	REPLICA: The object was replicated from the source grid.



StorageGRID does not support the `x-amz-replication-status` header.

PutObject

You can use the S3 PutObject request to add an object to a bucket.

Resolve conflicts

Conflicting client requests, such as two clients writing to the same key, are resolved on a "latest-wins" basis. The timing for the "latest-wins" evaluation is based on when the StorageGRID system completes a given request, and not on when S3 clients begin an operation.

Object size

The maximum *recommended* size for a single PutObject operation is 5 GiB (5,368,709,120 bytes). If you have objects that are larger than 5 GiB, use [multipart upload](#) instead.

The maximum *supported* size for a single PutObject operation is 5 TiB (5,497,558,138,880 bytes).



If you upgraded from StorageGRID 11.6 or earlier, the S3 PUT Object size too large alert will be triggered if you attempt to upload an object that exceeds 5 GiB. If you have a new installation of StorageGRID 11.7 or 11.8, the alert won't be triggered in this case. However, to align with the AWS S3 standard, future releases of StorageGRID won't support uploads of objects larger than 5 GiB.

User metadata size

Amazon S3 limits the size of user-defined metadata within each PUT request header to 2 KB. StorageGRID limits user metadata to 24 KiB. The size of user-defined metadata is measured by taking the sum of the number of bytes in the UTF-8 encoding of each key and value.

UTF-8 characters in user metadata

If a request includes (unescaped) UTF-8 values in the key name or value of user-defined metadata, StorageGRID behavior is undefined.

StorageGRID does not parse or interpret escaped UTF-8 characters included in the key name or value of user-defined metadata. Escaped UTF-8 characters are treated as ASCII characters:

- PutObject, CopyObject, GetObject, and HeadObject requests succeed if user-defined metadata includes escaped UTF-8 characters.
- StorageGRID does not return the `x-amz-missing-meta` header if the interpreted value of the key name or value includes unprintable characters.

Object tag limits

You can add tags to new objects when you upload them, or you can add them to existing objects. Both StorageGRID and Amazon S3 support up to 10 tags for each object. Tags associated with an object must have unique tag keys. A tag key can be up to 128 Unicode characters in length and tag values can be up to 256 Unicode characters in length. Key and values are case sensitive.

Object ownership

In StorageGRID, all objects are owned by the bucket owner account, including objects created by a non-owner account or an anonymous user.

Supported request headers

The following request headers are supported:

- Cache-Control
- Content-Disposition
- Content-Encoding

When you specify `aws-chunked` for `Content-Encoding` StorageGRID does not verify the following items:

- StorageGRID does not verify the `chunk-signature` against the chunk data.
- StorageGRID does not verify the value that you provide for `x-amz-decoded-content-length` against the object.
- Content-Language
- Content-Length
- Content-MD5
- Content-Type
- Expires
- Transfer-Encoding

Chunked transfer encoding is supported if `aws-chunked` payload signing is also used.

- `x-amz-checksum-sha256`
- `x-amz-meta-`, followed by a name-value pair containing user-defined metadata.

When specifying the name-value pair for user-defined metadata, use this general format:

```
x-amz-meta-name: value
```

If you want to use the **User defined creation time** option as the Reference time for an ILM rule, you must use `creation-time` as the name of the metadata that records when the object was created. For example:

```
x-amz-meta-creation-time: 1443399726
```

The value for `creation-time` is evaluated as seconds since January 1, 1970.



An ILM rule can't use both a **User defined creation time** for the Reference time and the Balanced or Strict ingest option. An error is returned when the ILM rule is created.

- `x-amz-tagging`
- S3 Object Lock request headers
 - `x-amz-object-lock-mode`
 - `x-amz-object-lock-retain-until-date`
 - `x-amz-object-lock-legal-hold`

If a request is made without these headers, the bucket default retention settings are used to calculate the object version mode and retain-until-date. See [Use S3 REST API to configure S3 Object Lock](#).

- SSE request headers:
 - `x-amz-server-side-encryption`
 - `x-amz-server-side-encryption-customer-key-MD5`
 - `x-amz-server-side-encryption-customer-key`
 - `x-amz-server-side-encryption-customer-algorithm`

See [Request headers for server-side encryption](#)

Unsupported request headers

The following request headers aren't supported:

- `x-amz-acl`
- `x-amz-sdk-checksum-algorithm`
- `x-amz-trailer`
- `x-amz-website-redirect-location`

The `x-amz-website-redirect-location` header returns `XNotImplemented`.

Storage class options

The `x-amz-storage-class` request header is supported. The value submitted for `x-amz-storage-class` affects how StorageGRID protects object data during ingest and not how many persistent copies of the object are stored in the StorageGRID system (which is determined by ILM).

If the ILM rule matching an ingested object uses the Strict ingest option, the `x-amz-storage-class` header has no effect.

The following values can be used for `x-amz-storage-class`:

- **STANDARD (Default)**
 - **Dual commit:** If the ILM rule specifies the Dual commit option for Ingest Behavior, as soon as an object is ingested a second copy of that object is created and distributed to a different Storage Node (dual commit). When the ILM is evaluated, StorageGRID determines if these initial interim copies satisfy the placement instructions in the rule. If they don't, new object copies might need to be made in different locations and the initial interim copies might need to be deleted.
 - **Balanced:** If the ILM rule specifies the Balanced option and StorageGRID can't immediately make all copies specified in the rule, StorageGRID makes two interim copies on different Storage Nodes.

If StorageGRID can immediately create all object copies specified in the ILM rule (synchronous placement), the `x-amz-storage-class` header has no effect.

- **REDUCED_REDUNDANCY**
 - **Dual commit:** If the ILM rule specifies the Dual commit option for Ingest Behavior, StorageGRID creates a single interim copy as the object is ingested (single commit).
 - **Balanced:** If the ILM rule specifies the Balanced option, StorageGRID makes a single interim copy only if the system can't immediately make all copies specified in the rule. If StorageGRID can perform synchronous placement, this header has no effect.

The `REDUCED_REDUNDANCY` option is best used when the ILM rule that matches the object creates a single replicated copy. In this case using `REDUCED_REDUNDANCY` eliminates the unnecessary creation and deletion of an extra object copy for every ingest operation.

Using the `REDUCED_REDUNDANCY` option is not recommended in other circumstances.

`REDUCED_REDUNDANCY` increases the risk of object data loss during ingest. For example, you might lose data if the single copy is initially stored on a Storage Node that fails before ILM evaluation can occur.



Having only one replicated copy for any time period puts data at risk of permanent loss. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.

Specifying `REDUCED_REDUNDANCY` only affects how many copies are created when an object is first ingested. It does not affect how many copies of the object are made when the object is evaluated by the active ILM policies, and does not result in data being stored at lower levels of redundancy in the StorageGRID system.



If you are ingesting an object into a bucket with S3 Object Lock enabled, the `REDUCED_REDUNDANCY` option is ignored. If you are ingesting an object into a legacy Compliant bucket, the `REDUCED_REDUNDANCY` option returns an error. StorageGRID will always perform a dual-commit ingest to ensure that compliance requirements are satisfied.

Request headers for server-side encryption

You can use the following request headers to encrypt an object with server-side encryption. The SSE and SSE-C options are mutually exclusive.

- **SSE:** Use the following header if you want to encrypt the object with a unique key managed by StorageGRID.

- `x-amz-server-side-encryption`

When the `x-amz-server-side-encryption` header isn't included in the `PutObject` request, the grid-wide [stored object encryption setting](#) is omitted from the `PutObject` response.

- **SSE-C:** Use all three of these headers if you want to encrypt the object with a unique key that you provide and manage.
 - `x-amz-server-side-encryption-customer-algorithm`: Specify AES256.
 - `x-amz-server-side-encryption-customer-key`: Specify your encryption key for the new object.
 - `x-amz-server-side-encryption-customer-key-MD5`: Specify the MD5 digest of the new object's encryption key.



The encryption keys you provide are never stored. If you lose an encryption key, you lose the corresponding object. Before using customer-provided keys to secure object data, review the considerations for [using server-side encryption](#).



If an object is encrypted with SSE or SSE-C, any bucket-level or grid-level encryption settings are ignored.

Versioning

If versioning is enabled for a bucket, a unique `versionId` is automatically generated for the version of the object being stored. This `versionId` is also returned in the response using the `x-amz-version-id` response header.

If versioning is suspended, the object version is stored with a null `versionId` and if a null version already exists it will be overwritten.

Signature calculations for the Authorization header

When using the `Authorization` header to authenticate requests, StorageGRID differs from AWS in the following ways:

- StorageGRID doesn't require `host` headers to be included within `CanonicalHeaders`.
- StorageGRID doesn't require `Content-Type` to be included within `CanonicalHeaders`.
- StorageGRID doesn't require `x-amz-*` headers to be included within `CanonicalHeaders`.



As a general best practice, always include these headers within `CanonicalHeaders` to ensure they are verified; however, if you exclude these headers, StorageGRID does not return an error.

For details, refer to [Signature Calculations for the Authorization Header: Transferring Payload in a Single](#)

Chunk (AWS Signature Version 4).

Related information

- [Manage objects with ILM](#)
- [Amazon Simple Storage Service API Reference: PutObject](#)

RestoreObject

You can use the S3 RestoreObject request to restore an object that is stored in a Cloud Storage Pool.

Supported request type

StorageGRID only supports RestoreObject requests to restore an object. It does not support the `SELECT` type of restoration. Select requests return `XNotImplemented`.

Versioning

Optionally, specify `versionId` to restore a specific version of an object in a versioned bucket. If you don't specify `versionId`, the most recent version of the object is restored

Behavior of RestoreObject on Cloud Storage Pool objects

If an object has been stored in a [Cloud Storage Pool](#), a RestoreObject request has the following behavior, based on the state of the object. See [HeadObject](#) for more details.



If an object is stored in a Cloud Storage Pool and one or more copies of the object also exist on the grid, there is no need to restore the object by issuing a RestoreObject request. Instead, the local copy can be retrieved directly, using a GetObject request.

State of object	Behavior of RestoreObject
Object ingested into StorageGRID but not yet evaluated by ILM, or object is not in a Cloud Storage Pool	403 Forbidden, InvalidObjectState
Object in Cloud Storage Pool but not yet transitioned to a non-retrievable state	200 OK No changes are made. Note: Before an object has been transitioned to a non-retrievable state, you can't change its <code>expiry-date</code> .

State of object	Behavior of RestoreObject
Object transitioned to a non-retrievable state	202 Accepted Restores a retrievable copy of the object to the Cloud Storage Pool for the number of days specified in the request body. At the end of this period, the object is returned to a non-retrievable state. Optionally, use the <code>Tier</code> request element to determine how long the restore job will take to finish (<code>Expedited</code>, <code>Standard</code>, or <code>Bulk</code>). If you don't specify <code>Tier</code>, the <code>Standard</code> tier is used. Important: If an object has been transitioned to S3 Glacier Deep Archive or the Cloud Storage Pool uses Azure Blob storage, you can't restore it using the <code>Expedited</code> tier. The following error is returned 403 Forbidden, InvalidTier: Retrieval option is not supported by this storage class.
Object in process of being restored from a non-retrievable state	409 Conflict, RestoreAlreadyInProgress
Object fully restored to the Cloud Storage Pool	200 OK Note: If an object has been restored to a retrievable state, you can change its <code>expiry-date</code> by reissuing the <code>RestoreObject</code> request with a new value for <code>Days</code>. The restoration date is updated relative to the time of the request.

SelectObjectContent

You can use the S3 `SelectObjectContent` request to filter the contents of an S3 object based on a simple SQL statement.

For more information see [Amazon Simple Storage Service API Reference: SelectObjectContent](#).

Before you begin

- The tenant account has the S3 Select permission.
- You have `s3:GetObject` permission for the object you want to query.
- The object you want to query must be in one of the following formats:
 - **CSV**. Can be used as is or compressed into GZIP or BZIP2 archives.
 - **Parquet**. Additional requirements for Parquet objects:
 - S3 Select supports only columnar compression using GZIP or Snappy. S3 Select doesn't support whole-object compression for Parquet objects.
 - S3 Select doesn't support Parquet output. You must specify the output format as CSV or JSON.
 - The maximum uncompressed row group size is 512 MB.
 - You must use the data types specified in the object's schema.
 - You can't use INTERVAL, JSON, LIST, TIME, or UUID logical types.
- Your SQL expression has a maximum length of 256 KB.

- Any record in the input or results has a maximum length of 1 MiB.

CSV request syntax example

```
POST /{Key+}?select&select-type=2 HTTP/1.1
Host: Bucket.s3.abc-company.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<SelectObjectContentRequest xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Expression>string</Expression>
  <ExpressionType>string</ExpressionType>
  <RequestProgress>
    <Enabled>boolean</Enabled>
  </RequestProgress>
  <InputSerialization>
    <CompressionType>GZIP</CompressionType>
    <CSV>
      <AllowQuotedRecordDelimiter>boolean</AllowQuotedRecordDelimiter>
      <Comments>#</Comments>
      <FieldDelimiter>\t</FieldDelimiter>
      <FileHeaderInfo>USE</FileHeaderInfo>
      <QuoteCharacter>'</QuoteCharacter>
      <QuoteEscapeCharacter>\\</QuoteEscapeCharacter>
      <RecordDelimiter>\n</RecordDelimiter>
    </CSV>
  </InputSerialization>
  <OutputSerialization>
    <CSV>
      <FieldDelimiter>string</FieldDelimiter>
      <QuoteCharacter>string</QuoteCharacter>
      <QuoteEscapeCharacter>string</QuoteEscapeCharacter>
      <QuoteFields>string</QuoteFields>
      <RecordDelimiter>string</RecordDelimiter>
    </CSV>
  </OutputSerialization>
  <ScanRange>
    <End>long</End>
    <Start>long</Start>
  </ScanRange>
</SelectObjectContentRequest>
```

Parquet request syntax example

```

POST /{Key+}?select&select-type=2 HTTP/1.1
Host: Bucket.s3.abc-company.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<SelectObjectContentRequest xmlns=http://s3.amazonaws.com/doc/2006-03-01/>
  <Expression>string</Expression>
  <ExpressionType>string</ExpressionType>
  <RequestProgress>
    <Enabled>boolean</Enabled>
  </RequestProgress>
  <InputSerialization>
    <CompressionType>GZIP</CompressionType>
    <PARQUET>
    </PARQUET>
  </InputSerialization>
  <OutputSerialization>
    <CSV>
      <FieldDelimiter>string</FieldDelimiter>
      <QuoteCharacter>string</QuoteCharacter>
      <QuoteEscapeCharacter>string</QuoteEscapeCharacter>
      <QuoteFields>string</QuoteFields>
      <RecordDelimiter>string</RecordDelimiter>
    </CSV>
  </OutputSerialization>
  <ScanRange>
    <End>long</End>
    <Start>long</Start>
  </ScanRange>
</SelectObjectContentRequest>

```

SQL query example

This query gets the state name, 2010 populations, estimated 2015 populations, and the percentage of change from US census data. Records in the file that aren't states are ignored.

```

SELECT STNAME, CENSUS2010POP, POPESTIMATE2015, CAST((POPESTIMATE2015 -
CENSUS2010POP) AS DECIMAL) / CENSUS2010POP * 100.0 FROM S3Object WHERE
NAME = STNAME

```

The first few lines of the file to be queried, SUB-EST2020_ALL.csv, look like this:


```
SUMLEV, STATE, COUNTY, PLACE, COUSUB, CONCIT, PRIMGEO_FLAG, FUNCSTAT, NAME, STNAME,
CENSUS2010POP,
ESTIMATESBASE2010, POPESTIMATE2010, POPESTIMATE2011, POPESTIMATE2012, POPESTIM
ATE2013, POPESTIMATE2014,
POPESTIMATE2015, POPESTIMATE2016, POPESTIMATE2017, POPESTIMATE2018, POPESTIMAT
E2019, POPESTIMATE042020,
POPESTIMATE2020
040, 01, 000, 00000, 00000, 00000, 0, A, Alabama, Alabama, 4779736, 4780118, 4785514, 4
799642, 4816632, 4831586,
4843737, 4854803, 4866824, 4877989, 4891628, 4907965, 4920706, 4921532
162, 01, 000, 00124, 00000, 00000, 0, A, Abbeville
city, Alabama, 2688, 2705, 2699, 2694, 2645, 2629, 2610, 2602,
2587, 2578, 2565, 2555, 2555, 2553
162, 01, 000, 00460, 00000, 00000, 0, A, Adamsville
city, Alabama, 4522, 4487, 4481, 4474, 4453, 4430, 4399, 4371,
4335, 4304, 4285, 4254, 4224, 4211
162, 01, 000, 00484, 00000, 00000, 0, A, Addison
town, Alabama, 758, 754, 751, 750, 745, 744, 742, 734, 734, 728,
725, 723, 719, 717
```

AWS-CLI usage example (CSV)

```
aws s3api select-object-content --endpoint-url https://10.224.7.44:10443
--no-verify-ssl --bucket 619c0755-9e38-42e0-a614-05064f74126d --key SUB-
EST2020_ALL.csv --expression-type SQL --input-serialization '{"CSV":
{"FileHeaderInfo": "USE", "Comments": "#", "QuoteEscapeCharacter": "\"",
"RecordDelimiter": "\n", "FieldDelimiter": ",", "QuoteCharacter": "\"",
"AllowQuotedRecordDelimiter": false}, "CompressionType": "NONE"}' --output
-serialization '{"CSV": {"QuoteFields": "ASNEEDED",
"QuoteEscapeCharacter": "#", "RecordDelimiter": "\n", "FieldDelimiter":
",", "QuoteCharacter": "\""}}' --expression "SELECT STNAME, CENSUS2010POP,
POPESTIMATE2015, CAST((POPESTIMATE2015 - CENSUS2010POP) AS DECIMAL) /
CENSUS2010POP * 100.0 FROM S3Object WHERE NAME = STNAME" changes.csv
```

The first few lines of the output file, `changes.csv`, look like this:

```
Alabama, 4779736, 4854803, 1.5705260708959658022953568983726297854
Alaska, 710231, 738430, 3.9703983633493891424057806544631253775
Arizona, 6392017, 6832810, 6.8959922978928247531256565807005832431
Arkansas, 2915918, 2979732, 2.1884703204959810255295244928012378949
California, 37253956, 38904296, 4.4299724839960620557988526104449148971
Colorado, 5029196, 5454328, 8.4532796097030221132761578590295546246
```

AWS-CLI usage example (Parquet)

```
aws s3api select-object-content --endpoint-url https://10.224.7.44:10443
--bucket 619c0755-9e38-42e0-a614-05064f74126d --key SUB-
EST2020_ALL.parquet --expression "SELECT STNAME, CENSUS2010POP,
POPESTIMATE2015, CAST((POPESTIMATE2015 - CENSUS2010POP) AS DECIMAL) /
CENSUS2010POP * 100.0 FROM S3Object WHERE NAME = STNAME" --expression-type
'SQL' --input-serialization '{"Parquet":{}}' --output-serialization
'{"CSV":{}}' changes.csv
```

The first few lines of the output file, changes.csv, look like this:

```
Alabama,4779736,4854803,1.5705260708959658022953568983726297854
Alaska,710231,738430,3.9703983633493891424057806544631253775
Arizona,6392017,6832810,6.8959922978928247531256565807005832431
Arkansas,2915918,2979732,2.1884703204959810255295244928012378949
California,37253956,38904296,4.4299724839960620557988526104449148971
Colorado,5029196,5454328,8.4532796097030221132761578590295546246
```

Operations for multipart uploads

Operations for multipart uploads

This section describes how StorageGRID supports operations for multipart uploads.

The following conditions and notes apply to all multipart upload operations:

- You should not exceed 1,000 concurrent multipart uploads to a single bucket because the results of ListMultipartUploads queries for that bucket might return incomplete results.
- StorageGRID enforces AWS size limits for multipart parts. S3 clients must follow these guidelines:
 - Each part in a multipart upload must be between 5 MiB (5,242,880 bytes) and 5 GiB (5,368,709,120 bytes).
 - The last part can be smaller than 5 MiB (5,242,880 bytes).
 - In general, part sizes should be as large as possible. For example, use part sizes of 5 GiB for a 100 GiB object. Because each part is considered a unique object, using large part sizes reduces StorageGRID metadata overhead.
 - For objects smaller than 5 GiB, consider using non-multipart upload instead.
- ILM is evaluated for each part of a multipart object as it is ingested and for the object as a whole when the multipart upload completes, if the ILM rule uses the Balanced or Strict [ingest option](#). You should be aware of how this affects object and part placement:
 - If ILM changes while an S3 multipart upload is in progress, some parts of the object might not meet current ILM requirements when the multipart upload completes. Any part that is not placed correctly is queued for ILM re-evaluation and moved to the correct location later.
 - When evaluating ILM for a part, StorageGRID filters on the size of the part, not the size of the object. This means that parts of an object can be stored in locations that don't meet ILM requirements for the

object as a whole. For example, if a rule specifies that all objects 10 GB or larger are stored at DC1 while all smaller objects are stored at DC2, each 1 GB part of a 10-part multipart upload is stored at DC2 at ingest. However, when ILM is evaluated for the object as a whole, all parts of the object are moved to DC1.

- All of the multipart upload operations support StorageGRID [consistency values](#).
- When an object is ingested using multipart upload, the [object segmentation threshold \(1 GiB\)](#) is not applied.
- As required, you can use [server-side encryption](#) with multipart uploads. To use SSE (server-side encryption with StorageGRID-managed keys), you include the `x-amz-server-side-encryption` request header in the `CreateMultipartUpload` request only. To use SSE-C (server-side encryption with customer-provided keys), you specify the same three encryption key request headers in the `CreateMultipartUpload` request and in each subsequent `UploadPart` request.

Operation	Implementation
<code>AbortMultipartUpload</code>	Implemented with all Amazon S3 REST API behavior. Subject to change without notice.
<code>CompleteMultipartUpload</code>	See CompleteMultipartUpload
<code>CreateMultipartUpload</code> (previously named <code>Initiate Multipart Upload</code>)	See CreateMultipartUpload
<code>ListMultipartUploads</code>	See ListMultipartUploads
<code>ListParts</code>	Implemented with all Amazon S3 REST API behavior. Subject to change without notice.
<code>UploadPart</code>	See UploadPart
<code>UploadPartCopy</code>	See UploadPartCopy

CompleteMultipartUpload

The `CompleteMultipartUpload` operation completes a multipart upload of an object by assembling the previously uploaded parts.



StorageGRID supports non-consecutive values in ascending order for the `partNumber` request parameter with `CompleteMultipartUpload`. The parameter can start with any value.

Resolve conflicts

Conflicting client requests, such as two clients writing to the same key, are resolved on a "latest-wins" basis. The timing for the "latest-wins" evaluation is based on when the StorageGRID system completes a given request, and not on when S3 clients begin an operation.

Supported request headers

The following request headers are supported:

- `x-amz-checksum-sha256`
- `x-amz-storage-class`

The `x-amz-storage-class` header affects how many object copies StorageGRID creates if the matching ILM rule specifies the [Dual commit or Balanced ingest option](#).

- `STANDARD`

(Default) Specifies a dual-commit ingest operation when the ILM rule uses the Dual commit option, or when the Balanced option falls back to creating interim copies.

- `REDUCED_REDUNDANCY`

Specifies a single-commit ingest operation when the ILM rule uses the Dual commit option, or when the Balanced option falls back to creating interim copies.



If you are ingesting an object into a bucket with S3 Object Lock enabled, the `REDUCED_REDUNDANCY` option is ignored. If you are ingesting an object into a legacy Compliant bucket, the `REDUCED_REDUNDANCY` option returns an error. StorageGRID will always perform a dual-commit ingest to ensure that compliance requirements are satisfied.



If a multipart upload is not completed within 15 days, the operation is marked as inactive and all associated data is deleted from the system.



The `ETag` value returned is not an MD5 sum of the data, but follows the Amazon S3 API implementation of the `ETag` value for multipart objects.

Unsupported request headers

The following request headers aren't supported:

- `x-amz-sdk-checksum-algorithm`
- `x-amz-trailer`

Versioning

This operation completes a multipart upload. If versioning is enabled for a bucket, the object version is created after completion of the multipart upload.

If versioning is enabled for a bucket, a unique `versionId` is automatically generated for the version of the object being stored. This `versionId` is also returned in the response using the `x-amz-version-id` response header.

If versioning is suspended, the object version is stored with a null `versionId` and if a null version already exists it will be overwritten.



When versioning is enabled for a bucket, completing a multipart upload always creates a new version, even if there are concurrent multipart uploads completed on the same object key. When versioning is not enabled for a bucket, it is possible to initiate a multipart upload and then have another multipart upload initiate and complete first on the same object key. On non-versioned buckets, the multipart upload that completes last takes precedence.

Failed replication, notification, or metadata notification

If the bucket where the multipart upload occurs is configured for a platform service, multipart upload succeeds even if the associated replication or notification action fails.

A tenant can trigger the failed replication or notification by updating the object's metadata or tags. A tenant can resubmit the existing values to avoid making unwanted changes.

Refer to [Troubleshoot platform services](#).

CreateMultipartUpload

The CreateMultipartUpload (previously named Initiate Multipart Upload) operation initiates a multipart upload for an object, and returns an upload ID.

The `x-amz-storage-class` request header is supported. The value submitted for `x-amz-storage-class` affects how StorageGRID protects object data during ingest and not how many persistent copies of the object are stored in the StorageGRID system (which is determined by ILM).

If the ILM rule matching an ingested object uses the Strict [ingest option](#), the `x-amz-storage-class` header has no effect.

The following values can be used for `x-amz-storage-class`:

- STANDARD (Default)
 - **Dual commit:** If the ILM rule specifies the Dual commit ingest option, as soon as an object is ingested a second copy of that object is created and distributed to a different Storage Node (dual commit). When the ILM is evaluated, StorageGRID determines if these initial interim copies satisfy the placement instructions in the rule. If they don't, new object copies might need to be made in different locations and the initial interim copies might need to be deleted.
 - **Balanced:** If the ILM rule specifies the Balanced option and StorageGRID can't immediately make all copies specified in the rule, StorageGRID makes two interim copies on different Storage Nodes.

If StorageGRID can immediately create all object copies specified in the ILM rule (synchronous placement), the `x-amz-storage-class` header has no effect.

- REDUCED_REDUNDANCY
 - **Dual commit:** If the ILM rule specifies the Dual commit option, StorageGRID creates a single interim copy as the object is ingested (single commit).
 - **Balanced:** If the ILM rule specifies the Balanced option, StorageGRID makes a single interim copy only if the system can't immediately make all copies specified in the rule. If StorageGRID can perform synchronous placement, this header has no effect.
The REDUCED_REDUNDANCY option is best used when the ILM rule that matches the object creates a single replicated copy. In this case using REDUCED_REDUNDANCY eliminates the unnecessary creation and deletion of an extra object copy for every ingest operation.

Using the `REDUCED_REDUNDANCY` option is not recommended in other circumstances.

`REDUCED_REDUNDANCY` increases the risk of object data loss during ingest. For example, you might lose data if the single copy is initially stored on a Storage Node that fails before ILM evaluation can occur.



Having only one replicated copy for any time period puts data at risk of permanent loss. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.

Specifying `REDUCED_REDUNDANCY` only affects how many copies are created when an object is first ingested. It does not affect how many copies of the object are made when the object is evaluated by the active ILM policies, and does not result in data being stored at lower levels of redundancy in the StorageGRID system.



If you are ingesting an object into a bucket with S3 Object Lock enabled, the `REDUCED_REDUNDANCY` option is ignored. If you are ingesting an object into a legacy Compliant bucket, the `REDUCED_REDUNDANCY` option returns an error. StorageGRID will always perform a dual-commit ingest to ensure that compliance requirements are satisfied.

Supported request headers

The following request headers are supported:

- `Content-Type`
- `x-amz-checksum-algorithm`

Currently, only the SHA256 value for `x-amz-checksum-algorithm` is supported.

- `x-amz-meta-`, followed by a name-value pair containing user-defined metadata

When specifying the name-value pair for user-defined metadata, use this general format:

```
x-amz-meta-_name_: `value`
```

If you want to use the **User defined creation time** option as the Reference time for an ILM rule, you must use `creation-time` as the name of the metadata that records when the object was created. For example:

```
x-amz-meta-creation-time: 1443399726
```

The value for `creation-time` is evaluated as seconds since January 1, 1970.



Adding `creation-time` as user-defined metadata is not allowed if you are adding an object to a bucket that has legacy Compliance enabled. An error will be returned.

- S3 Object Lock request headers:
 - `x-amz-object-lock-mode`

- `x-amz-object-lock-retain-until-date`
- `x-amz-object-lock-legal-hold`

If a request is made without these headers, the bucket default retention settings are used to calculate the object version retain-until-date.

[Use S3 REST API to configure S3 Object Lock](#)

- SSE request headers:

- `x-amz-server-side-encryption`
- `x-amz-server-side-encryption-customer-key-MD5`
- `x-amz-server-side-encryption-customer-key`
- `x-amz-server-side-encryption-customer-algorithm`

[Request headers for server-side encryption](#)



For information about how StorageGRID handles UTF-8 characters, see [PutObject](#).

Request headers for server-side encryption

You can use the following request headers to encrypt a multipart object with server-side encryption. The SSE and SSE-C options are mutually exclusive.

- **SSE:** Use the following header in the CreateMultipartUpload request if you want to encrypt the object with a unique key managed by StorageGRID. Don't specify this header in any of the UploadPart requests.

- `x-amz-server-side-encryption`

- **SSE-C:** Use all three of these headers in the CreateMultipartUpload request (and in each subsequent UploadPart request) if you want to encrypt the object with a unique key that you provide and manage.

- `x-amz-server-side-encryption-customer-algorithm`: Specify AES256.
- `x-amz-server-side-encryption-customer-key`: Specify your encryption key for the new object.
- `x-amz-server-side-encryption-customer-key-MD5`: Specify the MD5 digest of the new object's encryption key.



The encryption keys you provide are never stored. If you lose an encryption key, you lose the corresponding object. Before using customer-provided keys to secure object data, review the considerations for [using server-side encryption](#).

Unsupported request headers

The following request header isn't supported:

- `x-amz-website-redirect-location`

The `x-amz-website-redirect-location` header returns `XNotImplemented`.

Versioning

Multipart upload consists of separate operations for initiating the upload, listing uploads, uploading parts, assembling the uploaded parts, and completing the upload. Objects are created (and versioned if applicable) when the CompleteMultipartUpload operation is performed.

ListMultipartUploads

The ListMultipartUploads operation lists in-progress multipart uploads for a bucket.

The following request parameters are supported:

- encoding-type
- key-marker
- max-uploads
- prefix
- upload-id-marker
- Host
- Date
- Authorization

Versioning

Multipart upload consists of separate operations for initiating the upload, listing uploads, uploading parts, assembling the uploaded parts, and completing the upload. Objects are created (and versioned if applicable) when the CompleteMultipartUpload operation is performed.

UploadPart

The UploadPart operation uploads a part in a multipart upload for an object.

Supported request headers

The following request headers are supported:

- x-amz-checksum-sha256
- Content-Length
- Content-MD5

Request headers for server-side encryption

If you specified SSE-C encryption for the CreateMultipartUpload request, you must also include the following request headers in each UploadPart request:

- x-amz-server-side-encryption-customer-algorithm: Specify AES256.
- x-amz-server-side-encryption-customer-key: Specify the same encryption key that you provided in the CreateMultipartUpload request.

- `x-amz-server-side-encryption-customer-key-MD5`: Specify the same MD5 digest that you provided in the `CreateMultipartUpload` request.



The encryption keys you provide are never stored. If you lose an encryption key, you lose the corresponding object. Before using customer-provided keys to secure object data, review the considerations in [Use server-side encryption](#).

If you specified a SHA-256 checksum during the `CreateMultipartUpload` request, you must also include the following request header in each `UploadPart` request:

- `x-amz-checksum-sha256`: Specify the SHA-256 checksum for this part.

Unsupported request headers

The following request headers aren't supported:

- `x-amz-sdk-checksum-algorithm`
- `x-amz-trailer`

Versioning

Multipart upload consists of separate operations for initiating the upload, listing uploads, uploading parts, assembling the uploaded parts, and completing the upload. Objects are created (and versioned if applicable) when the `CompleteMultipartUpload` operation is performed.

UploadPartCopy

The `UploadPartCopy` operation uploads a part of an object by copying data from an existing object as the data source.

The `UploadPartCopy` operation is implemented with all Amazon S3 REST API behavior. Subject to change without notice.

This request reads and writes the object data specified in `x-amz-copy-source-range` within the StorageGRID system.

The following request headers are supported:

- `x-amz-copy-source-if-match`
- `x-amz-copy-source-if-none-match`
- `x-amz-copy-source-if-unmodified-since`
- `x-amz-copy-source-if-modified-since`

Request headers for server-side encryption

If you specified SSE-C encryption for the `CreateMultipartUpload` request, you must also include the following request headers in each `UploadPartCopy` request:

- `x-amz-server-side-encryption-customer-algorithm`: Specify AES256.
- `x-amz-server-side-encryption-customer-key`: Specify the same encryption key that you

provided in the CreateMultipartUpload request.

- `x-amz-server-side-encryption-customer-key-MD5`: Specify the same MD5 digest that you provided in the CreateMultipartUpload request.

If the source object is encrypted using a customer-provided key (SSE-C), you must include the following three headers in the UploadPartCopy request, so the object can be decrypted and then copied:

- `x-amz-copy-source-server-side-encryption-customer-algorithm`: Specify AES256.
- `x-amz-copy-source-server-side-encryption-customer-key`: Specify the encryption key you provided when you created the source object.
- `x-amz-copy-source-server-side-encryption-customer-key-MD5`: Specify the MD5 digest you provided when you created the source object.



The encryption keys you provide are never stored. If you lose an encryption key, you lose the corresponding object. Before using customer-provided keys to secure object data, review the considerations in [Use server-side encryption](#).

Versioning

Multipart upload consists of separate operations for initiating the upload, listing uploads, uploading parts, assembling the uploaded parts, and completing the upload. Objects are created (and versioned if applicable) when the CompleteMultipartUpload operation is performed.

Error responses

The StorageGRID system supports all standard S3 REST API error responses that apply. In addition, the StorageGRID implementation adds several custom responses.

Supported S3 API error codes

Name	HTTP status
AccessDenied	403 Forbidden
BadDigest	400 Bad Request
BucketAlreadyExists	409 Conflict
BucketNotEmpty	409 Conflict
IncompleteBody	400 Bad Request
InternalError	500 Internal Server Error
InvalidAccessKeyId	403 Forbidden
InvalidArgument	400 Bad Request

Name	HTTP status
InvalidBucketName	400 Bad Request
InvalidBucketState	409 Conflict
InvalidDigest	400 Bad Request
InvalidEncryptionAlgorithmError	400 Bad Request
InvalidPart	400 Bad Request
InvalidPartOrder	400 Bad Request
InvalidRange	416 Requested Range Not Satisfiable
InvalidRequest	400 Bad Request
InvalidStorageClass	400 Bad Request
InvalidTag	400 Bad Request
InvalidURI	400 Bad Request
KeyTooLong	400 Bad Request
MalformedXML	400 Bad Request
MetadataTooLarge	400 Bad Request
MethodNotAllowed	405 Method Not Allowed
MissingContentLength	411 Length Required
MissingRequestBodyError	400 Bad Request
MissingSecurityHeader	400 Bad Request
NoSuchBucket	404 Not Found
NoSuchKey	404 Not Found
NoSuchUpload	404 Not Found
NotImplemented	501 Not Implemented

Name	HTTP status
NoSuchBucketPolicy	404 Not Found
ObjectLockConfigurationNotFound	404 Not Found
PreconditionFailed	412 Precondition Failed
RequestTimeTooSkewed	403 Forbidden
ServiceUnavailable	503 Service Unavailable
SignatureDoesNotMatch	403 Forbidden
TooManyBuckets	400 Bad Request
UserKeyMustBeSpecified	400 Bad Request

StorageGRID custom error codes

Name	Description	HTTP status
XBucketLifecycleNotAllowed	Bucket lifecycle configuration is not allowed in a legacy Compliant bucket	400 Bad Request
XBucketPolicyParseException	Failed to parse received bucket policy JSON.	400 Bad Request
XComplianceConflict	Operation denied because of legacy Compliance settings.	403 Forbidden
XComplianceReducedRedundancyForbidden	Reduced redundancy is not allowed in legacy Compliant bucket	400 Bad Request
XMaxBucketPolicyLengthExceeded	Your policy exceeds the maximum allowed bucket policy length.	400 Bad Request
XMissingInternalRequestHeader	Missing a header of an internal request.	400 Bad Request
XNoSuchBucketCompliance	The specified bucket does not have legacy Compliance enabled.	404 Not Found
XNotAcceptable	The request contains one or more accept headers that could not be satisfied.	406 Not Acceptable
XNotImplemented	The request you provided implies functionality that is not implemented.	501 Not Implemented

StorageGRID custom operations

StorageGRID custom operations

The StorageGRID system supports custom operations that are added on to the S3 REST API.

The following table lists the custom operations supported by StorageGRID.

Operation	Description
GET Bucket consistency	Returns the consistency being applied to a particular bucket.
PUT Bucket consistency	Sets the consistency applied to a particular bucket.
GET Bucket last access time	Returns whether last access time updates are enabled or disabled for a particular bucket.
PUT Bucket last access time	Allows you to enable or disable last access time updates for a particular bucket.
DELETE Bucket metadata notification configuration	Deletes the metadata notification configuration XML associated with a particular bucket.
GET Bucket metadata notification configuration	Returns the metadata notification configuration XML associated with a particular bucket.
PUT Bucket metadata notification configuration	Configures the metadata notification service for a bucket.
GET Storage Usage	Tells you the total amount of storage in use by an account and for each bucket associated with the account.
Deprecated: CreateBucket with compliance settings	Deprecated and not supported: You can no longer create new buckets with Compliance enabled.
Deprecated: GET Bucket compliance	Deprecated but supported: Returns the compliance settings currently in effect for an existing legacy Compliant bucket.
Deprecated: PUT Bucket compliance	Deprecated but supported: Allows you to modify the compliance settings for an existing legacy Compliant bucket.

GET Bucket consistency

The GET Bucket consistency request allows you to determine the consistency being applied to a particular bucket.

The default consistency is set to guarantee read-after-write for newly created objects.

You must have the `s3:GetBucketConsistency` permission, or be account root, to complete this operation.

Request example

```
GET /bucket?x-ntap-sg-consistency HTTP/1.1
Date: date
Authorization: authorization string
Host: host
```

Response

In the response XML, `<Consistency>` will return one of the following values:

Consistency	Description
all	All nodes receive the data immediately, or the request will fail.
strong-global	Guarantees read-after-write consistency for all client requests across all sites.
strong-site	Guarantees read-after-write consistency for all client requests within a site.
read-after-new-write	(Default) Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.
available	Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that don't exist). Not supported for S3 FabricPool buckets.

Response example

```
HTTP/1.1 200 OK
Date: Fri, 18 Sep 2020 01:02:18 GMT
Connection: CLOSE
Server: StorageGRID/11.5.0
x-amz-request-id: 12345
Content-Length: 127
Content-Type: application/xml

<?xml version="1.0" encoding="UTF-8"?>
<Consistency xmlns="http://s3.storagegrid.com/doc/2015-02-01/">read-after-
new-write</Consistency>
```

Related information

[Consistency values](#)

PUT Bucket consistency

The PUT Bucket consistency request allows you to specify the consistency to apply to operations performed on a bucket.

The default consistency is set to guarantee read-after-write for newly created objects.

Before you begin

You must have the `s3:PutBucketConsistency` permission, or be account root, to complete this operation.

Request

The `x-ntap-sg-consistency` parameter must contain one of the following values:

Consistency	Description
all	All nodes receive the data immediately, or the request will fail.
strong-global	Guarantees read-after-write consistency for all client requests across all sites.
strong-site	Guarantees read-after-write consistency for all client requests within a site.
read-after-new-write	(Default) Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.
available	Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that don't exist). Not supported for S3 FabricPool buckets.

Note: In general, you should use the "Read-after-new-write" consistency. If requests aren't working correctly, change the application client behavior if possible. Or, configure the client to specify the consistency for each API request. Set the consistency at the bucket level only as a last resort.

Request example

```
PUT /bucket?x-ntap-sg-consistency=strong-global HTTP/1.1
Date: date
Authorization: authorization string
Host: host
```

Related information

[Consistency values](#)

GET Bucket last access time

The GET Bucket last access time request allows you to determine if last access time updates are enabled or disabled for individual buckets.

You must have the `s3:GetBucketLastAccessTime` permission, or be account root, to complete this operation.

Request example

```
GET /bucket?x-ntap-sg-lastaccesstime HTTP/1.1
Date: date
Authorization: authorization string
Host: host
```

Response example

This example shows that last access time updates are enabled for the bucket.

```
HTTP/1.1 200 OK
Date: Sat, 29 Nov 2015 01:02:18 GMT
Connection: CLOSE
Server: StorageGRID/10.3.0
x-amz-request-id: 12345
Content-Length: 127
Content-Type: application/xml

<?xml version="1.0" encoding="UTF-8"?>
<LastAccessTime xmlns="http://s3.storagegrid.com/doc/2015-02-01/">enabled
</LastAccessTime>
```

PUT Bucket last access time

The PUT Bucket last access time request allows you to enable or disable last access time updates for individual buckets. Disabling last access time updates improves performance, and is the default setting for all buckets created with version 10.3.0, or later.

You must have the `s3:PutBucketLastAccessTime` permission for a bucket, or be account root, to complete this operation.



Starting with StorageGRID version 10.3, updates to last access time are disabled by default for all new buckets. If you have buckets that were created using an earlier version of StorageGRID and you want to match the new default behavior, you must explicitly disable last access time updates for each of those earlier buckets. You can enable or disable updates to last access time using the PUT Bucket last access time request or from the details page for a bucket in the Tenant Manager. See [Enable or disable last access time updates](#).

If last access time updates are disabled for a bucket, the following behavior is applied to operations on the

bucket:

- GetObject, GetObjectAcl, GetObjectTagging, and HeadObject requests don't update last access time. The object is not added to queues for information lifecycle management (ILM) evaluation.
- CopyObject and PutObjectTagging requests that update only the metadata also update last access time. The object is added to queues for ILM evaluation.
- If updates to last access time are disabled for the source bucket, CopyObject requests don't update last access time for the source bucket. The object that was copied is not added to queues for ILM evaluation for the source bucket. However, for the destination, CopyObject requests always update last access time. The copy of the object is added to queues for ILM evaluation.
- CompleteMultipartUpload requests update last access time. The completed object is added to queues for ILM evaluation.

Request examples

This example enables last access time for a bucket.

```
PUT /bucket?x-ntap-sg-lastaccesstime=enabled HTTP/1.1
Date: date
Authorization: authorization string
Host: host
```

This example disables last access time for a bucket.

```
PUT /bucket?x-ntap-sg-lastaccesstime=disabled HTTP/1.1
Date: date
Authorization: authorization string
Host: host
```

DELETE Bucket metadata notification configuration

The DELETE Bucket metadata notification configuration request allows you to disable the search integration service for individual buckets by deleting the configuration XML.

You must have the `s3:DeleteBucketMetadataNotification` permission for a bucket, or be account root, to complete this operation.

Request example

This example shows disabling the search integration service for a bucket.

```
DELETE /test1?x-ntap-sg-metadata-notification HTTP/1.1
Date: date
Authorization: authorization string
Host: host
```

GET Bucket metadata notification configuration

The GET Bucket metadata notification configuration request allows you to retrieve the configuration XML used to configure search integration for individual buckets.

You must have the `s3:GetBucketMetadataNotification` permission, or be account root, to complete this operation.

Request example

This request retrieves the metadata notification configuration for the bucket named `bucket`.

```
GET /bucket?x-ntap-sg-metadata-notification HTTP/1.1
Date: date
Authorization: authorization string
Host: host
```

Response

The response body includes the metadata notification configuration for the bucket. The metadata notification configuration lets you determine how the bucket is configured for search integration. That is, it allows you to determine which objects are indexed, and which endpoints their object metadata is being sent to.

```
<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>rule-status</Status>
    <Prefix>key-prefix</Prefix>
    <Destination>
      <Urn>arn:aws:es:_region:account-
ID_:domain/_mydomain/myindex/mytype_</Urn>
    </Destination>
  </Rule>
  <Rule>
    <ID>Rule-2</ID>
    ...
  </Rule>
  ...
</MetadataNotificationConfiguration>
```

Each metadata notification configuration includes one or more rules. Each rule specifies the objects that it applies to and the destination where StorageGRID should send object metadata. Destinations must be specified using the URN of a StorageGRID endpoint.

Name	Description	Required
MetadataNotificationConfiguration	Container tag for rules used to specify the objects and destination for metadata notifications. Contains one or more Rule elements.	Yes
Rule	Container tag for a rule that identifies the objects whose metadata should be added to a specified index. Rules with overlapping prefixes are rejected. Included in the MetadataNotificationConfiguration element.	Yes
ID	Unique identifier for the rule. Included in the Rule element.	No
Status	Status can be 'Enabled' or 'Disabled'. No action is taken for rules that are disabled. Included in the Rule element.	Yes
Prefix	Objects that match the prefix are affected by the rule, and their metadata is sent to the specified destination. To match all objects, specify an empty prefix. Included in the Rule element.	Yes
Destination	Container tag for the destination of a rule. Included in the Rule element.	Yes

Name	Description	Required
Urn	URN of the destination where object metadata is sent. Must be the URN of a StorageGRID endpoint with the following properties: • <code>es</code> must be the third element. • The URN must end with the index and type where the metadata is stored, in the form <code>domain-name/myindex/mytype</code>. Endpoints are configured using the Tenant Manager or Tenant Management API. They take the following form: • <code>arn:aws:es:_region:account-ID_:domain/mydomain/myindex/mytype</code> • <code>urn:mysite:es:::mydomain/myindex/mytype</code> The endpoint must be configured before the configuration XML is submitted, or configuration will fail with a 404 error. Urn is included in the Destination element.	Yes

Response example

The XML included between the

`<MetadataNotificationConfiguration></MetadataNotificationConfiguration>` tags shows how integration with a search integration endpoint is configured for the bucket. In this example, object metadata is being sent to an Elasticsearch index named `current` and type named `2017` that is hosted in an AWS domain named `records`.

```
HTTP/1.1 200 OK
Date: Thu, 20 Jul 2017 18:24:05 GMT
Connection: KEEP-ALIVE
Server: StorageGRID/11.0.0
x-amz-request-id: 3832973499
Content-Length: 264
Content-Type: application/xml

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>Enabled</Status>
    <Prefix>2017</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-
1:3333333:domain/records/current/2017</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

Related information

[Use a tenant account](#)

PUT Bucket metadata notification configuration

The PUT Bucket metadata notification configuration request allows you to enable the search integration service for individual buckets. The metadata notification configuration XML that you supply in the request body specifies the objects whose metadata is sent to the destination search index.

You must have the `s3:PutBucketMetadataNotification` permission for a bucket, or be account root, to complete this operation.

Request

The request must include the metadata notification configuration in the request body. Each metadata notification configuration includes one or more rules. Each rule specifies the objects that it applies to, and the destination where StorageGRID should send object metadata.

Objects can be filtered on the prefix of the object name. For example, you could send metadata for objects with the prefix `/images` to one destination, and objects with the prefix `/videos` to another.

Configurations that have overlapping prefixes aren't valid, and are rejected when they are submitted. For example, a configuration that included one rule for objects with the prefix `test` and a second rule for objects with the prefix `test2` would not be allowed.

Destinations must be specified using the URN of a StorageGRID endpoint. The endpoint must exist when the metadata notification configuration is submitted, or the request fails as a `400 Bad Request`. The error

message states: Unable to save the metadata notification (search) policy. The specified endpoint URN does not exist: *URN*.

```
<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>rule-status</Status>
    <Prefix>key-prefix</Prefix>
    <Destination>
      <Urn>arn:aws:es:region:account-
ID:domain/mydomain/myindex/mytype</Urn>
    </Destination>
  </Rule>
  <Rule>
    <ID>Rule-2</ID>
    ...
  </Rule>
  ...
</MetadataNotificationConfiguration>
```

The table describes the elements in the metadata notification configuration XML.

Name	Description	Required
MetadataNotificationConfi guration	Container tag for rules used to specify the objects and destination for metadata notifications. Contains one or more Rule elements.	Yes
Rule	Container tag for a rule that identifies the objects whose metadata should be added to a specified index. Rules with overlapping prefixes are rejected. Included in the MetadataNotificationConfiguration element.	Yes
ID	Unique identifier for the rule. Included in the Rule element.	No
Status	Status can be 'Enabled' or 'Disabled'. No action is taken for rules that are disabled. Included in the Rule element.	Yes

Name	Description	Required
Prefix	Objects that match the prefix are affected by the rule, and their metadata is sent to the specified destination. To match all objects, specify an empty prefix. Included in the Rule element.	Yes
Destination	Container tag for the destination of a rule. Included in the Rule element.	Yes
Urn	URN of the destination where object metadata is sent. Must be the URN of a StorageGRID endpoint with the following properties: • <code>es</code> must be the third element. • The URN must end with the index and type where the metadata is stored, in the form <code>domain-name/myindex/mytype</code>. Endpoints are configured using the Tenant Manager or Tenant Management API. They take the following form: • <code>arn:aws:es:region:account-ID:domain/mydomain/myindex/mytype</code> • <code>urn:mysite:es:::mydomain/myindex/mytype</code> The endpoint must be configured before the configuration XML is submitted, or configuration will fail with a 404 error. Urn is included in the Destination element.	Yes

Request examples

This example shows enabling search integration for a bucket. In this example, object metadata for all objects is sent to the same destination.

```

PUT /test1?x-ntap-sg-metadata-notification HTTP/1.1
Date: date
Authorization: authorization string
Host: host

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>urn:sgws:es::sgws-notifications/test1/all</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

In this example, object metadata for objects that match the prefix `/images` is sent to one destination, while object metadata for objects that match the prefix `/videos` is sent to a second destination.

```

PUT /graphics?x-ntap-sg-metadata-notification HTTP/1.1
Date: date
Authorization: authorization string
Host: host

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Images-rule</ID>
    <Status>Enabled</Status>
    <Prefix>/images</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-1:33333333:domain/es-
domain/graphics/imagetype</Urn>
    </Destination>
  </Rule>
  <Rule>
    <ID>Videos-rule</ID>
    <Status>Enabled</Status>
    <Prefix>/videos</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-west-1:22222222:domain/es-
domain/graphics/videotype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```


JSON generated by the search integration service

When you enable the search integration service for a bucket, a JSON document is generated and sent to the destination endpoint each time object metadata or tags are added, updated, or deleted.

This example shows an example of the JSON that could be generated when an object with the key `SGWS/Tagging.txt` is created in a bucket named `test`. The `test` bucket is not versioned, so the `versionId` tag is empty.

```
{
  "bucket": "test",
  "key": "SGWS/Tagging.txt",
  "versionId": "",
  "accountId": "86928401983529626822",
  "size": 38,
  "md5": "3d6c7634a85436eee06d43415012855",
  "region": "us-east-1",
  "metadata": {
    "age": "25"
  },
  "tags": {
    "color": "yellow"
  }
}
```

Object metadata included in metadata notifications

The table lists all the fields that are included in the JSON document that is sent to the destination endpoint when search integration is enabled.

The document name includes the bucket name, object name, and version ID if present.

Type	Item name	Description
Bucket and object information	bucket	Name of the bucket
Bucket and object information	key	Object key name
Bucket and object information	versionID	Object version, for objects in versioned buckets
Bucket and object information	region	Bucket region, for example <code>us-east-1</code>
System metadata	size	Object size (in bytes) as visible to an HTTP client

Type	Item name	Description
System metadata	md5	Object hash
User metadata	metadata <i>key:value</i>	All user metadata for the object, as key-value pairs
Tags	tags <i>key:value</i>	All object tags defined for the object, as key-value pairs



For tags and user metadata, StorageGRID passes dates and numbers to Elasticsearch as strings or as S3 event notifications. To configure Elasticsearch to interpret these strings as dates or numbers, follow the Elasticsearch instructions for dynamic field mapping and for mapping date formats. You must enable the dynamic field mappings on the index before you configure the search integration service. After a document is indexed, you can't edit the document's field types in the index.

Related information

[Use a tenant account](#)

GET Storage Usage request

The GET Storage Usage request tells you the total amount of storage in use by an account, and for each bucket associated with the account.

The amount of storage used by an account and its buckets can be obtained by a modified ListBuckets request with the `x-ntap-sg-usage` query parameter. Bucket storage usage is tracked separately from the PUT and DELETE requests processed by the system. There might be some delay before the usage values match the expected values based on the processing of requests, particularly if the system is under heavy load.

By default, StorageGRID attempts to retrieve usage information using strong-global consistency. If strong-global consistency can't be achieved, StorageGRID attempts to retrieve the usage information at a strong-site consistency.

You must have the `s3:ListAllMyBuckets` permission, or be account root, to complete this operation.

Request example

```
GET /?x-ntap-sg-usage HTTP/1.1
Date: date
Authorization: authorization string
Host: host
```

Response example

This example shows an account that has four objects and 12 bytes of data in two buckets. Each bucket contains two objects and six bytes of data.

```
HTTP/1.1 200 OK
Date: Sat, 29 Nov 2015 00:49:05 GMT
Connection: KEEP-ALIVE
Server: StorageGRID/10.2.0
x-amz-request-id: 727237123
Content-Length: 427
Content-Type: application/xml

<?xml version="1.0" encoding="UTF-8"?>
<UsageResult xmlns="http://s3.storagegrid.com/doc/2015-02-01">
<CalculationTime>2014-11-19T05:30:11.000000Z</CalculationTime>
<ObjectCount>4</ObjectCount>
<DataBytes>12</DataBytes>
<Buckets>
<Bucket>
<Name>bucket1</Name>
<ObjectCount>2</ObjectCount>
<DataBytes>6</DataBytes>
</Bucket>
<Bucket>
<Name>bucket2</Name>
<ObjectCount>2</ObjectCount>
<DataBytes>6</DataBytes>
</Bucket>
</Buckets>
</UsageResult>
```

Versioning

Every object version stored will contribute to the `ObjectCount` and `DataBytes` values in the response. Delete markers aren't added to the `ObjectCount` total.

Related information

[Consistency values](#)

Deprecated bucket requests for legacy Compliance

Deprecated bucket requests for legacy Compliance

You might need to use the StorageGRID S3 REST API to manage buckets that were created using the legacy Compliance feature.

Compliance feature deprecated

The StorageGRID Compliance feature that was available in previous StorageGRID versions is deprecated and has been replaced by S3 Object Lock.

If you previously enabled the global Compliance setting, the global S3 Object Lock setting is enabled in StorageGRID 11.6. You can no longer create new buckets with Compliance enabled; however, as required, you can use the StorageGRID S3 REST API to manage any existing legacy Compliant buckets.

- [Use S3 REST API to configure S3 Object Lock](#)
- [Manage objects with ILM](#)
- [NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

Deprecated compliance requests:

- [Deprecated - PUT Bucket request modifications for compliance](#)

The SGCompliance XML element is deprecated. Previously, you could include this StorageGRID custom element in the optional XML request body of PUT Bucket requests to create a Compliant bucket.

- [Deprecated - GET Bucket compliance](#)

The GET Bucket compliance request is deprecated. However, you can continue to use this request to determine the compliance settings currently in effect for an existing legacy Compliant bucket.

- [Deprecated - PUT Bucket compliance](#)

The PUT Bucket compliance request is deprecated. However, you can continue to use this request to modify the compliance settings for an existing legacy Compliant bucket. For example, you can place an existing bucket on legal hold or increase its retention period.

Deprecated: CreateBucket request modifications for compliance

The SGCompliance XML element is deprecated. Previously, you could include this StorageGRID custom element in the optional XML request body of CreateBucket requests to create a Compliant bucket.



The StorageGRID Compliance feature that was available in previous StorageGRID versions is deprecated and has been replaced by S3 Object Lock. See the following for more details:

- [Use S3 REST API to configure S3 Object Lock](#)
- [NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

You can no longer create new buckets with Compliance enabled. The following error message is returned if you attempt to use the CreateBucket request modifications for compliance to create a new Compliant bucket:

```
The Compliance feature is deprecated.  
Contact your StorageGRID administrator if you need to create new Compliant  
buckets.
```

Deprecated: GET Bucket compliance request

The GET Bucket compliance request is deprecated. However, you can continue to use this request to determine the compliance settings currently in effect for an existing legacy

Compliant bucket.



The StorageGRID Compliance feature that was available in previous StorageGRID versions is deprecated and has been replaced by S3 Object Lock. See the following for more details:

- [Use S3 REST API to configure S3 Object Lock](#)
- [NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

You must have the `s3:GetBucketCompliance` permission, or be account root, to complete this operation.

Request example

This example request allows you to determine the compliance settings for the bucket named `mybucket`.

```
GET /mybucket/?x-ntap-sg-compliance HTTP/1.1
Date: date
Authorization: authorization string
Host: host
```

Response example

In the response XML, `<SGCompliance>` lists the compliance settings in effect for the bucket. This example response shows the compliance settings for a bucket in which each object will be retained for one year (525,600 minutes), starting from when the object is ingested into the grid. There is currently no legal hold on this bucket. Each object will be automatically deleted after one year.

```
HTTP/1.1 200 OK
Date: date
Connection: connection
Server: StorageGRID/11.1.0
x-amz-request-id: request ID
Content-Length: length
Content-Type: application/xml

<SGCompliance>
  <RetentionPeriodMinutes>525600</RetentionPeriodMinutes>
  <LegalHold>false</LegalHold>
  <AutoDelete>true</AutoDelete>
</SGCompliance>
```

Name	Description
RetentionPeriodMinutes	The length of the retention period for objects added to this bucket, in minutes. The retention period starts when the object is ingested into the grid.

Name	Description
LegalHold	• True: This bucket is currently under a legal hold. Objects in this bucket can't be deleted until the legal hold is lifted, even if their retention period has expired. • False: This bucket is not currently under a legal hold. Objects in this bucket can be deleted when their retention period expires.
AutoDelete	• True: The objects in this bucket will be deleted automatically when their retention period expires, unless the bucket is under a legal hold. • False: The objects in this bucket will not be deleted automatically when the retention period expires. You must delete these objects manually if you need to delete them.

Error responses

If the bucket was not created to be compliant, the HTTP status code for the response is 404 Not Found, with an S3 error code of `XNoSuchBucketCompliance`.

Deprecated: PUT Bucket compliance request

The PUT Bucket compliance request is deprecated. However, you can continue to use this request to modify the compliance settings for an existing legacy Compliant bucket. For example, you can place an existing bucket on legal hold or increase its retention period.



The StorageGRID Compliance feature that was available in previous StorageGRID versions is deprecated and has been replaced by S3 Object Lock. See the following for more details:

- [Use S3 REST API to configure S3 Object Lock](#)
- [NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

You must have the `s3:PutBucketCompliance` permission, or be account root, to complete this operation.

You must specify a value for every field of the compliance settings when issuing a PUT Bucket compliance request.

Request example

This example request modifies the compliance settings for the bucket named `mybucket`. In this example, objects in `mybucket` will now be retained for two years (1,051,200 minutes) instead of one year, starting from when the object is ingested into the grid. There is no legal hold on this bucket. Each object will be automatically deleted after two years.

```

PUT /mybucket/?x-ntap-sg-compliance HTTP/1.1
Date: date
Authorization: authorization name
Host: host
Content-Length: 152

<SGCompliance>
  <RetentionPeriodMinutes>1051200</RetentionPeriodMinutes>
  <LegalHold>false</LegalHold>
  <AutoDelete>true</AutoDelete>
</SGCompliance>

```

Name	Description
RetentionPeriodMinutes	The length of the retention period for objects added to this bucket, in minutes. The retention period starts when the object is ingested into the grid. Important When specifying a new value for RetentionPeriodMinutes, you must specify a value that is equal to or greater than the bucket's current retention period. After the bucket's retention period is set, you can't decrease that value; you can only increase it.
LegalHold	• True: This bucket is currently under a legal hold. Objects in this bucket can't be deleted until the legal hold is lifted, even if their retention period has expired. • False: This bucket is not currently under a legal hold. Objects in this bucket can be deleted when their retention period expires.
AutoDelete	• True: The objects in this bucket will be deleted automatically when their retention period expires, unless the bucket is under a legal hold. • False: The objects in this bucket will not be deleted automatically when the retention period expires. You must delete these objects manually if you need to delete them.

Consistency for compliance settings

When you update the compliance settings for an S3 bucket with a PUT Bucket compliance request, StorageGRID attempts to update the bucket's metadata across the grid. By default, StorageGRID uses the **Strong-global** consistency to guarantee that all data center sites and all Storage Nodes that contain bucket metadata have read-after-write consistency for the changed compliance settings.

If StorageGRID can't achieve the **Strong-global** consistency because a data center site or multiple Storage Nodes at a site are unavailable, the HTTP status code for the response is 503 *Service Unavailable*.

If you receive this response, you must contact the grid administrator to ensure that the required storage services are made available as soon as possible. If the grid administrator is unable to make enough of the

Storage Nodes at each site available, technical support might direct you to retry the failed request by forcing the **Strong-site** consistency.



Never force the **Strong-site** consistency for PUT bucket compliance unless you have been directed to do so by technical support and unless you understand the potential consequences of using this level.

When the consistency is reduced to **Strong-site**, StorageGRID guarantees that updated compliance settings will have read-after-write consistency only for client requests within a site. This means that the StorageGRID system might temporarily have multiple, inconsistent settings for this bucket until all sites and Storage Nodes are available. The inconsistent settings can result in unexpected and undesired behavior. For example, if you are placing a bucket under a legal hold and you force a lower consistency, the bucket's previous compliance settings (that is, legal hold off) might continue to be in effect at some data center sites. As a result, objects that you think are on legal hold might be deleted when their retention period expires, either by the user or by AutoDelete, if enabled.

To force the use of the **Strong-site** consistency, reissue the PUT Bucket compliance request and include the Consistency-Control HTTP request header, as follows:

```
PUT /mybucket/?x-ntap-sg-compliance HTTP/1.1
Consistency-Control: strong-site
```

Error responses

- If the bucket was not created to be compliant, the HTTP status code for the response is 404 Not Found.
- If `RetentionPeriodMinutes` in the request is less than the bucket's current retention period, the HTTP status code is 400 Bad Request.

Related information

[Deprecated: PUT Bucket request modifications for compliance](#)

Bucket and group access policies

Use bucket and group access policies

StorageGRID uses the Amazon Web Services (AWS) policy language to allow S3 tenants to control access to buckets and objects within those buckets. The StorageGRID system implements a subset of the S3 REST API policy language. Access policies for the S3 API are written in JSON.

Access policy overview

There are two kinds of access policies supported by StorageGRID.

- **Bucket policies**, which are managed using the `GetBucketPolicy`, `PutBucketPolicy`, and `DeleteBucketPolicy` S3 API operations or the Tenant Manager or Tenant Management API. Bucket policies are attached to buckets, so they are configured to control access by users in the bucket owner account or other accounts to the bucket and the objects in it. A bucket policy applies to only one bucket and possibly multiple groups.
- **Group policies**, which are configured using the Tenant Manager or Tenant Management API. Group

policies are attached to a group in the account, so they are configured to allow that group to access specific resources owned by that account. A group policy applies to only one group and possibly multiple buckets.



There is no difference in priority between group and bucket policies.

StorageGRID bucket and group policies follow a specific grammar defined by Amazon. Inside each policy is an array of policy statements, and each statement contains the following elements:

- Statement ID (Sid) (optional)
- Effect
- Principal/NotPrincipal
- Resource/NotResource
- Action/NotAction
- Condition (optional)

Policy statements are built using this structure to specify permissions: Grant <Effect> to allow/deny <Principal> to perform <Action> on <Resource> when <Condition> applies.

Each policy element is used for a specific function:

Element	Description
Sid	The Sid element is optional. The Sid is only intended as a description for the user. It is stored but not interpreted by the StorageGRID system.
Effect	Use the Effect element to establish whether the specified operations are allowed or denied. You must identify operations you allow (or deny) on buckets or objects using the supported Action element keywords.
Principal/NotPrincipal	You can allow users, groups, and accounts to access specific resources and perform specific actions. If no S3 signature is included in the request, anonymous access is allowed by specifying the wildcard character (*) as the principal. By default, only the account root has access to resources owned by the account. You only need to specify the Principal element in a bucket policy. For group policies, the group to which the policy is attached is the implicit Principal element.
Resource/NotResource	The Resource element identifies buckets and objects. You can allow or deny permissions to buckets and objects using the Amazon Resource Name (ARN) to identify the resource.
Action/NotAction	The Action and Effect elements are the two components of permissions. When a group requests a resource, they are either granted or denied access to the resource. Access is denied unless you specifically assign permissions, but you can use explicit deny to override a permission granted by another policy.

Element	Description
Condition	The Condition element is optional. Conditions allow you to build expressions to determine when a policy should be applied.

In the Action element, you can use the wildcard character (*) to specify all operations, or a subset of operations. For example, this Action matches permissions such as `s3:GetObject`, `s3:PutObject`, and `s3:DeleteObject`.

```
s3:*Object
```

In the Resource element, you can use the wildcard characters (*) and (?). While the asterisk (*) matches 0 or more characters, the question mark (?) matches any single character.

In the Principal element, wildcard characters aren't supported except to set anonymous access, which grants permission to everyone. For example, you set the wildcard (*) as the Principal value.

```
"Principal": "*"

```

```
"Principal":{"AWS": "*"

```

In the following example, the statement is using the Effect, Principal, Action, and Resource elements. This example shows a complete bucket policy statement that uses the Effect "Allow" to give the Principals, the admin group `federated-group/admin` and the finance group `federated-group/finance`, permissions to perform the Action `s3:ListBucket` on the bucket named `mybucket` and the Action `s3:GetObject` on all objects inside that bucket.

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": [
          "arn:aws:iam::27233906934684427525:federated-group/admin",
          "arn:aws:iam::27233906934684427525:federated-group/finance"
        ]
      },
      "Action": [
        "s3:ListBucket",
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::mybucket",
        "arn:aws:s3:::mybucket/*"
      ]
    }
  ]
}
```

The bucket policy has a size limit of 20,480 bytes, and the group policy has a size limit of 5,120 bytes.

Consistency for policies

By default, any updates you make to group policies are eventually consistent. When a group policy becomes consistent, the changes can take an additional 15 minutes to take effect, because of policy caching. By default, any updates you make to bucket policies are strongly consistent.

As required, you can change the consistency guarantees for bucket policy updates. For example, you might want a change to a bucket policy to be available during a site outage.

In this case, you can either set the `Consistency-Control` header in the `PutBucketPolicy` request, or you can use the `PUT Bucket` consistency request. When a bucket policy becomes consistent, the changes can take an additional 8 seconds to take effect, because of policy caching.



If you set the consistency to a different value to work around a temporary situation, be sure to set the bucket-level setting back to its original value when you are done. Otherwise, all future bucket requests will use the modified setting.

Use ARN in policy statements

In policy statements, the ARN is used in `Principal` and `Resource` elements.

- Use this syntax to specify the S3 resource ARN:

```
arn:aws:s3:::bucket-name
arn:aws:s3:::bucket-name/object_key
```

- Use this syntax to specify the identity resource ARN (users and groups):

```
arn:aws:iam::account_id:root
arn:aws:iam::account_id:user/user_name
arn:aws:iam::account_id:group/group_name
arn:aws:iam::account_id:federated-user/user_name
arn:aws:iam::account_id:federated-group/group_name
```

Other considerations:

- You can use the asterisk (*) as a wildcard to match zero or more characters inside the object key.
- International characters, which can be specified in the object key, should be encoded using JSON UTF-8 or using JSON `\u` escape sequences. Percent-encoding is not supported.

[RFC 2141 URN Syntax](#)

The HTTP request body for the `PutBucketPolicy` operation must be encoded with `charset=UTF-8`.

Specify resources in a policy

In policy statements, you can use the `Resource` element to specify the bucket or object for which permissions are allowed or denied.

- Each policy statement requires a `Resource` element. In a policy, resources are denoted by the element `Resource`, or alternatively, `NotResource` for exclusion.
- You specify resources with an S3 resource ARN. For example:

```
"Resource": "arn:aws:s3:::mybucket/*"
```

- You can also use policy variables inside the object key. For example:

```
"Resource": "arn:aws:s3:::mybucket/home/${aws:username}/*"
```

- The resource value can specify a bucket that does not yet exist when a group policy is created.

Specify principals in a policy

Use the `Principal` element to identify the user, group, or tenant account that is allowed/denied access to the resource by the policy statement.

- Each policy statement in a bucket policy must include a `Principal` element. Policy statements in a group policy don't need the `Principal` element because the group is understood to be the principal.

- In a policy, principals are denoted by the element "Principal," or alternatively "NotPrincipal" for exclusion.
- Account-based identities must be specified using an ID or an ARN:

```
"Principal": { "AWS": "account_id"}
"Principal": { "AWS": "identity_arn" }
```

- This example uses the tenant account ID 27233906934684427525, which includes the account root and all users in the account:

```
"Principal": { "AWS": "27233906934684427525" }
```

- You can specify just the account root:

```
"Principal": { "AWS": "arn:aws:iam::27233906934684427525:root" }
```

- You can specify a specific federated user ("Alex"):

```
"Principal": { "AWS": "arn:aws:iam::27233906934684427525:federated-
user/Alex" }
```

- You can specify a specific federated group ("Managers"):

```
"Principal": { "AWS": "arn:aws:iam::27233906934684427525:federated-
group/Managers" }
```

- You can specify an anonymous principal:

```
"Principal": "*" 
```

- To avoid ambiguity, you can use the user UUID instead of the username:

```
arn:aws:iam::27233906934684427525:user-uuid/de305d54-75b4-431b-adb2-
eb6b9e546013
```

For example, suppose Alex leaves the organization and the username `Alex` is deleted. If a new Alex joins the organization and is assigned the same `Alex` username, the new user might unintentionally inherit the permissions granted to the original user.

- The principal value can specify a group/user name that does not yet exist when a bucket policy is created.

Specify permissions in a policy

In a policy, the Action element is used to allow/deny permissions to a resource. There are a set of permissions that you can specify in a policy, which are denoted by the element "Action," or alternatively, "NotAction" for exclusion. Each of these elements maps to specific S3 REST API operations.

The tables lists the permissions that apply to buckets and the permissions that apply to objects.



Amazon S3 now uses the s3:PutReplicationConfiguration permission for both the PutBucketReplication and DeleteBucketReplication actions. StorageGRID uses separate permissions for each action, which matches the original Amazon S3 specification.



A delete is performed when a put is used to overwrite an existing value.

Permissions that apply to buckets

Permissions	S3 REST API operations	Custom for StorageGRID
s3:CreateBucket	CreateBucket	Yes. Note: Use in group policy only.
s3>DeleteBucket	DeleteBucket	
s3>DeleteBucketMetadataNotification	DELETE Bucket metadata notification configuration	Yes
s3>DeleteBucketPolicy	DeleteBucketPolicy	
s3>DeleteReplicationConfiguration	DeleteBucketReplication	Yes, separate permissions for PUT and DELETE
s3:GetBucketAcl	GetBucketAcl	
s3:GetBucketCompliance	GET Bucket compliance (deprecated)	Yes
s3:GetBucketConsistency	GET Bucket consistency	Yes
s3:GetBucketCORS	GetBucketCors	
s3:GetEncryptionConfiguration	GetBucketEncryption	
s3:GetBucketLastAccessTime	GET Bucket last access time	Yes
s3:GetBucketLocation	GetBucketLocation	

Permissions	S3 REST API operations	Custom for StorageGRID
s3:GetBucketMetadataNotification	GET Bucket metadata notification configuration	Yes
s3:GetBucketNotification	GetBucketNotificationConfiguration	
s3:GetBucketObjectLockConfiguration	GetObjectLockConfiguration	
s3:GetBucketPolicy	GetBucketPolicy	
s3:GetBucketTagging	GetBucketTagging	
s3:GetBucketVersioning	GetBucketVersioning	
s3:GetLifecycleConfiguration	GetBucketLifecycleConfiguration	
s3:GetReplicationConfiguration	GetBucketReplication	
s3:ListAllMyBuckets	- ListBuckets - GET Storage Usage	Yes, for GET Storage Usage. Note: Use in group policy only.
s3:ListBucket	- ListObjects - HeadBucket - RestoreObject	
s3:ListBucketMultipartUploads	- ListMultipartUploads - RestoreObject	
s3:ListBucketVersions	GET Bucket versions	
s3:PutBucketCompliance	PUT Bucket compliance (deprecated)	Yes
s3:PutBucketConsistency	PUT Bucket consistency	Yes
s3:PutBucketCORS	- DeleteBucketCors† - PutBucketCors	
s3:PutEncryptionConfiguration	- DeleteBucketEncryption - PutBucketEncryption	

Permissions	S3 REST API operations	Custom for StorageGRID
s3:PutBucketLastAccessTime	PUT Bucket last access time	Yes
s3:PutBucketMetadataNotification	PUT Bucket metadata notification configuration	Yes
s3:PutBucketNotification	PutBucketNotificationConfiguration	
s3:PutBucketObjectLockConfiguration	- CreateBucket with the <code>x-amz-bucket-object-lock-enabled: true</code> request header (also requires the <code>s3:CreateBucket</code> permission) - PutObjectLockConfiguration	
s3:PutBucketPolicy	PutBucketPolicy	
s3:PutBucketTagging	- DeleteBucketTagging† - PutBucketTagging	
s3:PutBucketVersioning	PutBucketVersioning	
s3:PutLifecycleConfiguration	- DeleteBucketLifecycle† - PutBucketLifecycleConfiguration	
s3:PutReplicationConfiguration	PutBucketReplication	Yes, separate permissions for PUT and DELETE

Permissions that apply to objects

Permissions	S3 REST API operations	Custom for StorageGRID
s3:AbortMultipartUpload	- AbortMultipartUpload - RestoreObject	
s3:BypassGovernanceRetention	- DeleteObject - DeleteObjects - PutObjectRetention	
s3>DeleteObject	- DeleteObject - DeleteObjects - RestoreObject	

Permissions	S3 REST API operations	Custom for StorageGRID
s3:DeleteObjectTagging	DeleteObjectTagging	
s3:DeleteObjectVersionTagging	DeleteObjectTagging (a specific version of the object)	
s3:DeleteObjectVersion	DeleteObject (a specific version of the object)	
s3:GetObject	• GetObject • HeadObject • RestoreObject • SelectObjectContent	
s3:GetObjectAcl	GetObjectAcl	
s3:GetObjectLegalHold	GetObjectLegalHold	
s3:GetObjectRetention	GetObjectRetention	
s3:GetObjectTagging	GetObjectTagging	
s3:GetObjectVersionTagging	GetObjectTagging (a specific version of the object)	
s3:GetObjectVersion	GetObject (a specific version of the object)	
s3:ListMultipartUploadParts	ListParts, RestoreObject	
s3:PutObject	• PutObject • CopyObject • RestoreObject • CreateMultipartUpload • CompleteMultipartUpload • UploadPart • UploadPartCopy	
s3:PutObjectLegalHold	PutObjectLegalHold	
s3:PutObjectRetention	PutObjectRetention	
s3:PutObjectTagging	PutObjectTagging	

Permissions	S3 REST API operations	Custom for StorageGRID
s3:PutObjectVersionTagging	PutObjectTagging (a specific version of the object)	
s3:PutOverwriteObject	• PutObject • CopyObject • PutObjectTagging • DeleteObjectTagging • CompleteMultipartUpload	Yes
s3:RestoreObject	RestoreObject	

Use PutOverwriteObject permission

The s3:PutOverwriteObject permission is a custom StorageGRID permission that applies to operations that create or update objects. The setting of this permission determines whether the client can overwrite an object's data, user-defined metadata, or S3 object tagging.

Possible settings for this permission include:

- **Allow:** The client can overwrite an object. This is the default setting.
- **Deny:** The client can't overwrite an object. When set to Deny, the PutOverwriteObject permission works as follows:
 - If an existing object is found at the same path:
 - The object's data, user-defined metadata, or S3 object tagging can't be overwritten.
 - Any ingest operations in progress are cancelled, and an error is returned.
 - If S3 versioning is enabled, the Deny setting prevents PutObjectTagging or DeleteObjectTagging operations from modifying the TagSet for an object and its noncurrent versions.
 - If an existing object is not found, this permission has no effect.
- When this permission is not present, the effect is the same as if Allow were set.



If the current S3 policy allows overwrite, and the PutOverwriteObject permission is set to Deny, the client can't overwrite an object's data, user-defined metadata, or object tagging. In addition, if the **Prevent client modification** checkbox is selected (**CONFIGURATION > Security settings > Network and objects**), that setting overrides the setting of the PutOverwriteObject permission.

Specify conditions in a policy

Conditions define when a policy will be in effect. Conditions consist of operators and key-value pairs.

Conditions use key-value pairs for evaluation. A Condition element can contain multiple conditions, and each condition can contain multiple key-value pairs. The condition block uses the following format:

```
Condition: {
  condition_type: {
    condition_key: condition_values
```

In the following example, the `IpAddress` condition uses the `SourceIp` condition key.

```
"Condition": {
  "IpAddress": {
    "aws:SourceIp": "54.240.143.0/24"
    ...
  },
  ...
```

Supported condition operators

Condition operators are categorized as follows:

- String
- Numeric
- Boolean
- IP address
- Null check

Condition operators	Description
StringEquals	Compares a key to a string value based on exact matching (case sensitive).
StringNotEquals	Compares a key to a string value based on negated matching (case sensitive).
StringEqualsIgnoreCase	Compares a key to a string value based on exact matching (ignores case).
StringNotEqualsIgnoreCase	Compares a key to a string value based on negated matching (ignores case).
StringLike	Compares a key to a string value based on exact matching (case sensitive). Can include * and ? wildcard characters.
StringNotLike	Compares a key to a string value based on negated matching (case sensitive). Can include * and ? wildcard characters.
NumericEquals	Compares a key to a numeric value based on exact matching.

Condition operators	Description
NumericNotEquals	Compares a key to a numeric value based on negated matching.
NumericGreaterThan	Compares a key to a numeric value based on "greater than" matching.
NumericGreaterThanEquals	Compares a key to a numeric value based on "greater than or equals" matching.
NumericLessThan	Compares a key to a numeric value based on "less than" matching.
NumericLessThanEquals	Compares a key to a numeric value based on "less than or equals" matching.
Bool	Compares a key to a Boolean value based on "true or false" matching.
IpAddress	Compares a key to an IP address or range of IP addresses.
NotIpAddress	Compares a key to an IP address or range of IP addresses based on negated matching.
Null	Checks if a condition key is present in the current request context.

Supported condition keys

Condition keys	Actions	Description
aws:SourceIp	IP operators	Will compare to the IP address from which the request was sent. Can be used for bucket or object operations. Note: If the S3 request was sent through the Load Balancer service on Admin Nodes and Gateways Nodes, this will compare to the IP address upstream of the Load Balancer service. Note: If a third-party, non-transparent load balancer is used, this will compare to the IP address of that load balancer. Any <code>X-Forwarded-For</code> header will be ignored because its validity can't be ascertained.
aws:username	Resource/Identity	Will compare to the sender's username from which the request was sent. Can be used for bucket or object operations.
s3:delimiter	s3:ListBucket and s3:ListBucketVersions permissions	Will compare to the delimiter parameter specified in a ListObjects or ListObjectVersions request.

Condition keys	Actions	Description
s3:ExistingObjectTag/<tag-key>	s3:DeleteObjectTagging s3:DeleteObjectVersionTagging s3:GetObject s3:GetObjectAcl s3:GetObjectTagging s3:GetObjectVersion s3:GetObjectVersionAcl s3:GetObjectVersionTagging s3:PutObjectAcl s3:PutObjectTagging s3:PutObjectVersionAcl s3:PutObjectVersionTagging	Will require that the existing object has the specific tag key and value.
s3:max-keys	s3:ListBucket and s3:ListBucketVersions permissions	Will compare to the max-keys parameter specified in a ListObjects or ListObjectVersions request.
s3:object-lock-remaining-retention-days	s3:PutObject	Compares to the retain-until-date specified in the x-amz-object-lock-retain-until-date request header or computed from the bucket default retention period to make sure that these values are within the allowable range for the following requests: • PutObject • CopyObject • CreateMultipartUpload
s3:object-lock-remaining-retention-days	s3:PutObjectRetention	Compares to the retain-until-date specified in the PutObjectRetention request to ensure that it is within the allowable range.
s3:prefix	s3:ListBucket and s3:ListBucketVersions permissions	Will compare to the prefix parameter specified in a ListObjects or ListObjectVersions request.

Condition keys	Actions	Description
s3:RequestObjectTag/<tag-key>	s3:PutObject s3:PutObjectTagging s3:PutObjectVersionTagging	Will require a specific tag key and value when the object request includes tagging.

Specify variables in a policy

You can use variables in policies to populate policy information when it is available. You can use policy variables in the `Resource` element and in string comparisons in the `Condition` element.

In this example, the variable `${aws:username}` is part of the `Resource` element:

```
"Resource": "arn:aws:s3:::bucket-name/home/${aws:username}/*"
```

In this example, the variable `${aws:username}` is part of the condition value in the condition block:

```
"Condition": {
  "StringLike": {
    "s3:prefix": "${aws:username}/*"
    ...
  },
  ...
}
```

Variable	Description
<code>\${aws:SourceIp}</code>	Uses the <code>SourceIp</code> key as the provided variable.
<code>\${aws:username}</code>	Uses the <code>username</code> key as the provided variable.
<code>\${s3:prefix}</code>	Uses the service-specific <code>prefix</code> key as the provided variable.
<code>\${s3:max-keys}</code>	Uses the service-specific <code>max-keys</code> key as the provided variable.
<code>\${*}</code>	Special character. Uses the character as a literal <code>*</code> character.
<code>\${?}</code>	Special character. Uses the character as a literal <code>?</code> character.
<code>\${\$}</code>	Special character. Uses the character as a literal <code>\$</code> character.

Create policies requiring special handling

Sometimes a policy can grant permissions that are dangerous for security or dangerous for continued operations, such as locking out the root user of the account. The StorageGRID S3 REST API implementation is less restrictive during policy validation than Amazon, but equally strict during policy evaluation.

Policy description	Policy type	Amazon behavior	StorageGRID behavior
Deny self any permissions to the root account	Bucket	Valid and enforced, but root user account retains permission for all S3 bucket policy operations	Same
Deny self any permissions to user/group	Group	Valid and enforced	Same
Allow a foreign account group any permission	Bucket	Invalid principal	Valid, but permissions for all S3 bucket policy operations return a 405 Method Not Allowed error when allowed by a policy
Allow a foreign account root or user any permission	Bucket	Valid, but permissions for all S3 bucket policy operations return a 405 Method Not Allowed error when allowed by a policy	Same
Allow everyone permissions to all actions	Bucket	Valid, but permissions for all S3 bucket policy operations return a 405 Method Not Allowed error for the foreign account root and users	Same
Deny everyone permissions to all actions	Bucket	Valid and enforced, but root user account retains permission for all S3 bucket policy operations	Same
Principal is a non-existent user or group	Bucket	Invalid principal	Valid
Resource is a non-existent S3 bucket	Group	Valid	Same
Principal is a local group	Bucket	Invalid principal	Valid

Policy description	Policy type	Amazon behavior	StorageGRID behavior
Policy grants a non-owner account (including anonymous accounts) permissions to put objects.	Bucket	Valid. Objects are owned by the creator account, and the bucket policy does not apply. The creator account must grant access permissions for the object using object ACLs.	Valid. Objects are owned by the bucket owner account. Bucket policy applies.

Write-once-read-many (WORM) protection

You can create write-once-read-many (WORM) buckets to protect data, user-defined object metadata, and S3 object tagging. You configure the WORM buckets to allow the creation of new objects and to prevent overwrites or deletion of existing content. Use one of the approaches described here.

To ensure that overwrites are always denied, you can:

- From the Grid Manager, go to **CONFIGURATION > Security > Security settings > Network and objects**, and select the **Prevent client modification** checkbox.
- Apply the following rules and S3 policies:
 - Add a PutOverwriteObject DENY operation to the S3 policy.
 - Add a DeleteObject DENY operation to the S3 policy.
 - Add a PutObject ALLOW operation to the S3 policy.



Setting DeleteObject to DENY in an S3 policy does not prevent ILM from deleting objects when a rule such as "zero copies after 30 days" exists.



Even when all of these rules and policies are applied, they don't guard against concurrent writes (see Situation A). They do guard against sequential completed overwrites (see Situation B).

Situation A: Concurrent writes (not guarded against)

```
/mybucket/important.doc
PUT#1 ---> OK
PUT#2 -----> OK
```

Situation B: Sequential completed overwrites (guarded against)

```
/mybucket/important.doc
PUT#1 -----> PUT#2 ---X (denied)
```

Related information

- [How StorageGRID ILM rules manage objects](#)
- [Example bucket policies](#)
- [Example group policies](#)

- [Manage objects with ILM](#)
- [Use a tenant account](#)

Example bucket policies

Use the examples in this section to build StorageGRID access policies for buckets.

Bucket policies specify the access permissions for the bucket that the policy is attached to. You configure a bucket policy by using the S3 PutBucketPolicy API through one of these tools:

- [Tenant Manager](#).
- AWS CLI using this command (refer to [Operations on buckets](#)):

```
> aws s3api put-bucket-policy --bucket examplebucket --policy
file://policy.json
```

Example: Allow everyone read-only access to a bucket

In this example, everyone, including anonymous, is allowed to list objects in the bucket and perform GetObject operations on all objects in the bucket. All other operations will be denied. Note that this policy might not be particularly useful because no one except the account root has permissions to write to the bucket.

```
{
  "Statement": [
    {
      "Sid": "AllowEveryoneReadOnlyAccess",
      "Effect": "Allow",
      "Principal": "*",
      "Action": [ "s3:GetObject", "s3:ListBucket" ],
      "Resource":
[ "arn:aws:s3:::examplebucket", "arn:aws:s3:::examplebucket/*" ]
    }
  ]
}
```

Example: Allow everyone in one account full access, and everyone in another account read-only access to a bucket

In this example, everyone in one specified account is allowed full access to a bucket, while everyone in another specified account is only permitted to List the bucket and perform GetObject operations on objects in the bucket beginning with the `shared/` object key prefix.



In StorageGRID, objects created by a non-owner account (including anonymous accounts) are owned by the bucket owner account. The bucket policy applies to these objects.

```

{
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "95390887230002558202"
      },
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::examplebucket",
        "arn:aws:s3:::examplebucket/*"
      ]
    },
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "31181711887329436680"
      },
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::examplebucket/shared/*"
    },
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "31181711887329436680"
      },
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::examplebucket",
      "Condition": {
        "StringLike": {
          "s3:prefix": "shared/*"
        }
      }
    }
  ]
}

```

Example: Allow everyone read-only access to a bucket and full access by specified group

In this example, everyone including anonymous, is allowed to List the bucket and perform GetObject operations on all objects in the bucket, while only users belonging the group `Marketing` in the specified account are allowed full access.

```

{
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::95390887230002558202:federated-
group/Marketing"
      },
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::examplebucket",
        "arn:aws:s3:::examplebucket/*"
      ]
    },
    {
      "Effect": "Allow",
      "Principal": "*",
      "Action": ["s3:ListBucket", "s3:GetObject"],
      "Resource": [
        "arn:aws:s3:::examplebucket",
        "arn:aws:s3:::examplebucket/*"
      ]
    }
  ]
}

```

Example: Allow everyone read and write access to a bucket if client in IP range

In this example, everyone, including anonymous, is allowed to List the bucket and perform any Object operations on all objects in the bucket, provided that the requests come from a specified IP range (54.240.143.0 to 54.240.143.255, except 54.240.143.188). All other operations will be denied, and all requests outside of the IP range will be denied.

```

{
  "Statement": [
    {
      "Sid": "AllowEveryoneReadWriteAccessIfInSourceIpRange",
      "Effect": "Allow",
      "Principal": "*",
      "Action": [ "s3:*Object", "s3:ListBucket" ],
      "Resource":
[ "arn:aws:s3:::examplebucket", "arn:aws:s3:::examplebucket/*" ],
      "Condition": {
        "IpAddress": { "aws:SourceIp": "54.240.143.0/24" },
        "NotIpAddress": { "aws:SourceIp": "54.240.143.188" }
      }
    }
  ]
}

```

Example: Allow full access to a bucket exclusively by a specified federated user

In this example, the federated user Alex is allowed full access to the `examplebucket` bucket and its objects. All other users, including 'root', are explicitly denied all operations. Note however that 'root' is never denied permissions to Put/Get/DeleteBucketPolicy.

```

{
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::95390887230002558202:federated-user/Alex"
      },
      "Action": [
        "s3:*"
      ],
      "Resource": [
        "arn:aws:s3:::examplebucket",
        "arn:aws:s3:::examplebucket/*"
      ]
    },
    {
      "Effect": "Deny",
      "NotPrincipal": {
        "AWS": "arn:aws:iam::95390887230002558202:federated-user/Alex"
      },
      "Action": [
        "s3:*"
      ],
      "Resource": [
        "arn:aws:s3:::examplebucket",
        "arn:aws:s3:::examplebucket/*"
      ]
    }
  ]
}

```

Example: PutOverwriteObject permission

In this example, the `Deny` Effect for `PutOverwriteObject` and `DeleteObject` ensures that no one can overwrite or delete the object's data, user-defined metadata, and S3 object tagging.

```

{
  "Statement": [
    {
      "Effect": "Deny",
      "Principal": "*",
      "Action": [
        "s3:PutOverwriteObject",
        "s3:DeleteObject",
        "s3:DeleteObjectVersion"
      ],
      "Resource": "arn:aws:s3:::wormbucket/*"
    },
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::95390887230002558202:federated-
group/SomeGroup"
      },
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::wormbucket"
    },
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::95390887230002558202:federated-
group/SomeGroup"
      },
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::wormbucket/*"
    }
  ]
}

```

Example group policies

Use the examples in this section to build StorageGRID access policies for groups.

Group policies specify the access permissions for the group that the policy is attached to. There is no `Principal` element in the policy because it is implicit. Group policies are configured using the Tenant Manager or the API.

Example: Set group policy using Tenant Manager

When you add or edit a group in the Tenant Manager, you can select a group policy to determine which S3 access permissions the members of this group will have. See [Create groups for an S3 tenant](#).

- **No S3 Access:** Default option. Users in this group don't have access to S3 resources, unless access is granted with a bucket policy. If you select this option, only the root user will have access to S3 resources by default.
- **Read Only Access:** Users in this group have read-only access to S3 resources. For example, users in this group can list objects and read object data, metadata, and tags. When you select this option, the JSON string for a read-only group policy appears in the text box. You can't edit this string.
- **Full Access:** Users in this group have full access to S3 resources, including buckets. When you select this option, the JSON string for a full-access group policy appears in the text box. You can't edit this string.
- **Ransomware Mitigation:** This sample policy applies to all buckets for this tenant. Users in this group can perform common actions, but can't permanently delete objects from buckets that have object versioning enabled.

Tenant Manager users who have the Manage all buckets permission can override this group policy. Limit the Manage all buckets permission to trusted users, and use Multi-Factor Authentication (MFA) where available.

- **Custom:** Users in the group are granted the permissions you specify in the text box.

Example: Allow group full access to all buckets

In this example, all members of the group are permitted full access to all buckets owned by the tenant account unless explicitly denied by bucket policy.

```
{
  "Statement": [
    {
      "Action": "s3:*",
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::*"
    }
  ]
}
```

Example: Allow group read-only access to all buckets

In this example, all members of the group have read-only access to S3 resources, unless explicitly denied by the bucket policy. For example, users in this group can list objects and read object data, metadata, and tags.

```

{
  "Statement": [
    {
      "Sid": "AllowGroupReadOnlyAccess",
      "Effect": "Allow",
      "Action": [
        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "s3:ListBucketVersions",
        "s3:GetObject",
        "s3:GetObjectTagging",
        "s3:GetObjectVersion",
        "s3:GetObjectVersionTagging"
      ],
      "Resource": "arn:aws:s3:::*"
    }
  ]
}

```

Example: Allow group members full access to only their "folder" in a bucket

In this example, members of the group are only permitted to list and access their specific folder (key prefix) in the specified bucket. Note that access permissions from other group policies and the bucket policy should be considered when determining the privacy of these folders.


```
{
  "Statement": [
    {
      "Sid": "AllowListBucketOfASpecificUserPrefix",
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::department-bucket",
      "Condition": {
        "StringLike": {
          "s3:prefix": "${aws:username}/*"
        }
      }
    },
    {
      "Sid": "AllowUserSpecificActionsOnlyInTheSpecificUserPrefix",
      "Effect": "Allow",
      "Action": "s3:*Object",
      "Resource": "arn:aws:s3:::department-bucket/${aws:username}/*"
    }
  ]
}
```

S3 operations tracked in the audit logs

Audit messages are generated by StorageGRID services and stored in text log files. You can review the S3-specific audit messages in the audit log to get details about bucket and object operations.

Bucket operations tracked in the audit logs

- CreateBucket
- DeleteBucket
- DeleteBucketTagging
- DeleteObjects
- GetBucketTagging
- HeadBucket
- ListObjects
- ListObjectVersions
- PUT Bucket compliance
- PutBucketTagging
- PutBucketVersioning

Object operations tracked in the audit logs

- CompleteMultipartUpload
- CopyObject
- DeleteObject
- GetObject
- HeadObject
- PutObject
- RestoreObject
- SelectObject
- UploadPart (when an ILM rule uses Balanced or Strict ingest)
- UploadPartCopy (when an ILM rule uses Balanced or Strict ingest)

Related information

- [Access audit log file](#)
- [Client write audit messages](#)
- [Client read audit messages](#)

Use Swift REST API (end of life)

Use Swift REST API

Support for the Swift API has reached end of life and will be removed in a future release.



Swift details have been removed from this version of the doc site. See [StorageGRID 11.8: Use Swift REST API](#).

Monitor and troubleshoot a StorageGRID system

Monitor StorageGRID system

Monitor a StorageGRID system

Monitor your StorageGRID system regularly to ensure it is performing as expected.

Before you begin

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



To change units for the storage values displayed in the Grid Manager, select the user drop-down in the upper right of the Grid Manager, then select **User preferences**.

About this task

These instructions describe how to:

- [View and manage the dashboard](#)
- [View the Nodes page](#)
- [Monitor these aspects of the system regularly](#):
 - [System health](#)
 - [Storage capacity](#)
 - [Information lifecycle management](#)
 - [Networking and system resources](#)
 - [Tenant activity](#)
 - [Load balancing operations](#)
 - [Grid federation connections](#)
- [Manage alerts](#)
- [View log files](#)
- [Configure audit messages and log destinations](#)
- [Use an external syslog server](#) to collect audit information
- [Use SNMP for monitoring](#)
- [Obtain additional StorageGRID data](#), including metrics and diagnostics


View and manage the dashboard

You can use the dashboard to monitor system activities at a glance. You can create custom dashboards to monitor your implementation of StorageGRID.



To change units for the storage values displayed in the Grid Manager, select the user drop-down in the upper right of the Grid Manager, then select **User preferences**.

Your dashboard might be different based on system configuration.



View the dashboard

The dashboard consists of tabs that contain specific information about the StorageGRID system. Each tab contains categories of information displayed on cards.

You can use the system-provided dashboard as is. Additionally, you can create custom dashboards that contain only the tabs and cards that are relevant to monitoring your implementation of StorageGRID.

The system-provided dashboard tabs contain cards with the following types of information:

Tab on system-provided dashboard	Contains
Overview	General information about the grid, such as active alerts, space usage, and total objects in the grid.
Performance	Space usage, storage used over time, S3 operations, request duration, error rate.
Storage	Tenant quota usage and logical space usage. Forecasts of space usage for user data and metadata.
ILM	Information lifecycle management queue and evaluation rate.

Tab on system-provided dashboard	Contains
Nodes	CPU, data, and memory usage by node. S3 operations by node. Node to site distribution.

Some of the cards can be maximized for easier viewing. Select the maximize icon  in the upper right corner of the card. To close a maximized card, select the minimize icon  or select **Close**.

Manage dashboards

If you have Root access (see [Admin group permissions](#)), you can perform the following management tasks for dashboards:

- Create a custom dashboard from scratch. You can use custom dashboards to control which StorageGRID information is displayed and how that information is organized.
- Clone a dashboard to create custom dashboards.
- Set an active dashboard for a user. The active dashboard can be the system-provided dashboard or a custom dashboard.
- Set a default dashboard, which is what all users see unless they activate their own dashboard.
- Edit a dashboard name.
- Edit a dashboard to add or remove tabs and cards. You can have a minimum of 1 and a maximum of 20 tabs.
- Remove a dashboard.



If you have any other permission besides Root access, you can only set an active dashboard.

To manage dashboards, select **Actions > Manage dashboards**.



Configure dashboards

To create a new dashboard by cloning the active dashboard, select **Actions > Clone active dashboard**.

To edit or clone an existing dashboard, select **Actions > Manage dashboards**.



The system-provided dashboard can't be edited or removed.

When configuring a dashboard, you can:

- Add or remove tabs
- Rename tabs and give new tabs unique names
- Add, remove, or rearrange (drag) cards for each tab

- Select the size for individual cards by selecting **S**, **M**, **L** or **XL** at the top of the card

Configure dashboard

Overview Performance Storage ILM Nodes + Add tab

Tab name
Overview

Select cards

S M L

Health status

License
1

License

M L XL

Data space usage breakdown

3.50 MB (0%) of 3.09 TB used overall

Site name	Data storage usage	Used space	Total space
Data Center 1	0%	1.79 MB	1.24 TB
Data Center 2	0%	921.11 KB	926.62 GB
Data Center 3	0%	790.21 KB	926.62 GB

View the Nodes page

View the Nodes page

When you need more detailed information about your StorageGRID system than the dashboard provides, you can use the Nodes page to view metrics for the entire grid, each site in the grid, and each node at a site.

The Nodes table lists summary information for the entire grid, each site, and each node. If a node is disconnected or has an active alert, an icon appears next to the node name. If the node is connected and has no active alerts, no icon is shown.



When a node is not connected to the grid, such as during upgrade or a disconnected state, certain metrics might be unavailable or excluded from site and grid totals. After a node reconnects to the grid, wait several minutes for the values to stabilize.



To change units for the storage values displayed in the Grid Manager, select the user drop-down in the upper right of the Grid Manager, then select **User preferences**.



The screenshots shown are examples. Your results might vary depending on your StorageGRID version.

Nodes

View the list and status of sites and grid nodes.

Search...		Total node count: 12		
Name	Type	Object data used	Object metadata used	CPU usage
StorageGRID Webscale Deployment	Grid	0%	0%	—
^ DC1	Site	0%	0%	—
 DC1-ADM1	Primary Admin Node	—	—	6%
 DC1-ARC1	Archive Node	—	—	1%
 DC1-G1	Gateway Node	—	—	3%
DC1-S1	Storage Node	0%	0%	6%
DC1-S2	Storage Node	0%	0%	8%
DC1-S3	Storage Node	0%	0%	4%

Connection state icons

If a node is disconnected from the grid, either of the following icons appears next to the node name.

Icon	Description	Action required
	Not connected - Unknown For an unknown reason, a node is disconnected or services on the node are unexpectedly down. For example, a service on the node might be stopped, or the node might have lost its network connection because of a power failure or unexpected outage. The Unable to communicate with node alert might also be triggered. Other alerts might also be active.	Requires immediate attention. Select each alert and follow the recommended actions. For example, you might need to restart a service that has stopped or restart the host for the node. Note: A node might appear as Unknown during managed shutdown operations. You can ignore the Unknown state in these cases.

Icon	Description	Action required
	Not connected - Administratively down For an expected reason, node is not connected to grid. For example, the node, or services on the node, has been gracefully shut down, the node is rebooting, or the software is being upgraded. One or more alerts might also be active. Based on the underlying issue, these nodes often go back online with no intervention.	Determine if any alerts are affecting this node. If one or more alerts are active, Select each alert and follow the recommended actions.

If a node is disconnected from the grid, it might have an underlying alert, but only the "Not connected" icon appears. To see the active alerts for a node, select the node.

Alert icons

If there is an active alert for a node, one of the following icons appears next to the node name:

 **Critical:** An abnormal condition exists that has stopped the normal operations of a StorageGRID node or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved.

 **Major:** An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service.

 **Minor:** The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that don't clear on their own to ensure they don't result in a more serious problem.

View details for a system, site, or node

To filter the information shown in the Nodes table, enter a search string in the **Search** field. You can search by system name, display name, or type (for example, enter **gat** to quickly locate all Gateway Nodes).

To view the information for the grid, site, or node:

- Select the grid name to see an aggregate summary of the statistics for your entire StorageGRID system.
- Select a specific data center site to see an aggregate summary of the statistics for all nodes at that site.
- Select a specific node to view detailed information for that node.

View the Overview tab

The Overview tab provides basic information about each node. It also shows any alerts currently affecting the node.

The Overview tab is shown for all nodes.

Node Information

The Node Information section of the Overview tab lists basic information about the node.

NYC-ADM1 (Primary Admin Node)

Overview

Hardware

Network

Storage

Load balancer

Tasks

Node information

Display name:	NYC-ADM1
System name:	DC1-ADM1
Type:	Primary Admin Node
ID:	3adb1aa8-9c7a-4901-8074-47054aa06ae6
Connection state:	 Connected
Software version:	11.7.0
IP addresses:	10.96.105.85 - eth0 (Grid Network)

Show additional IP addresses 

The overview information for a node includes the following:

- **Display name** (shown only if the node has been renamed): The current display name for the node. Use the [Rename grid, sites, and nodes](#) procedure to update this value.
- **System name**: The name you entered for the node during installation. System names are used for internal StorageGRID operations and can't be changed.
- **Type**: The type of node — Admin Node, primary Admin Node, Storage Node, or Gateway Node.
- **ID**: The unique identifier for the node, which is also referred to as the UUID.
- **Connection state**: One of three states. The icon for the most severe state is shown.

- **Unknown** : For an unknown reason, the node is not connected to the grid, or one or more services are unexpectedly down. For example, the network connection between nodes has been lost, the power is down, or a service is down. The **Unable to communicate with node** alert might also be triggered. Other alerts might be active as well. This situation requires immediate attention.



A node might appear as Unknown during managed shutdown operations. You can ignore the Unknown state in these cases.

- **Administratively down** : The node is not connected to the grid for an expected reason. For

example, the node, or services on the node, has been gracefully shut down, the node is rebooting, or the software is being upgraded. One or more alerts might also be active.

◦ **Connected** : The node is connected to the grid.

• **Storage used:** For Storage Nodes only.

◦ **Object data:** The percentage of the total usable space for object data that has been used on the Storage Node.

◦ **Object metadata:** The percentage of the total allowed space for object metadata that has been used on the Storage Node.

• **Software version:** The version of StorageGRID that is installed on the node.

• **HA groups:** For Admin Node and Gateway Nodes only. Shown if a network interface on the node is included in a high availability group and whether that interface is the Primary interface.

• **IP addresses:** The node's IP addresses. Click **Show additional IP addresses** to view the node's IPv4 and IPv6 addresses and interface mappings.

Alerts

The Alerts section of the Overview tab lists any [alerts currently affecting this node that have not been silenced](#). Select the alert name to view additional details and recommended actions.

Alerts			
Alert name 	Severity  	Time triggered 	Current values
Low installed node memory  The amount of installed memory on a node is low.	 Critical	11 hours ago 	Total RAM size: 8.37 GB

Alerts are also included for [node connection states](#).


View the Hardware tab

The Hardware tab displays CPU utilization and memory usage for each node, and additional hardware information about appliances.




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

The Hardware tab is shown for all nodes.



To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.

To see details for CPU utilization and memory usage, position your cursor over each graph.



If the node is an appliance node, this tab also includes a section with more information about the appliance hardware.

View information about appliance Storage Nodes

The Nodes page lists information about service health and all computational, disk device, and network resources for each appliance Storage Node. You can also see memory, storage hardware, controller firmware version, network resources, network interfaces, network addresses, and receive and transmit data.

Steps

1. From the Nodes page, select an appliance Storage Node.
2. Select **Overview**.

The Node information section of the Overview tab displays summary information for the node, such as the node's name, type, ID, and connection state. The list of IP addresses includes the name of the interface for each address, as follows:

- **eth**: The Grid Network, Admin Network, or Client Network.
- **hic**: One of the physical 10, 25, or 100 GbE ports on the appliance. These ports can be bonded together and connected to the StorageGRID Grid Network (eth0) and Client Network (eth2).
- **mtc**: One of the physical 1 GbE ports on the appliance. One or more mtc interfaces are bonded to form the StorageGRID Admin Network interface (eth1). You can leave other mtc interfaces available for temporary local connectivity for a technician in the data center.

DC2-SGA-010-096-106-021 (Storage Node) [🔗](#)

✕

OverviewHardwareNetworkStorageObjectsILMTasks

Node information [?](#)

Name: DC2-SGA-010-096-106-021

Type: Storage Node

ID: f0890e03-4c72-401f-ae92-245511a38e51

Connection state: ✔ Connected

Storage used:

Object data7%

Object metadata5%

Software version: 11.6.0 (build 20210915.1941.afce2d9)

IP addresses: 10.96.106.21 - eth0 (Grid Network)

[Hide additional IP addresses ^](#)

Interface ⌵	IP address ⌵
eth0 (Grid Network)	10.96.106.21
eth0 (Grid Network)	fe80::2a0:98ff:fe64:6582
hic2	10.96.106.21
hic4	10.96.106.21
mtc2	169.254.0.1

Alerts

Alert name [⌵](#)

Severity [?](#) [⌵](#)

Time triggered [⌵](#)

Current values

[ILM placement unachievable \[🔗\]\(#\)](#)


! Major

2 hours ago [?](#)

A placement instruction in an ILM rule cannot be achieved for certain objects.

The Alerts section of the Overview tab displays any active alerts for the node.

3. Select **Hardware** to see more information about the appliance.
 - a. View the CPU Utilization and Memory graphs to determine the percentages of CPU and memory usage over time. To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.



- b. Scroll down to view the table of components for the appliance. This table contains information such as the model name of the appliance; controller names, serial numbers, and IP addresses; and the status of each component.



Some fields, such as Compute controller BMC IP and Compute hardware, appear only for appliances with that feature.

Components for the storage shelves, and expansion shelves if they are part of the installation, appear in a separate table below the appliance table.

StorageGRID Appliance

Appliance model:	SG6060	
Storage controller name:	StorageGRID-Lab79-SG6060-7-134	
Storage controller A management IP:	10.2	
Storage controller B management IP:	10.2	
Storage controller WWID:	6d039ea0000173e50000000065b7b761	
Storage appliance chassis serial number:	721924500068	
Storage controller firmware version:	08.53.00.09	
Storage controller SANtricity OS version:	11.50.3R2	
Storage controller NVSRAM version:	N280X-853834-DG1	
Storage hardware:	Nominal	
Storage controller failed drive count:	0	
Storage controller A:	Nominal	
Storage controller B:	Nominal	
Storage controller power supply A:	Nominal	
Storage controller power supply B:	Nominal	
Storage data drive type:	NL-SAS HDD	
Storage data drive size:	4.00 TB	
Storage RAID mode:	DDP16	
Storage connectivity:	Nominal	
Overall power supply:	Degraded	
Compute controller BMC IP:	10.2	
Compute controller serial number:	721917500060	
Compute hardware:	Needs Attention	
Compute controller CPU temperature:	Nominal	
Compute controller chassis temperature:	Nominal	
Compute controller power supply A:	Failed	
Compute controller power supply B:	Nominal	

Storage shelves

Shelf chassis serial number	Shelf ID	Shelf status	IOM status	Power supply status	Drawer status	Fan status
721924500068	99	Nominal	N/A	Nominal	Nominal	Nominal

Field in the Appliance table	Description
Appliance model	The model number for this StorageGRID appliance shown in SANtricity OS.
Storage controller name	The name for this StorageGRID appliance shown in SANtricity OS.
Storage controller A management IP	IP address for management port 1 on storage controller A. You use this IP to access SANtricity OS to troubleshoot storage issues.
Storage controller B management IP	IP address for management port 1 on storage controller B. You use this IP to access SANtricity OS to troubleshoot storage issues. Some appliance models don't have a storage controller B.
Storage controller WWID	The worldwide identifier of the storage controller shown in SANtricity OS.

Field in the Appliance table	Description
Storage appliance chassis serial number	The chassis serial number of the appliance.
Storage controller firmware version	The version of the firmware on the storage controller for this appliance.
Storage controller SANtricity OS version	The SANtricity OS version of storage controller A.
Storage controller NVSRAM version	NVSRAM version of the storage controller as reported by SANtricity System Manager. For the SG6060 and SG6160, if there is an NVSRAM version mismatch between the two controllers, the version of controller A displays. If controller A is not installed or operational, the version of controller B displays.
Storage hardware	The overall status of the storage controller hardware. If SANtricity System Manager reports a status of Needs Attention for the storage hardware, the StorageGRID system also reports this value. If the status is "needs attention," first check the storage controller using SANtricity OS. Then, ensure that no other alerts exist that apply to the compute controller.
Storage controller failed drive count	The number of drives that aren't optimal.
Storage controller A	The status of storage controller A.
Storage controller B	The status of storage controller B. Some appliance models don't have a storage controller B.
Storage controller power supply A	The status of power supply A for the storage controller.
Storage controller power supply B	The status of power supply B for the storage controller.
Storage data drive type	The type of drives in the appliance, such as HDD (hard drive) or SSD (solid state drive).

Field in the Appliance table	Description
Storage data drive size	The effective size of one data drive. For the SG6160, the size of the cache drive also displays. Note: For nodes with expansion shelves, use the Data drive size for each shelf instead. Effective drive size might differ by shelf.
Storage RAID mode	The RAID mode configured for the appliance.
Storage connectivity	The storage connectivity state.
Overall power supply	The status of all power supplies for the appliance.
Compute controller BMC IP	The IP address of the baseboard management controller (BMC) port in the compute controller. You use this IP to connect to the BMC interface to monitor and diagnose the appliance hardware. This field is not displayed for appliance models that don't contain a BMC.
Compute controller serial number	The serial number of the compute controller.
Compute hardware	The status of the compute controller hardware. This field is not displayed for appliance models that don't have separate compute hardware and storage hardware.
Compute controller CPU temperature	The temperature status of the compute controller's CPU.
Compute controller chassis temperature	The temperature status of the compute controller.

Column in the Storage shelves table	Description
Shelf chassis serial number	The serial number for the storage shelf chassis.
Shelf ID	The numeric identifier for the storage shelf. • 99: Storage controller shelf • 0: First expansion shelf • 1: Second expansion shelf Note: Expansion shelves apply only to the SG6060 and SG6160.

Column in the Storage shelves table	Description
Shelf status	The overall status of the storage shelf.
IOM status	The status of the input/output modules (IOMs) in any expansion shelves. N/A if this is not an expansion shelf.
Power supply status	The overall status of the power supplies for the storage shelf.
Drawer status	The status of the drawers in the storage shelf. N/A if the shelf does not contain drawers.
Fan status	The overall status of the cooling fans in the storage shelf.
Drive slots	The total number of drive slots in the storage shelf.
Data drives	The number of drives in the storage shelf that are used for data storage.
Data drive size	The effective size of one data drive in the storage shelf.
Cache drives	The number of drives in the storage shelf that are used as cache.
Cache drive size	The size of the smallest cache drive in the storage shelf. Normally, cache drives are all the same size.
Configuration status	The configuration status of the storage shelf.

c. Confirm that all statuses are "Nominal."

If a status is not "Nominal," review any current alerts. You can also use SANtricity System Manager to learn more about some of these hardware values. See the instructions for installing and maintaining your appliance.

4. Select **Network** to view information for each network.

The Network Traffic graph provides a summary of overall network traffic.



a. Review the Network Interfaces section.

Network interfaces						
Name	Hardware address	Speed	Duplex	Auto-negotiation	Link status	
eth0	00:50:56:A7:66:75	10 Gigabit	Full	Off	Up	

Use the following table with the values in the **Speed** column in the Network Interfaces table to determine whether the 10/25-GbE network ports on the appliance were configured to use active/backup mode or LACP mode.



The values shown in the table assume all four links are used.

Link mode	Bond mode	Individual HIC link speed (hic1, hic2, hic3, hic4)	Expected Grid/Client Network speed (eth0,eth2)
Aggregate	LACP	25	100
Fixed	LACP	25	50
Fixed	Active/Backup	25	25
Aggregate	LACP	10	40
Fixed	LACP	10	20
Fixed	Active/Backup	10	10

See [Configure network links](#) for more information about configuring the 10/25-GbE ports.

b. Review the Network Communication section.

The Receive and Transmit tables show how many bytes and packets have been received and sent across each network as well as other receive and transmit metrics.

Network communication



Receive

Interface	Data	Packets	Errors	Dropped	Frame overruns	Frames
eth0	2.89 GB	19,421,503	0	24,032	0	0

Transmit

Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	3.64 GB	18,494,381	0	0	0	0

5. Select **Storage** to view graphs that show the percentages of storage used over time for object data and object metadata, as well as information about disk devices, volumes, and object stores.



- a. Scroll down to view the amounts of available storage for each volume and object store.

The Worldwide Name for each disk matches the volume world-wide identifier (WWID) that appears when you view standard volume properties in SANtricity OS (the management software connected to the appliance's storage controller).

To help you interpret disk read and write statistics related to volume mount points, the first portion of the name shown in the **Name** column of the Disk Devices table (that is, *sdc*, *sdd*, *sde*, and so on) matches the value shown in the **Device** column of the Volumes table.

Disk devices					
Name	World Wide Name	I/O load	Read rate	Write rate	
croot(8:1,sda1)	N/A	0.04%	0 bytes/s	3 KB/s	
cvloc(8:2,sda2)	N/A	0.67%	0 bytes/s	50 KB/s	
sdc(8:16,sdb)	N/A	0.03%	0 bytes/s	4 KB/s	
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s	
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s	

Volumes					
Mount point	Device	Status	Size	Available	Write cache status
/	croot	Online	21.00 GB	14.75 GB	Unknown
/var/local	cvloc	Online	85.86 GB	84.05 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

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

View information about appliance Admin Nodes and Gateway Nodes

The Nodes page lists information about service health and all computational, disk device, and network resources for each services appliance that is used as an Admin Node or a Gateway Node. You can also see memory, storage hardware, network resources, network interfaces, network addresses, and receive and

transmit data.

Steps

- 1. From the Nodes page, select an appliance Admin Node or an appliance Gateway Node.
- 2. Select **Overview**.

The Node information section of the Overview tab displays summary information for the node, such as the node’s name, type, ID, and connection state. The list of IP addresses includes the name of the interface for each address, as follows:

- **adllb** and **adlli**: Shown if active/backup bonding is used for the Admin Network interface
- **eth**: The Grid Network, Admin Network, or Client Network.
- **hic**: One of the physical 10, 25, or 100 GbE ports on the appliance. These ports can be bonded together and connected to the StorageGRID Grid Network (eth0) and Client Network (eth2).
- **mtc**: One of the physical 1-GbE ports on the appliance. One or more mtc interfaces are bonded to form the Admin Network interface (eth1). You can leave other mtc interfaces available for temporary local connectivity for a technician in the data center.

10-224-6-199-ADM1 (Primary Admin Node)

Overview

Hardware

Network

Storage

Load balancer

Tasks

SANtricity System Manager

Node information

Name:

10-224-6-199-ADM1

Type:

Primary Admin Node

ID:

6fdc1890-ca0a-4493-acdd-72ed317d95fb

Connection state:

Connected

Software version:

11.6.0 (build 20210928.1321.6687ee3)

IP addresses:

172.16.6.199 - eth0 (Grid Network)

10.224.6.199 - eth1 (Admin Network)


47.47.7.241 - eth2 (Client Network)

Hide additional IP addresses

Interface 	IP address
eth2 (Client Network)	47.47.7.241
eth2 (Client Network)	fd20:332:332:0:e42:a1ff:fe86:b5b0
eth2 (Client Network)	fe80::e42:a1ff:fe86:b5b0
hic1	47.47.7.241
hic2	47.47.7.241
hic3	47.47.7.241

The Alerts section of the Overview tab displays any active alerts for the node.

3. Select **Hardware** to see more information about the appliance.
 - a. View the CPU Utilization and Memory graphs to determine the percentages of CPU and memory usage over time. To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.



- b. Scroll down to view the table of components for the appliance. This table contains information such as the model name, serial number, controller firmware version, and the status of each component.

StorageGRID Appliance		
Appliance model: ?	SG100	
Storage controller failed drive count: ?	0	
Storage data drive type: ?	SSD	
Storage data drive size: ?	960.20 GB	
Storage RAID mode: ?	RAID1 [healthy]	
Storage connectivity: ?	Nominal	
Overall power supply: ?	Nominal	
Compute controller BMC IP: ?	10.60.8.38	
Compute controller serial number: ?	372038000093	
Compute hardware: ?	Nominal	
Compute controller CPU temperature: ?	Nominal	
Compute controller chassis temperature: ?	Nominal	
Compute controller power supply A: ?	Nominal	
Compute controller power supply B: ?	Nominal	

Field in the Appliance table	Description
Appliance model	The model number for this StorageGRID appliance.
Storage controller failed drive count	The number of drives that aren't optimal.
Storage data drive type	The type of drives in the appliance, such as HDD (hard drive) or SSD (solid state drive).
Storage data drive size	The effective size of one data drive.
Storage RAID mode	The RAID mode for the appliance.
Overall power supply	The status of all power supplies in the appliance.
Compute controller BMC IP	The IP address of the baseboard management controller (BMC) port in the compute controller. You can use this IP to connect to the BMC interface to monitor and diagnose the appliance hardware. This field is not displayed for appliance models that don't contain a BMC.
Compute controller serial number	The serial number of the compute controller.
Compute hardware	The status of the compute controller hardware.
Compute controller CPU temperature	The temperature status of the compute controller's CPU.
Compute controller chassis temperature	The temperature status of the compute controller.

c. Confirm that all statuses are "Nominal."

If a status is not "Nominal," review any current alerts.

4. Select **Network** to view information for each network.

The Network Traffic graph provides a summary of overall network traffic.



a. Review the Network Interfaces section.

Network interfaces						
Name ?	Hardware address ?	Speed ?	Duplex ?	Auto-negotiation ?	Link status ?	
eth0	0C:42:A1:86:B5:B0	100 Gigabit	Full	Off	Up	
eth1	B4:A9:FC:71:68:36	Gigabit	Full	Off	Up	
eth2	0C:42:A1:86:B5:B0	100 Gigabit	Full	Off	Up	
hic1	0C:42:A1:86:B5:B0	25 Gigabit	Full	On	Up	
hic2	0C:42:A1:86:B5:B0	25 Gigabit	Full	On	Up	
hic3	0C:42:A1:86:B5:B0	25 Gigabit	Full	On	Up	
hic4	0C:42:A1:86:B5:B0	25 Gigabit	Full	On	Up	
mtc1	B4:A9:FC:71:68:36	Gigabit	Full	On	Up	
mtc2	B4:A9:FC:71:68:35	Gigabit	Full	On	Up	

Use the following table with the values in the **Speed** column in the Network Interfaces table to determine whether the four 40/100-GbE network ports on the appliance were configured to use active/backup mode or LACP mode.



The values shown in the table assume all four links are used.

Link mode	Bond mode	Individual HIC link speed (hic1, hic2, hic3, hic4)	Expected Grid/Client Network speed (eth0, eth2)
Aggregate	LACP	100	400
Fixed	LACP	100	200
Fixed	Active/Backup	100	100
Aggregate	LACP	40	160
Fixed	LACP	40	80
Fixed	Active/Backup	40	40

b. Review the Network Communication section.

The Receive and Transmit tables show how many bytes and packets have been received and sent across each network as well as other receive and transmission metrics.

Network communication							
Receive							
Interface ?	Data ?	Packets ?	Errors ?	Dropped ?	Frame overruns ?	Frames ?	
eth0	2.89 GB	19,421,503	0	24,032	0	0	
Transmit							
Interface ?	Data ?	Packets ?	Errors ?	Dropped ?	Collisions ?	Carrier ?	
eth0	3.64 GB	18,494,381	0	0	0	0	

5. Select **Storage** to view information about the disk devices and volumes on the services appliance.

[Overview](#)[Hardware](#)[Network](#)[Storage](#)[Load balancer](#)[Tasks](#)

Disk devices

Name ? ⬆ ⬇	World Wide Name ? ⬆ ⬇	I/O load ? ⬆ ⬇	Read rate ? ⬆ ⬇	Write rate ? ⬆ ⬇
croot(8:1,sda1)	N/A	0.02%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.03%	0 bytes/s	6 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	84.63 GB 	Unknown

View the Network tab


The Network tab displays a graph showing the network traffic received and sent across all of the network interfaces on the node, site, or grid.

The Network tab is shown for all nodes, each site, and the entire grid.

To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.

For nodes, the Network interfaces table provides information about each node's physical network ports. The Network communications table provides details about each node's receive and transmit operations and any driver reported fault counters.

DC1-S2 (Storage Node)

[Overview](#)[Hardware](#)[Network](#)[Storage](#)[Objects](#)[ILM](#)[Tasks](#)[1 hour](#)[1 day](#)[1 week](#)[1 month](#)[Custom](#)

Network interfaces

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

Network communication

Receive

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

Transmit

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

Related information

[Monitor network connections and performance](#)

View the Storage tab

The Storage tab summarizes storage availability and other storage metrics.


The Storage tab is shown for all nodes, each site, and the entire grid.

Storage used graphs

For Storage Nodes, each site, and the entire grid, the Storage tab includes graphs showing how much storage has been used by object data and object metadata over time.



When a node is not connected to the grid, such as during upgrade or a disconnected state, certain metrics might be unavailable or excluded from site and grid totals. After a node reconnects to the grid, wait several minutes for the values to stabilize.



Disk devices, Volumes, and Object stores tables

For all nodes, the Storage tab contains details for the disk devices and volumes on the node. For Storage Nodes, the Object Stores table provides information about each storage volume.

Disk devices

Name  	World Wide Name  	I/O load  	Read rate  	Write rate  
croot(8:1,sda1)	N/A	0.04%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.67%	0 bytes/s	50 KB/s
sdc(8:16,sdb)	N/A	0.03%	0 bytes/s	4 KB/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s

Volumes

Mount point  	Device  	Status  	Size  	Available  	Write cache status  
/	croot	Online	21.00 GB	14.75 GB 	Unknown
/var/local	cvloc	Online	85.86 GB	84.05 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

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

Related information

[Monitor storage capacity](#)

View the Objects tab

The Objects tab provides information about [S3 ingest and retrieve rates](#).

The Objects tab is shown for each Storage Node, each site, and the entire grid. For Storage Nodes, the Objects tab also provides object counts and information about metadata queries and background verification.

DC1-S1 (Storage Node) [🔗](#)

[Overview](#)[Hardware](#)[Network](#)[Storage](#)[Objects](#)[ILM](#)[Tasks](#)[1 hour](#)[1 day](#)[1 week](#)[1 month](#)[Custom](#)

Object counts

Total objects: [?](#) 1,295

Lost objects: [?](#) 0

S3 buckets and Swift containers: [?](#) 161

Metadata store queries

Average latency: [?](#) 10.00 milliseconds

Queries - successful: [?](#) 14,587

Queries - failed (timed out): [?](#) 0

Queries - failed (consistency level unmet): [?](#) 0

Verification

Status: [?](#) No errors

Percent complete: [?](#) 47.14%

Average stat time: [?](#) 0.00 microseconds

Objects verified: [?](#) 0

Object verification rate: [?](#) 0.00 objects / second

Data verified: [?](#) 0 bytes

Data verification rate: [?](#) 0.00 bytes / second

Missing objects: [?](#) 0

Corrupt objects: [?](#) 0

Corrupt objects unidentified: [?](#) 0


Quarantined objects: [?](#) 0

View the ILM tab

The ILM tab provides information about information lifecycle management (ILM) operations.

The ILM tab is shown for each Storage Node, each site, and the entire grid. For each site and the grid, the ILM tab shows a graph of the ILM queue over time. For the grid, this tab also provides the estimated time to complete a full ILM scan of all objects.

For Storage Nodes, the ILM tab provides details about ILM evaluation and background verification for erasure-coded objects.



Related information

- [Monitor information lifecycle management](#)
- [Administer StorageGRID](#)

Use the Tasks tab

The Tasks tab is shown for all nodes. You can use this tab to rename or reboot a node or to put an appliance node into maintenance mode.

For the complete requirements and instructions for each option on this tab, see the following:

- [Rename grid, sites, and nodes](#)
- [Reboot grid node](#)
- [Place appliance into maintenance mode](#)

View the Load balancer tab

The Load Balancer tab includes performance and diagnostic graphs related to the operation of the Load Balancer service.

The Load Balancer tab is shown for Admin Nodes and Gateway Nodes, each site, and the entire grid. For each site, the Load Balancer tab provides an aggregate summary of the statistics for all nodes at that site. For the entire grid, the Load Balancer tab provides an aggregate summary of the statistics for all sites.

If there is no I/O being run through the Load Balancer service, or there is no load balancer configured, the graphs display "No data."



Request traffic

This graph provides a 3-minute moving average of the throughput of data transmitted between load balancer endpoints and the clients making the requests, in bits per second.



This value is updated at the completion of each request. As a result, this value might differ from the real-time throughput at low request rates or for very long-lived requests. You can look at the Network tab to get a more realistic view of the current network behavior.

Incoming request rate

This graph provides a 3-minute moving average of the number of new requests per second, broken down by request type (GET, PUT, HEAD, and DELETE). This value is updated when the headers of a new request have been validated.

Average request duration (non-error)

This graph provides a 3-minute moving average of request durations, broken down by request type (GET, PUT, HEAD, and DELETE). Each request duration starts when a request header is parsed by the Load Balancer service and ends when the complete response body is returned to the client.

Error response rate

This graph provides a 3-minute moving average of the number of error responses returned to clients per second, broken down by the error response code.



Related information

- [Monitor load balancing operations](#)
- [Administer StorageGRID](#)

View the Platform services tab

The Platform services tab provides information about any S3 platform service operations at a site.

The Platform services tab is shown for each site. This tab provides information about S3 platform services, such as CloudMirror replication and the search integration service. Graphs on this tab display metrics such as the number of pending requests, request completion rate, and request failure rate.

[Network](#)[Storage](#)[Objects](#)[ILM](#)[Platform services](#)[Load balancer](#)[1 hour](#)[1 day](#)[1 week](#)[1 month](#)[Custom](#)

For more information about S3 platform services, including troubleshooting details, see the [instructions for administering StorageGRID](#).

View the Manage drives tab

The Manage drives tab enables you to access details and perform troubleshooting and maintenance tasks on drives in the appliances that support this feature.

Using the Manage drives tab, you can do the following:

- View a layout of the data storage drives in the appliance

- View a table that lists each drive location, type, status, firmware version, and serial number
- Perform troubleshooting and maintenance functions on each drive

To access the Manage drives tab, you must have the [Storage appliance administrator or Root access permission](#).

For information about using the Manage drives tab, see [Use the Manage drives tab](#).

View the SANtricity System Manager tab (E-Series only)

The SANtricity System Manager tab enables you to access SANtricity System Manager without having to configure or connect the management port of the storage appliance. You can use this tab to review hardware diagnostic and environmental information as well as issues related to the drives.



Accessing SANtricity System Manager from the Grid Manager is generally meant only to monitor appliance hardware and configure E-Series AutoSupport. Many features and operations within SANtricity System Manager such as upgrading firmware don't apply to monitoring your StorageGRID appliance. To avoid issues, always follow the hardware maintenance instructions for your appliance. To upgrade SANtricity firmware, see the [Maintenance configuration procedures](#) for your storage appliance.



The SANtricity System Manager tab is shown only for storage appliance nodes using E-Series hardware.

Using SANtricity System Manager, you can do the following:

- View performance data such as storage array level performance, I/O latency, storage controller CPU utilization, and throughput.
- Check hardware component status.
- Perform support functions including viewing diagnostic data, and configuring E-Series AutoSupport.



To use SANtricity System Manager to configure a proxy for E-Series AutoSupport, see [Send E-Series AutoSupport packages through StorageGRID](#).

To access SANtricity System Manager through Grid Manager, you must have the [Storage appliance administrator or Root access permission](#).




You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager.

The tab displays the home page of SANtricity System Manager.

Use SANtricity System Manager to monitor and manage the hardware components in this storage appliance. From SANtricity System Manager, you can review hardware diagnostic and environmental information as well as issues related to the drives.

Note: Many features and operations within SANtricity Storage Manager do not apply to your StorageGRID appliance. To avoid issues, always follow the hardware installation and maintenance instructions for your appliance model.

Open [SANtricity System Manager](#) in a new browser tab.



You can use the SANtricity System Manager link to open the SANtricity System Manager in a new browser window for easier viewing.

To see details for storage array level performance and capacity usage, position your cursor over each graph.

For more details on viewing the information accessible from the SANtricity System Manager tab, see [NetApp E-Series and SANtricity documentation](#).

Information to monitor regularly

What and when to monitor

Even though the StorageGRID system can continue to operate when errors occur or parts of the grid are unavailable, you should monitor and address potential issues before they affect the grid’s efficiency or availability.

Before you begin

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

About monitoring tasks

A busy system generates large amounts of information. The following list provides guidance about the most important information to monitor on an ongoing basis.

What to monitor	Frequency
System health status	Daily
Rate at which Storage Node object and metadata capacity is being consumed	Weekly
Information lifecycle management operations	Weekly
Networking and system resources	Weekly
Tenant activity	Weekly
S3 client operations	Weekly
Load balancing operations	After the initial configuration and after any configuration changes
Grid federation connections	Weekly


Monitor system health

Monitor the overall health of your StorageGRID system on a daily basis.

About this task

The StorageGRID system can continue to operate when parts of the grid are unavailable. Potential issues indicated by alerts aren’t necessarily issues with system operations. Investigate issues summarized on the Health status card of the Grid Manager Dashboard.

To be notified of alerts as soon as they are triggered, you can [set up email notifications for alerts](#) or [configure SNMP traps](#).



When issues exist, links appear that allow you to view additional details:

Link	Appears when...
Grid details	Any nodes are disconnected (connection state Unknown or Administratively Down).
Current alerts (Critical, Major, Minor)	Alerts are currently active .
Recently resolved alerts	Alerts triggered in the past week are now resolved .
License	There is an issue with the software license for this StorageGRID system. You can update license information as needed .

Monitor node connection states

If one or more nodes are disconnected from the grid, critical StorageGRID operations might be affected. Monitor node connection states and address any issues promptly.


Icon	Description	Action required
	Not connected - Unknown For an unknown reason, a node is disconnected or services on the node are unexpectedly down. For example, a service on the node might be stopped, or the node might have lost its network connection because of a power failure or unexpected outage. The Unable to communicate with node alert might also be triggered. Other alerts might also be active.	Requires immediate attention. Select each alert and follow the recommended actions. For example, you might need to restart a service that has stopped or restart the host for the node. Note: A node might appear as Unknown during managed shutdown operations. You can ignore the Unknown state in these cases.
	Not connected - Administratively down For an expected reason, node is not connected to grid. For example, the node, or services on the node, has been gracefully shut down, the node is rebooting, or the software is being upgraded. One or more alerts might also be active. Based on the underlying issue, these nodes often go back online with no intervention.	Determine if any alerts are affecting this node. If one or more alerts are active, select each alert and follow the recommended actions.
	Connected The node is connected to the grid.	No action required.

View current and resolved alerts

Current alerts: When an alert is triggered, an alert icon is displayed on the dashboard. An alert icon is also displayed for the node on the Nodes page. If [alert email notifications are configured](#), an email notification will also be sent, unless the alert has been silenced.

Resolved alerts: You can search and view a history of alerts that have been resolved.

Optionally, you have watched the video: [Video: Alerts overview](#)



The following table describes the information shown in the Grid Manager for current and resolved alerts.

Column header	Description
Name or title	The name of the alert and its description.
Severity	The severity of the alert. For current alerts, if multiple alerts are grouped the title row shows how many instances of that alert are occurring at each severity.  Critical: An abnormal condition exists that has stopped the normal operations of a StorageGRID node or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved.  Major: An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service.  Minor: The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that don't clear on their own to ensure they don't result in a more serious problem.
Time triggered	Current alerts: The date and time the alert was triggered in your local time and in UTC. If multiple alerts are grouped, the title row shows times for the most recent instance of the alert (<i>newest</i>) and the oldest instance of the alert (<i>oldest</i>). Resolved alerts: How long ago the alert was triggered.
Site/Node	The name of the site and node where the alert is occurring or has occurred.
Status	Whether the alert is active, silenced, or resolved. If multiple alerts are grouped and All alerts is selected in the drop-down, the title row shows how many instances of that alert are active and how many instances have been silenced.
Time resolved (resolved alerts only)	How long ago the alert was resolved.

Column header	Description
Current values or <i>data values</i>	The value of the metric that caused the alert to be triggered. For some alerts, additional values are shown to help you understand and investigate the alert. For example, the values shown for a Low object data storage alert include the percentage of disk space used, the total amount of disk space, and the amount of disk space used. Note: If multiple current alerts are grouped, current values aren't shown in the title row.
Triggered values (resolved alerts only)	The value of the metric that caused the alert to be triggered. For some alerts, additional values are shown to help you understand and investigate the alert. For example, the values shown for a Low object data storage alert include the percentage of disk space used, the total amount of disk space, and the amount of disk space used.

Steps

1. Select the **Current alerts** or **Resolved alerts** link to view a list of alerts in those categories. You can also view the details for an alert by selecting **Nodes > node > Overview** and then selecting the alert from the Alerts table.

By default, current alerts are shown as follows:

- The most recently triggered alerts are shown first.
- Multiple alerts of the same type are shown as a group.
- Alerts that have been silenced aren't shown.
- For a specific alert on a specific node, if the thresholds are reached for more than one severity, only the most severe alert is shown. That is, if alert thresholds are reached for the minor, major, and critical severities, only the critical alert is shown.

The Current alerts page is refreshed every two minutes.

2. To expand groups of alerts, select the down caret ▼. To collapse individual alerts in a group, select the up caret ▲, or select the group's name.
3. To display individual alerts instead of groups of alerts, clear the **Group alerts** checkbox.
4. To sort current alerts or alert groups, select the up/down arrows ⬆⬇ in each column header.
 - When **Group alerts** is selected, both the alert groups and the individual alerts within each group are sorted. For example, you might want to sort the alerts in a group by **Time triggered** to find the most recent instance of a specific alert.
 - When **Group alerts** is cleared, the entire list of alerts is sorted. For example, you might want to sort all alerts by **Node/Site** to see all alerts affecting a specific node.
5. To filter current alerts by status (**All alerts**, **Active**, or **Silenced**, use the drop-down menu at the top of the table.

See [Silence alert notifications](#).

6. To sort resolved alerts:
 - Select a time period from the **When triggered** drop-down menu.

- Select one or more severities from the **Severity** drop-down menu.
 - Select one or more default or custom alert rules from the **Alert rule** drop-down menu to filter on resolved alerts related to a specific alert rule.
 - Select one or more nodes from the **Node** drop-down menu to filter on resolved alerts related to a specific node.
7. To view details for a specific alert, select the alert. A dialog box provides details and recommended actions for the alert you selected.
 8. (Optional) For a specific alert, select silence this alert to silence the alert rule that caused this alert to be triggered.

You must have the [Manage alerts or Root access permission](#) to silence an alert rule.



Be careful when deciding to silence an alert rule. If an alert rule is silenced, you might not detect an underlying problem until it prevents a critical operation from completing.

9. To view the current conditions for the alert rule:

- a. From the alert details, select **View conditions**.

A pop-up appears, listing the Prometheus expression for each defined severity.

- b. To close the pop-up, click anywhere outside of the pop-up.

10. Optionally, select **Edit rule** to edit the alert rule that caused this alert to be triggered.

You must have the [Manage alerts or Root access permission](#) to edit an alert rule.



Be careful when deciding to edit an alert rule. If you change trigger values, you might not detect an underlying problem until it prevents a critical operation from completing.

11. To close the alert details, select **Close**.

Monitor storage capacity

Monitor the total usable space available to ensure that the StorageGRID system does not run out of storage space for objects or for object metadata.

StorageGRID stores object data and object metadata separately, and reserves a specific amount of space for a distributed Cassandra database that contains object metadata. Monitor the total amount of space consumed for objects and for object metadata, as well as trends in the amount of space consumed for each. This will enable you to plan ahead for the addition of nodes and avoid any service outages.

You can [view storage capacity information](#) for the entire grid, for each site, and for each Storage Node in your StorageGRID system.

Monitor storage capacity for the entire grid


Monitor the overall storage capacity for your grid to ensure that adequate free space remains for object data and object metadata. Understanding how storage capacity changes over time can help you plan to add Storage Nodes or storage volumes before the grid's usable storage capacity is consumed.

The Grid Manager dashboard lets you quickly assess how much storage is available for the entire grid and for each data center. The Nodes page provides more detailed values for object data and object metadata.


Steps

- 1. Assess how much storage is available for the entire grid and for each data center.
 - a. Select **Dashboard > Overview**.
 - b. Note the values on the Data space usage breakdown and the Metadata allowed space usage breakdown cards. Each card lists a percentage of storage usage, the capacity of used space, and the total space available or allowed by site.


 The summary does not include archival media.



- c. Note the chart on the Storage over time card. Use the time period drop-down to help you determine how quickly storage is consumed.



2. Use the Nodes page for additional details on how much storage has been used and how much storage remains available on the grid for object data and object metadata.
 - a. Select **NODES**.
 - b. Select **grid** > **Storage**.



- c. Position your cursor over the **Storage used - object data** and the **Storage used - object metadata** charts to see how much object storage and object metadata storage is available for the entire grid, and how much has been used over time.



The total values for a site or the grid don't include nodes that have not reported metrics for at least five minutes, such as offline nodes.

3. Plan to perform an expansion to add Storage Nodes or storage volumes before the grid's usable storage capacity is consumed.

When planning the timing of an expansion, consider how long it will take to procure and install additional storage.



If your ILM policy uses erasure coding, you might prefer to expand when existing Storage Nodes are approximately 70% full to reduce the number of nodes that must be added.


For more information about planning a storage expansion, see the [instructions for expanding StorageGRID](#).

Monitor storage capacity for each Storage Node

Monitor the total usable space for each Storage Node to ensure that the node has enough space for new object data.

About this task

Usable space is the amount of storage space available to store objects. The total usable space for a Storage Node is calculated by adding together the available space on all object stores within the node.



$$\text{Total Usable Space} = \text{Usable Space 0} + \text{Usable Space 1} + \text{Usable Space 2}$$

Steps

1. Select **NODES > Storage Node > Storage**.


The graphs and tables for the node appear.

2. 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.



3. Review the Available values in the Volumes and Object stores tables, below the graphs.



To view graphs of these values, click the chart icons  in the Available columns.

Disk devices

Name ? ⇅	World Wide Name ? ⇅	I/O load ? ⇅	Read rate ? ⇅	Write rate ? ⇅
croot(8:1,sda1)	N/A	0.04%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.67%	0 bytes/s	50 KB/s
sdc(8:16,sdb)	N/A	0.03%	0 bytes/s	4 KB/s
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	82 bytes/s
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	82 bytes/s

Volumes

Mount point ? ⇅	Device ? ⇅	Status ? ⇅	Size ? ⇅	Available ? ⇅	Write cache status ? ⇅
/	croot	Online	21.00 GB	14.75 GB 	Unknown
/var/local	cvloc	Online	85.86 GB	84.05 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

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

- Monitor the values over time to estimate the rate at which usable storage space is being consumed.
- To maintain normal system operations, add Storage Nodes, add storage volumes, or archive object data before usable space is consumed.

When planning the timing of an expansion, consider how long it will take to procure and install additional storage.



If your ILM policy uses erasure coding, you might prefer to expand when existing Storage Nodes are approximately 70% full to reduce the number of nodes that must be added.

For more information about planning a storage expansion, see the [instructions for expanding StorageGRID](#).

The [Low object data storage](#) alert is triggered when insufficient space remains for storing object data on a Storage Node.

Monitor object metadata capacity for each Storage Node

Monitor the metadata usage for each Storage Node to ensure that adequate space remains available for essential database operations. You must add new Storage Nodes at each site before object metadata exceeds 100% of the allowed metadata space.

About this task


StorageGRID maintains three copies of object metadata at each site to provide redundancy and to protect object metadata from loss. The three copies are evenly distributed across all Storage Nodes at each site using the space reserved for metadata on storage volume 0 of each Storage Node.

In some cases, the grid's object metadata capacity might be consumed faster than its object storage capacity. For example, if you typically ingest large numbers of small objects, you might need to add Storage Nodes to increase metadata capacity even though sufficient object storage capacity remains.

Some of the factors that can increase metadata usage include the size and quantity of user metadata and tags, the total number of parts in a multipart upload, and the frequency of changes to ILM storage locations.

Steps

1. Select **NODES > Storage Node > Storage**.
2. Position your cursor over the Storage used - object metadata graph to see the values for a specific time.



Used (%)

The percentage of the allowed metadata space that has been used on this Storage Node.

Prometheus metrics: `storagegrid_storage_utilization_metadata_bytes` and `storagegrid_storage_utilization_metadata_allowed_bytes`

Used

The bytes of the allowed metadata space that have been used on this Storage Node.

Prometheus metric: `storagegrid_storage_utilization_metadata_bytes`

Allowed

The space allowed for object metadata on this Storage Node. To learn how this value is determined for each Storage Node, see the [full description of Allowed metadata space](#).

Prometheus metric: `storagegrid_storage_utilization_metadata_allowed_bytes`

Actual reserved

The actual space reserved for metadata on this Storage Node. Includes the allowed space and the required space for essential metadata operations. To learn how this value is calculated for each Storage Node, see the [full description of Actual reserved space for metadata](#).

Prometheus metric will be added in a future release.



The total values for a site or the grid don't include nodes that have not reported metrics for at least five minutes, such as offline nodes.

3. If the **Used (%)** value is 70% or higher, expand your StorageGRID system by adding Storage Nodes to each site.



The **Low metadata storage** alert is triggered when the **Used (%)** value reaches certain thresholds. Undesirable results can occur if object metadata uses more than 100% of the allowed space.

When you add the new nodes, the system automatically rebalances object metadata across all Storage Nodes within the site. See the [instructions for expanding a StorageGRID system](#).

Monitor space usage forecasts


Monitor space usage forecasts for user data and metadata to estimate when you will need to [expand a grid](#).

If you notice that the rate of consumption changes over time, select a shorter range from the **Averaged over** pull-down to reflect only the most recent ingest patterns. If you notice seasonal patterns, select a longer range.

If you have a new StorageGRID installation, allow data and metadata to accumulate before evaluating the space usage forecasts.

Steps

1. On the dashboard, select **Storage**.
2. View the dashboard cards, Forecast of data usage by storage pool and Forecast of metadata usage by site.
3. Use these values to estimate when you will need to add new Storage Nodes for data and metadata storage.



Monitor information lifecycle management

The information lifecycle management (ILM) system provides data management for all objects stored on the grid. You must monitor ILM operations to understand if the grid can handle the current load, or if more resources are needed.

About this task

The StorageGRID system manages objects by applying the active ILM policies. The ILM policies and associated ILM rules determine how many copies are made, the type of copies that are created, where copies are placed, and the length of time each copy is retained.

Object ingest and other object-related activities can exceed the rate at which StorageGRID can evaluate ILM, causing the system to queue objects whose ILM placement instructions can't be fulfilled in near real time. You should monitor whether StorageGRID is keeping up with client actions.

Use Grid Manager dashboard tab

Steps

Use the ILM tab on the Grid Manager dashboard to monitor ILM operations:

1. Sign in to the Grid Manager.
2. From the dashboard, select the ILM tab and note the values on the ILM queue (Objects) card and ILM evaluation rate card.

Temporary spikes in the ILM queue (Objects) card on the dashboard are to be expected. But if the queue continues to increase and never declines, the grid needs more resources to operate efficiently: either more Storage Nodes, or, if the ILM policy places objects in remote locations, more network bandwidth.

Use the NODES page

Steps

Additionally, investigate ILM queues using the **NODES** page:



The charts on the **NODES** page will be replaced with the corresponding dashboard cards in a future StorageGRID release.

1. Select **NODES**.
2. Select **grid name > ILM**.
3. Position your cursor over the ILM queue graph to see the value of following attributes at a given point in

time:

- **Objects queued (from client operations):** The total number of objects awaiting ILM evaluation because of client operations (for example, ingest).
- **Objects queued (from all operations):** The total number of objects awaiting ILM evaluation.
- **Scan rate (objects/sec):** The rate at which objects in the grid are scanned and queued for ILM.
- **Evaluation rate (objects/sec):** The current rate at which objects are being evaluated against the ILM policy in the grid.

4. In the ILM Queue section, look at the following attributes.



The ILM queue section is included for the grid only. This information is not shown on the ILM tab for a site or Storage Node.

- **Scan period - estimated:** The estimated time to complete a full ILM scan of all objects.



A full scan does not guarantee that ILM has been applied to all objects.

- **Repairs attempted:** The total number of object repair operations for replicated data that have been attempted. This count increments each time a Storage Node tries to repair a high-risk object. High-risk ILM repairs are prioritized if the grid becomes busy.



The same object repair might increment again if replication failed after the repair.

These attributes can be useful when you are monitoring the progress of Storage Node volume recovery. If the number of Repairs attempted has stopped increasing and a full scan has been completed, the repair has probably completed.

Monitor networking and system resources

The integrity and bandwidth of the network between nodes and sites, and the resource usage by individual grid nodes, are critical to efficient operations.

Monitor network connections and performance

Network connectivity and bandwidth are especially important if your information lifecycle management (ILM) policy copies replicated objects between sites or stores erasure-coded objects using a scheme that provides site-loss protection. If the network between sites is not available, network latency is too high, or network bandwidth is insufficient, some ILM rules might not be able to place objects where expected. This can lead to ingest failures (when the Strict ingest option is selected for ILM rules), or to poor ingest performance and ILM backlogs.

Use the Grid Manager to monitor connectivity and network performance, so you can address any issues promptly.

Additionally, consider [creating network traffic classification policies](#) so that you can monitor traffic related to specific tenants, buckets, subnets, or load balancer endpoints. You can set traffic limiting policies as needed.

Steps

1. Select **NODES**.

The Nodes page appears. Each node in the grid is listed in table format.

DASHBOARD

ALERTS ✓

Current

Resolved

Silences

Rules

Email setup

NODES

TENANTS

ILM

CONFIGURATION

MAINTENANCE

SUPPORT

Nodes

View the list and status of sites and grid nodes.

Search...

Total node count: 14

Name	Type	Object data used	Object metadata used	CPU usage
StorageGRID Deployment	Grid	0%	0%	—
^ Data Center 1	Site	0%	0%	—
✓ DC1-ADM1	Primary Admin Node	—	—	21%
✓ DC1-ARC1	Archive Node	—	—	8%
✓ DC1-G1	Gateway Node	—	—	10%
✓ DC1-S1	Storage Node	0%	0%	29%

2. Select the grid name, a specific data center site, or a grid node, and then select the **Network** tab.

The Network Traffic graph provides a summary of overall network traffic for the grid as a whole, the data center site, or for the node.



- a. If you selected a grid node, scroll down to review the **Network Interfaces** section of the page.

Network interfaces					
Name	Hardware address	Speed	Duplex	Auto-negotiation	Link status
eth0	00:50:56:A7:66:75	10 Gigabit	Full	Off	Up

- b. For grid nodes, scroll down to review the **Network Communication** section of the page.

The Receive and Transmit tables show how many bytes and packets have been received and sent across each network as well as other receive and transmission metrics.

Network communication						
Receive						
Interface ?	Data ?	Packets ?	Errors ?	Dropped ?	Frame overruns ?	Frames ?
eth0	2.89 GB	19,421,503	0	24,032	0	0
Transmit						
Interface ?	Data ?	Packets ?	Errors ?	Dropped ?	Collisions ?	Carrier ?
eth0	3.64 GB	18,494,381	0	0	0	0

3. Use the metrics associated with your traffic classification policies to monitor network traffic.

- Select **CONFIGURATION > Network > Traffic classification**.

The Traffic Classification Policies page appears, and the existing policies are listed in the table.

Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

+ Create Edit Remove Metrics		
Name	Description	ID
☐ ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
☒ Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b
Displaying 2 traffic classification policies.		

- To view graphs that show the networking metrics associated with a policy, select the radio button to the left of the policy, and then click **Metrics**.
- Review the graphs to understand the network traffic associated with the policy.

If a traffic classification policy is designed to limit network traffic, analyze how often traffic is limited and decide if the policy continues to meet your needs. From time to time, [adjust each traffic classification policy as needed](#).

Related information

- [View the Network tab](#)
- [Monitor node connection states](#)


Monitor node-level resources

Monitor individual grid nodes to check their resource usage levels. If nodes are consistently overloaded, more nodes might be required for efficient operations.

Steps

- From the **NODES** page, select the node.

2. Select the **Hardware** tab to display graphs of CPU Utilization and Memory Usage.



3. To display a different time interval, select one of the controls above the chart or graph. You can display the information available for intervals of 1 hour, 1 day, 1 week, or 1 month. You can also set a custom interval, which allows you to specify date and time ranges.
4. If the node is hosted on a storage appliance or a services appliance, scroll down to view the tables of components. The status of all components should be "Nominal." Investigate components that have any other status.

Related information

- [View information about appliance Storage Nodes](#)
- [View information about appliance Admin Nodes and Gateway Nodes](#)

Monitor tenant activity

All S3 client activity is associated with StorageGRID tenant accounts. You can use the Grid Manager to monitor the storage usage or network traffic for all tenants or a specific tenant. You can use the audit log or Grafana dashboards to gather more detailed information about how tenants are using StorageGRID.

Before you begin

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

View all tenants

The Tenants page shows basic information for all current tenant accounts.

Steps

1. Select **TENANTS**.

2. Review the information shown on the Tenant pages.

The Logical space used, Quota usage, Quota, and Object count are listed for each tenant. If a quota is not set for a tenant, the Quota usage and Quota fields contain a dash (—).



The space used values are estimates. These estimates are affected by the timing of ingests, network connectivity, and node status.

Tenants

View information for each tenant account. Depending on the timing of ingests, network connectivity, and node status, the usage data shown might be out of date. To view more recent values, select the tenant name.

[Create](#) [Export to CSV](#) [Actions](#) Displaying 5 results

☐	Name	Logical space used	Quota utilization	Quota	Object count	Sign in/Copy URL
☐	Tenant 01	2.00 GB	10%	20.00 GB	100	→ 📄
☐	Tenant 02	85.00 GB	85%	100.00 GB	500	→ 📄
☐	Tenant 03	500.00 TB	50%	1.00 PB	10,000	→ 📄
☐	Tenant 04	475.00 TB	95%	500.00 TB	50,000	→ 📄
☐	Tenant 05	5.00 GB	—	—	500	→ 📄

- Optionally, sign in to a tenant account by selecting the sign-in link [→](#) in the **Sign in/Copy URL** column.
- Optionally, copy the URL for a tenant's sign-in page by selecting the copy URL link [📄](#) in the **Sign in/Copy URL** column.
- Optionally, select **Export to CSV** to view and export a `.csv` file containing the usage values for all tenants.

You are prompted to open or save the `.csv` file.

The contents of the `.csv` file look like the following example:

Tenant ID	Display Name	Space Used (Bytes)	Quota utilization (%)	Quota (Bytes)	Object Count	Protocol
12659822378459233654	Tenant 01	2000000000	10	20000000000	100	S3
99658234112547853685	Tenant 02	85000000000	85	1100000000	500	S3
03521145586975586321	Tenant 03	60500000000	50	150000	10000	S3
44251365987569885632	Tenant 04	4750000000	95	140000000	50000	S3
36521587546689565123	Tenant 05	5000000000	Infinity		500	S3

You can open the `.csv` file in a spreadsheet application or use it in automation.

- If no objects are listed, optionally, select **Actions > Delete** to remove one or more tenants. See [Delete tenant account](#).

You can't remove a tenant account if the account includes any buckets or containers.


View a specific tenant

You can view details for a specific tenant.

Steps

- 1. Select the tenant name from the Tenants page.

The tenant details page appears.



- 2. Review the tenant overview at the top of the page.


This section of the details page provides summary information for the tenant, including the tenant's object count, quota usage, logical space used, and quota setting.

- 3. From the **Space breakdown** tab, review the **Space consumption** chart.

This chart shows the total space consumption for all of the tenant's S3 buckets.

If a quota was set for this tenant, the amount of quota used and remaining is displayed in text (for example, 85.00 GB of 100 GB used). If no quota was set, the tenant has an unlimited quota, and the text includes only an amount of space used (for example, 85.00 GB used). The bar chart shows the percentage of quota in each bucket or container. If the tenant has exceeded the storage quota by more than 1% and by at least 1 GB, the chart shows the total quota and the excess amount.

You can place your cursor over the bar chart to see the storage used by each bucket or container. You can place your cursor over the free space segment to see the amount of storage quota remaining.



Quota usage is based on internal estimates and might be exceeded in some cases. For example, StorageGRID checks the quota when a tenant starts uploading objects and rejects new ingests if the tenant has exceeded the quota. However, StorageGRID does not take into account the size of the current upload when determining if the quota has been exceeded. If objects are deleted, a tenant might be temporarily prevented from uploading new objects until the quota usage is recalculated. Quota usage calculations can take 10 minutes or longer.



A tenant's quota usage indicates the total amount of object data the tenant has uploaded to StorageGRID (logical size). The quota usage does not represent the space used to store copies of those objects and their metadata (physical size).



You can enable the **Tenant quota usage high** alert rule to determine if tenants are consuming their quotas. If enabled, this alert is triggered when a tenant has used 90% of its quota. For instructions, see [Edit alert rules](#).

4. From the **Space breakdown** tab, review the **Bucket details**.

This table lists the S3 buckets for the tenant. Space used is the total amount of object data in the bucket or container. This value does not represent the storage space required for ILM copies and object metadata.

5. Optionally, select **Export to CSV** to view and export a .csv file containing the usage values for each bucket or container.

The contents of an individual S3 tenant's .csv file look like the following example:

Tenant ID	Bucket Name	Space Used (Bytes)	Number of Objects
64796966429038923647	bucket-01	88717711	14
64796966429038923647	bucket-02	21747507	11
64796966429038923647	bucket-03	15294070	3

You can open the .csv file in a spreadsheet application or use it in automation.

6. Optionally, select the **Allowed features** tab to see a list of the permissions and features that are enabled for the tenant. See [Edit tenant account](#) if you need to change any of these settings.
7. If the tenant has the **Use grid federation connection** permission, optionally select the **Grid federation** tab to learn more about the connection.

See [What is grid federation?](#) and [Manage the permitted tenants for grid federation](#).

View network traffic

If traffic classification policies are in place for a tenant, review the network traffic for that tenant.

Steps

1. Select **CONFIGURATION > Network > Traffic classification**.

The Traffic Classification Policies page appears, and the existing policies are listed in the table.

2. Review the list of policies to identify the ones that apply to a specific tenant.
3. To view metrics associated with a policy, select the radio button to the left of the policy, and select **Metrics**.
4. Analyze the graphs to determine how often the policy is limiting traffic and whether you need to adjust the policy.

See [Manage traffic classification policies](#) for more information.

Use the audit log

Optionally, you can use the audit log for more granular monitoring of a tenant's activities.

For instance, you can monitor the following types of information:

- Specific client operations, such as PUT, GET, or DELETE
- Object sizes
- The ILM rule applied to objects
- The source IP of client requests

Audit logs are written to text files that you can analyze using your choice of log analysis tool. This allows you to better understand client activities, or to implement sophisticated chargeback and billing models.

See [Review audit logs](#) for more information.

Use Prometheus metrics

Optionally, use Prometheus metrics to report on tenant activity.

- In the Grid Manager, select **SUPPORT > Tools > Metrics**. You can use existing dashboards, such as S3 Overview, to review client activities.



The tools available on the Metrics page are primarily intended for use by technical support. Some features and menu items within these tools are intentionally non-functional.

- From the top of the Grid Manager, select the help icon and select **API documentation**. You can use the metrics in the Metrics section of the Grid Management API to create custom alert rules and dashboards for tenant activity.

See [Review support metrics](#) for more information.

Monitor S3 client operations

You can monitor object ingest and retrieval rates as well as metrics for object counts, queries, and verification. You can view the number of successful and failed attempts by

client applications to read, write, and modify objects in the StorageGRID system.

Before you begin

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


Steps

1. From the dashboard, select the **Performance** tab.
2. Refer to the S3 charts, which summarize the number of client operations performed by Storage Nodes and the number of API requests received by Storage Nodes during the selected time frame.
3. Select **NODES** to access the Nodes page.
4. From the Nodes home page (grid level), select the **Objects** tab.

The chart shows S3 ingest and retrieve rates for your entire StorageGRID system in bytes per second and the amount of data ingested or retrieved. You can select a time interval or apply a custom interval.

5. To see information for a particular Storage Node, select the node from the list on the left, and select the **Objects** tab.

The chart shows the ingest and retrieve rates for the node. The tab also includes metrics for object counts, metadata queries, and verification operations.



Monitor load balancing operations

If you are using a load balancer to manage client connections to StorageGRID, you should monitor load balancing operations after you configure the system initially and after you make any configuration changes or perform an expansion.

About this task

You can use the Load Balancer service on Admin Nodes or Gateway Nodes or an external third-party load balancer to distribute client requests across multiple Storage Nodes.

After configuring load balancing, you should confirm that object ingest and retrieval operations are being evenly distributed across Storage Nodes. Evenly distributed requests ensure that StorageGRID remains responsive to client requests under load and can help maintain client performance.

If you configured a high availability (HA) group of Gateway Nodes or Admin Nodes in active-backup mode, only one node in the group actively distributes client requests.

For more information, see [Configure S3 client connections](#).

Steps

1. If S3 clients connect using the Load Balancer service, check that Admin Nodes or Gateway Nodes are actively distributing traffic as you expect:
 - a. Select **NODES**.
 - b. Select a Gateway Node or Admin Node.
 - c. On the **Overview** tab, check if a node interface is in an HA group and if the node interface has the role of Primary.

Nodes with the role of Primary and nodes that aren't in an HA group should be actively distributing requests to clients.

- d. For each node that should be actively distributing client requests, select the [Load Balancer tab](#).
 - e. Review the chart of Load Balancer Request Traffic for the last week to ensure that the node has been actively distributing requests.

Nodes in an active-backup HA group might take the Backup role from time to time. During that time the nodes don't distribute client requests.

- f. Review the chart of Load Balancer Incoming Request Rate for the last week to review the object throughput of the node.
 - g. Repeat these steps for each Admin Node or Gateway Node in the StorageGRID system.
 - h. Optionally, use traffic classification policies to view a more detailed analysis of traffic being served by the Load Balancer service.
2. Verify that these requests are being evenly distributed to Storage Nodes.
 - a. Select **Storage Node > LDR > HTTP**.
 - b. Review the number of **Currently Established Incoming Sessions**.
 - c. Repeat for each Storage Node in the grid.

The number of sessions should be roughly equal across all Storage Nodes.

Monitor grid federation connections

You can monitor basic information about all [grid federation connections](#), detailed information about a specific connection, or Prometheus metrics about cross-grid replication operations. You can monitor a connection from either grid.

Before you begin

- You are signed in to the Grid Manager on either grid using a [supported web browser](#).
- You have the [Root access permission](#) for the grid you are signed in to.

View all connections

The Grid federation page shows basic information about all grid federation connections and about all tenant accounts that are permitted to use grid federation connections.

Steps


1. Select **CONFIGURATION** > **System** > **Grid federation**.

The Grid federation page appears.

2. To see basic information for all connections on this grid, select the **Connections** tab.

From this tab, you can:

- [Create a new connection](#).
- Select an existing connection to [edit or test](#).



The screenshot shows the 'Grid federation' page with the 'Connections' tab selected. The page includes a header with a 'Learn more about grid federation' link, a description of grid federation, and two tabs: 'Connections' (active) and 'Permitted tenants'. Below the tabs are buttons for 'Add connection', 'Upload verification file', and 'Actions', along with a search bar. A table displays one connection: 'Grid 1 - Grid 2' with remote hostname '10.96.130.76' and status 'Connected'.

Connection name	Remote hostname	Connection status
Grid 1 - Grid 2	10.96.130.76	Connected

3. To see basic information for all tenant accounts on this grid that have the **Use grid federation connection** permission, select the **Permitted tenants** tab.

From this tab, you can:

- [View the details page for each permitted tenant](#).
- View the details page for each connection. See [View a specific connection](#).
- Select a permitted tenant and [remove the permission](#).
- Check for cross-grid replication errors and clear the last error, if any. See [Troubleshoot grid federation errors](#).

Grid federation [Learn more about grid federation](#)

You can use grid federation to clone tenant accounts and replicate their objects between two StorageGRID systems. Grid federation uses a trusted and secure connection between Admin and Gateway Nodes in two discrete StorageGRID systems.

[Connections](#)
[Permitted tenants](#)

[Remove permission](#)
[Clear error](#)

Displaying one result

Tenant name	Connection name	Connection status	Remote grid hostname	Last error
Tenant A	Grid 1 - Grid 2	Connected	10.96.130.76	Check for errors

View a specific connection

You can view details for a specific grid federation connection.

Steps

1. Select either tab from the Grid federation page and then select the connection name from the table.

From the details page for the connection, you can:

- See basic status information about the connection, including the local and remote hostnames, port, and connection status.
- Select a connection to [edit](#), [test](#), or [remove](#).

2. When viewing a specific connection, select the **Permitted tenants** tab to view details about the permitted tenants for the connection.

From this tab, you can:

- [View the details page for each permitted tenant](#).
- [Remove a tenant's permission](#) to use the connection.
- Check for cross-grid replication errors and clear the last error. See [Troubleshoot grid federation errors](#).

Grid 1 - Grid 2

Local hostname (this grid):

10.96.130.64

Port:

23000

Remote hostname (other grid):

10.96.130.76

Connection status:

✓

Connected

Edit

Download file

Test connection

Remove

Permitted tenants

Certificates

Remove permission

Clear error

Search...

Displaying one result

Tenant name	Last error
Tenant A	Check for errors

3. When viewing a specific connection, select the **Certificates** tab to view the system-generated server and client certificates for this connection.

From this tab, you can:

- [Rotate connection certificates](#).
- Select **Server** or **Client** to view or download the associated certificate or copy the certificate PEM.

Grid A-Grid B

Local hostname (this grid):10.96.106.230

Port:23000

Remote hostname (other grid):10.96.104.230

Connection status:

Connected

Edit

Download file

Test connection

Remove

Permitted tenants

Certificates

Rotate certificates

Server

Client

Download certificate

Copy certificate PEM

Metadata

Subject DN:/C=US/ST=California/L=Sunnyvale/O=NetApp Inc./OU=NetApp StorageGRID/CN=10.96.106.230

Serial number:30:81:B8:DD:AE:B2:86:0A

Issuer DN:/C=US/ST=California/L=Sunnyvale/O=NetApp Inc./OU=NetApp StorageGRID/CN=GPT

Issued on:2022-10-04T02:21:18.000Z

Expires on:2024-10-03T19:05:13.000Z

SHA-1 fingerprint:92:7A:03:AF:6D:1C:94:8C:33:24:08:84:F9:2B:01:23:7D:BE:F2:DF

SHA-256 fingerprint:54:97:3E:77:EB:D3:6A:0F:8F:EE:72:83:D0:39:86:02:32:A5:60:9D:6F:C0:A2:3C:76:DA:3F:4D:FF:64:5D:60

Alternative names:IP Address:10.96.106.230

Certificate PEM

-----BEGIN CERTIFICATE-----

MIIGdTCCBF2gAwIBAgIIMIG43a6yhgowDQYJKoZIhvcNAQENBQAwdzELMAkGA1UE


BhMCVVMhxEzARBgNVBAGhNCKNhbg1mb3JuaWExEjAQBGNVBAcMCVN1bm55dmFsZTEU

NETAPPSUNNYVAZBm55dmFsZTEU

-----END CERTIFICATE-----

You can use the Cross-Grid Replication dashboard in Grafana to view Prometheus metrics about cross-grid replication operations on your grid.

1. From the Grid Manager, select **SUPPORT** > **Tools** > **Metrics**.



- For detailed instructions, see [Review support metrics](#).

3. To retry replication of objects that failed to replicate, see [Identify and retry failed replication operations](#).

Manage alerts

Manage alerts

The alert system provides an easy-to-use interface for detecting, evaluating, and resolving the issues that can occur during StorageGRID operation.

Alerts are triggered at specific severity levels when alert rule conditions evaluate as true. When an alert is triggered, the following actions occur:

- An alert severity icon is shown on the dashboard in the Grid Manager, and the count of Current Alerts is incremented.
- The alert is shown on the **NODES** summary page and on the **NODES > node > Overview** tab.
- An email notification is sent, assuming you have configured an SMTP server and provided email addresses for the recipients.
- An Simple Network Management Protocol (SNMP) notification is sent, assuming you have configured the StorageGRID SNMP agent.


You can create custom alerts, edit or disable alerts, and manage alert notifications.

To learn more:

- Review the video: [Video: Alerts overview](#)



- Review the video: [Video: Custom alerts](#)



- See the [Alerts reference](#).

View alert rules

Alert rules define the conditions that trigger [specific alerts](#). StorageGRID includes a set of default alert rules, which you can use as is or modify, or you can create custom alert rules.

You can view the list of all default and custom alert rules to learn which conditions will trigger each alert and to see whether any alerts are disabled.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Manage alerts](#) or [Root access permission](#).
- Optionally, you have watched the video: [Video: Alerts overview](#)



Steps

1. Select **ALERTS > Rules**.

The Alert Rules page appears.

Alert Rules [Learn more](#)

Alert rules define which conditions trigger specific alerts.

You can edit the conditions for default alert rules to better suit your environment, or create custom alert rules that use your own conditions for triggering alerts.

[+ Create custom rule](#) [Edit rule](#) [Remove custom rule](#)

Name	Conditions	Type	Status
☐ Appliance battery expired The battery in the appliance's storage controller has expired.	storagegrid_appliance_component_failure(type="REC_EXPIRED_BATTERY") Major > 0	Default	Enabled
☐ Appliance battery failed The battery in the appliance's storage controller has failed.	storagegrid_appliance_component_failure(type="REC_FAILED_BATTERY") Major > 0	Default	Enabled
☐ Appliance battery has insufficient learned capacity The battery in the appliance's storage controller has insufficient learned capacity.	storagegrid_appliance_component_failure(type="REC_BATTERY_WARN") Major > 0	Default	Enabled
☐ Appliance battery near expiration The battery in the appliance's storage controller is nearing expiration.	storagegrid_appliance_component_failure(type="REC_BATTERY_NEAR_EXPIRATION") Major > 0	Default	Enabled
☐ Appliance battery removed The battery in the appliance's storage controller is missing.	storagegrid_appliance_component_failure(type="REC_REMOVED_BATTERY") Major > 0	Default	Enabled
☐ Appliance battery too hot The battery in the appliance's storage controller is overheated.	storagegrid_appliance_component_failure(type="REC_BATTERY_OVERTEMP") Major > 0	Default	Enabled
☐ Appliance cache backup device failed A persistent cache backup device has failed.	storagegrid_appliance_component_failure(type="REC_CACHE_BACKUP_DEVICE_FAILED") Major > 0	Default	Enabled
☐ Appliance cache backup device insufficient capacity There is insufficient cache backup device capacity.	storagegrid_appliance_component_failure(type="REC_CACHE_BACKUP_DEVICE_INSUFFICIENT_CAPACITY") Major > 0	Default	Enabled
☐ Appliance cache backup device write-protected A cache backup device is write-protected.	storagegrid_appliance_component_failure(type="REC_CACHE_BACKUP_DEVICE_WRITE_PROTECTED") Major > 0	Default	Enabled
☐ Appliance cache memory size mismatch The two controllers in the appliance have different cache sizes.	storagegrid_appliance_component_failure(type="REC_CACHE_MEM_SIZE_MISMATCH") Major > 0	Default	Enabled

Displaying 62 alert rules.

2. Review the information in the alert rules table:

Column header	Description
Name	The unique name and description of the alert rule. Custom alert rules are listed first, followed by default alert rules. The alert rule name is the subject for email notifications.
Conditions	The Prometheus expressions that determine when this alert is triggered. An alert can be triggered at one or more of the following severity levels, but a condition for each severity is not required. • Critical : An abnormal condition exists that has stopped the normal operations of a StorageGRID node or service. You must address the underlying issue immediately. Service disruption and loss of data might result if the issue is not resolved. • Major : An abnormal condition exists that is either affecting current operations or approaching the threshold for a critical alert. You should investigate major alerts and address any underlying issues to ensure that the abnormal condition does not stop the normal operation of a StorageGRID node or service. • Minor : The system is operating normally, but an abnormal condition exists that could affect the system's ability to operate if it continues. You should monitor and resolve minor alerts that don't clear on their own to ensure they don't result in a more serious problem.
Type	The type of alert rule: • Default: An alert rule provided with the system. You can disable a default alert rule or edit the conditions and duration for a default alert rule. You can't remove a default alert rule. • Default*: A default alert rule that includes an edited condition or duration. As required, you can easily revert a modified condition back to the original default. • Custom: An alert rule that you created. You can disable, edit, and remove custom alert rules.
Status	Whether this alert rule is currently enabled or disabled. The conditions for disabled alert rules aren't evaluated, so no alerts are triggered.


Create custom alert rules

You can create custom alert rules to define your own conditions for triggering alerts.

Before you begin

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

- You are familiar with the [commonly used Prometheus metrics](#).
- You understand the [syntax of Prometheus queries](#).
- Optionally, you have watched the video: [Video: Custom alerts](#).



About this task

StorageGRID does not validate custom alerts. If you decide to create custom alert rules, follow these general guidelines:

- Look at the conditions for the default alert rules, and use them as examples for your custom alert rules.
- If you define more than one condition for an alert rule, use the same expression for all conditions. Then, change the threshold value for each condition.
- Carefully check each condition for typos and logic errors.
- Use only the metrics listed in the Grid Management API.
- When testing an expression using the Grid Management API, be aware that a "successful" response might be an empty response body (no alert triggered). To see if the alert is actually triggered, you can temporarily set a threshold to a value you expect to be true currently.

For example, to test the expression `node_memory_MemTotal_bytes < 24000000000`, first execute `node_memory_MemTotal_bytes >= 0` and ensure you get the expected results (all nodes return a value). Then, change the operator and the threshold back to the intended values and execute again. No results indicate there are no current alerts for this expression.

- Don't assume a custom alert is working unless you have validated that the alert is triggered when expected.

Steps

1. Select **ALERTS > Rules**.

The Alert Rules page appears.

2. Select **Create custom rule**.

The Create Custom Rule dialog box appears.

Create Custom Rule

Enabled ☒

Unique Name

Description

Recommended Actions
(optional)

Conditions

Minor

Major

Critical

Enter the amount of time a condition must continuously remain in effect before an alert is triggered.

Duration

5

minutes

Cancel

Save

3. Select or clear the **Enabled** checkbox to determine if this alert rule is currently enabled.

If an alert rule is disabled, its expressions aren't evaluated and no alerts are triggered.

4. Enter the following information:

Field	Description
Unique Name	A unique name for this rule. The alert rule name is shown on the Alerts page and is also the subject for email notifications. Names for alert rules can be between 1 and 64 characters.
Description	A description of the problem that is occurring. The description is the alert message shown on the Alerts page and in email notifications. Descriptions for alert rules can be between 1 and 128 characters.

Field	Description
Recommended Actions	Optionally, the recommended actions to take when this alert is triggered. Enter recommended actions as plain text (no formatting codes). Recommended actions for alert rules can be between 0 and 1,024 characters.

5. In the Conditions section, enter a Prometheus expression for one or more of the alert severity levels.

A basic expression is usually of the form:

```
[metric] [operator] [value]
```

Expressions can be any length, but appear on a single line in the user interface. At least one expression is required.

This expression causes an alert to be triggered if the amount of installed RAM for a node is less than 24,000,000,000 bytes (24 GB).

```
node_memory_MemTotal_bytes < 24000000000
```

To see available metrics and to test Prometheus expressions, select the help icon  and follow the link to the Metrics section of the Grid Management API.

6. In the **Duration** field, enter the amount of time a condition must continuously remain in effect before the alert is triggered, and select a unit of time.

To trigger an alert immediately when a condition becomes true, enter **0**. Increase this value to prevent temporary conditions from triggering alerts.

The default is 5 minutes.

7. Select **Save**.

The dialog box closes, and the new custom alert rule appears in the Alert Rules table.

Edit alert rules

You can edit an alert rule to change the trigger conditions, For a custom alert rule, you can also update the rule name, description, and recommended actions.

Before you begin

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

About this task

When you edit a default alert rule, you can change the conditions for minor, major, and critical alerts; and the duration. When you edit a custom alert rule, you can also edit the rule's name, description, and recommended actions.



Be careful when deciding to edit an alert rule. If you change trigger values, you might not detect an underlying problem until it prevents a critical operation from completing.

Steps

1. Select **ALERTS > Rules**.

The Alert Rules page appears.

2. Select the radio button for the alert rule you want to edit.
3. Select **Edit rule**.

The Edit Rule dialog box appears. This example shows a default alert rule—the Unique Name, Description, and Recommended Actions fields are disabled and can't be edited.

Edit Rule - Low installed node memory

Enabled ☒

Unique Name

Low installed node memory

Description

The amount of installed memory on a node is low.

Recommended Actions (optional)

Increase the amount of RAM available to the virtual machine or Linux host. Check the threshold value for the major alert to determine the default minimum requirement for a StorageGRID node.

See the instructions for your platform:

- [VMware installation](#)
- [Red Hat Enterprise Linux or CentOS installation](#)
- [Ubuntu or Debian installation](#)

Conditions ?

Minor

Major

node_memory_MemTotal_bytes < 24000000000

Critical

node_memory_MemTotal_bytes <= 12000000000

Enter the amount of time a condition must continuously remain in effect before an alert is triggered.

Duration

2

minutes

Cancel

Save

4. Select or clear the **Enabled** checkbox to determine if this alert rule is currently enabled.

If an alert rule is disabled, its expressions aren't evaluated and no alerts are triggered.

If you disable the alert rule for a current alert, you must wait a few minutes for the alert to no longer appear as an active alert.

In general, disabling a default alert rule is not recommended. If an alert rule is disabled, you might not detect an underlying problem until it prevents a critical operation from completing.

1038

5. For custom alert rules, update the following information as required.



You can't edit this information for default alert rules.

Field	Description
Unique Name	A unique name for this rule. The alert rule name is shown on the Alerts page and is also the subject for email notifications. Names for alert rules can be between 1 and 64 characters.
Description	A description of the problem that is occurring. The description is the alert message shown on the Alerts page and in email notifications. Descriptions for alert rules can be between 1 and 128 characters.
Recommended Actions	Optionally, the recommended actions to take when this alert is triggered. Enter recommended actions as plain text (no formatting codes). Recommended actions for alert rules can be between 0 and 1,024 characters.

6. In the Conditions section, enter or update the Prometheus expression for one or more of the alert severity levels.



If you want to restore a condition for an edited default alert rule back to its original value, select the three dots to the right of the modified condition.

Conditions

Minor	
Major	node_memory_MemTotal_bytes < 24000000000
Critical	node_memory_MemTotal_bytes <= 14000000000



If you update the conditions for a current alert, your changes might not be implemented until the previous condition is resolved. The next time one of the conditions for the rule is met, the alert will reflect the updated values.

A basic expression is usually of the form:

```
[metric] [operator] [value]
```

Expressions can be any length, but appear on a single line in the user interface. At least one expression is required.

This expression causes an alert to be triggered if the amount of installed RAM for a node is less than 24,000,000,000 bytes (24 GB).

```
node_memory_MemTotal_bytes < 24000000000
```

7. In the **Duration** field, enter the amount of time a condition must continuously remain in effect before the

alert is triggered, and select the unit of time.

To trigger an alert immediately when a condition becomes true, enter **0**. Increase this value to prevent temporary conditions from triggering alerts.

The default is 5 minutes.

8. Select **Save**.

If you edited a default alert rule, **Default*** appears in the Type column. If you disabled a default or custom alert rule, **Disabled** appears in the **Status** column.

Disable alert rules

You can change the enabled/disabled state for a default or custom alert rule.

Before you begin

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

About this task

When an alert rule is disabled, its expressions aren't evaluated and no alerts are triggered.



In general, disabling a default alert rule is not recommended. If an alert rule is disabled, you might not detect an underlying problem until it prevents a critical operation from completing.

Steps

1. Select **ALERTS > Rules**.

The Alert Rules page appears.

2. Select the radio button for the alert rule you want to disable or enable.

3. Select **Edit rule**.

The Edit Rule dialog box appears.

4. Select or clear the **Enabled** checkbox to determine if this alert rule is currently enabled.

If an alert rule is disabled, its expressions aren't evaluated and no alerts are triggered.



If you disable the alert rule for a current alert, you must wait a few minutes for the alert to no longer display as an active alert.

5. Select **Save**.

Disabled appears in the **Status** column.

Remove custom alert rules

You can remove a custom alert rule if you no longer want to use it.

Before you begin

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

Steps

1. Select **ALERTS > Rules**.

The Alert Rules page appears.

2. Select the radio button for the custom alert rule you want to remove.

You can't remove a default alert rule.

3. Select **Remove custom rule**.

A confirmation dialog box appears.

4. Select **OK** to remove the alert rule.

Any active instances of the alert will be resolved within 10 minutes.

Manage alert notifications

Set up SNMP notifications for alerts

If you want StorageGRID to send SNMP notifications when alerts occur, you must enable the StorageGRID SNMP agent and configure one or more trap destinations.

You can use the **CONFIGURATION > Monitoring > SNMP agent** option in the Grid Manager or the SNMP endpoints for the Grid Management API to enable and configure the StorageGRID SNMP agent. The SNMP agent supports all three versions of the SNMP protocol.

To learn how to configure the SNMP agent, see [Use SNMP monitoring](#).

After you configure the StorageGRID SNMP agent, two types of event-driven notifications can be sent:

- Traps are notifications sent by the SNMP agent that don't require acknowledgment by the management system. Traps serve to notify the management system that something has happened within StorageGRID, such as an alert being triggered. Traps are supported in all three versions of SNMP.
- Informs are similar to traps, but they require acknowledgment by the management system. If the SNMP agent does not receive an acknowledgment within a certain amount of time, it resends the inform until an acknowledgment is received or the maximum retry value has been reached. Informs are supported in SNMPv2c and SNMPv3.

Trap and inform notifications are sent when a default or custom alert is triggered at any severity level. To suppress SNMP notifications for an alert, you must configure a silence for the alert. See [Silence alert notifications](#).

If your StorageGRID deployment includes multiple Admin Nodes, the primary Admin Node is the preferred sender for alert notifications, AutoSupport packages, and SNMP traps and informs. If the primary Admin Node becomes unavailable, notifications are temporarily sent by other Admin Nodes. See [What is an Admin Node?](#).

Set up email notifications for alerts

If you want email notifications to be sent when alerts occur, you must provide information about your SMTP server. You must also enter email addresses for the recipients of alert notifications.

Before you begin

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

About this task

The email setup used for alert notifications is not used for AutoSupport packages. However, you can use the same email server for all notifications.

If your StorageGRID deployment includes multiple Admin Nodes, the primary Admin Node is the preferred sender for alert notifications, AutoSupport packages, and SNMP traps and informs. If the primary Admin Node becomes unavailable, notifications are temporarily sent by other Admin Nodes. See [What is an Admin Node?](#).

Steps

1. Select **ALERTS > Email setup**.

The Email Setup page appears.

2. Select the **Enable Email Notifications** checkbox to indicate that you want notification emails to be sent when alerts reach configured thresholds.

The Email (SMTP) Server, Transport Layer Security (TLS), Email Addresses, and Filters sections appear.

3. In the Email (SMTP) Server section, enter the information StorageGRID needs to access your SMTP server.

If your SMTP server requires authentication, you must provide both a username and a password.

Field	Enter
Mail Server	The fully qualified domain name (FQDN) or IP address of the SMTP server.
Port	The port used to access the SMTP server. Must be between 1 and 65535.
Username (optional)	If your SMTP server requires authentication, enter the username to authenticate with.
Password (optional)	If your SMTP server requires authentication, enter the password to authenticate with.

4. In the Email Addresses section, enter email addresses for the sender and for each recipient.
 - a. For the **Sender Email Address**, specify a valid email address to use as the From address for alert notifications.

For example: `storagegrid-alerts@example.com`

- b. In the Recipients section, enter an email address for each email list or person who should receive an email when an alert occurs.

Select the plus icon **+** to add recipients.

5. If Transport Layer Security (TLS) is required for communications with the SMTP server, select **Require TLS** in the Transport Layer Security (TLS) section.

- a. In the **CA Certificate** field, provide the CA certificate that will be used to verify the identify of the SMTP server.

You can copy and paste the contents into this field, or select **Browse** and select the file.

You must provide a single file that contains the certificates from each intermediate issuing certificate authority (CA). The file should contain each of the PEM-encoded CA certificate files, concatenated in certificate chain order.

- b. Select the **Send Client Certificate** checkbox if your SMTP email server requires email senders to provide client certificates for authentication.
- c. In the **Client Certificate** field, provide the PEM-encoded client certificate to send to the SMTP server.

You can copy and paste the contents into this field, or select **Browse** and select the file.

- d. In the **Private Key** field, enter the private key for the client certificate in unencrypted PEM encoding.

You can copy and paste the contents into this field, or select **Browse** and select the file.



If you need to edit the email setup, select the pencil icon  to update this field.

6. In the Filters section, select which alert severity levels should result in email notifications, unless the rule for a specific alert has been silenced.

Severity	Description
Minor, major, critical	An email notification is sent when the minor, major, or critical condition for an alert rule is met.
Major, critical	An email notification is sent when the major or critical condition for an alert rule is met. Notifications aren't sent for minor alerts.
Critical only	An email notification is sent only when the critical condition for an alert rule is met. Notifications aren't sent for minor or major alerts.

7. When you are ready to test your email settings, perform these steps:

- a. Select **Send Test Email**.

A confirmation message appears, indicating that a test email was sent.

- b. Check the inboxes of all email recipients and confirm that a test email was received.



If the email is not received within a few minutes or if the **Email notification failure** alert is triggered, check your settings and try again.

c. Sign in to any other Admin Nodes and send a test email to verify connectivity from all sites.



When you test alert notifications, you must sign in to every Admin Node to verify connectivity. This is in contrast to testing AutoSupport packages, where all Admin Nodes send the test email.

8. Select **Save**.

Sending a test email does not save your settings. You must select **Save**.

The email settings are saved.

Information included in alert email notifications

After you configure the SMTP email server, email notifications are sent to the designated recipients when an alert is triggered, unless the alert rule is suppressed by a silence. See [Silence alert notifications](#).

Email notifications include the following information:

NetApp StorageGRID

Low object data storage (6 alerts) 1

The space available for storing object data is low. 2

Recommended actions 3

Perform an expansion procedure. You can add storage volumes (LUNs) to existing Storage Nodes, or you can add new Storage Nodes. See the instructions for expanding a StorageGRID system.

DC1-S1-226

Node	DC1-S1-226 4
Site	DC1 225-230
Severity	Minor
Time triggered	Fri Jun 28 14:43:27 UTC 2019
Job	storagegrid
Service	ldr

DC1-S2-227

Node	DC1-S2-227
Site	DC1 225-230
Severity	Minor
Time triggered	Fri Jun 28 14:43:27 UTC 2019
Job	storagegrid
Service	ldr

5

Sent from: DC1-ADM1-225

Callout	Description
1	The name of the alert, followed by the number of active instances of this alert.
2	The description of the alert.
3	Any recommended actions for the alert.
4	Details about each active instance of the alert, including the node and site affected, the alert severity, the UTC time when the alert rule was triggered, and the name of the affected job and service.
5	The hostname of the Admin Node that sent the notification.

How alerts are grouped

To prevent an excessive number of email notifications from being sent when alerts are triggered, StorageGRID attempts to group multiple alerts in the same notification.

Refer to the following table for examples of how StorageGRID groups multiple alerts in email notifications.

Behavior	Example
Each alert notification applies only to alerts that have the same name. If two alerts with different names are triggered at the same time, two email notifications are sent.	- Alert A is triggered on two nodes at the same time. Only one notification is sent. - Alert A is triggered on node 1, and Alert B is triggered on node 2 at the same time. Two notifications are sent—one for each alert.
For a specific alert on a specific node, if the thresholds are reached for more than one severity, a notification is sent only for the most severe alert.	Alert A is triggered and the minor, major, and critical alert thresholds are reached. One notification is sent for the critical alert.
The first time an alert is triggered, StorageGRID waits 2 minutes before sending a notification. If other alerts with the same name are triggered during that time, StorageGRID groups all of the alerts in the initial notification.	- Alert A is triggered on node 1 at 08:00. No notification is sent. - Alert A is triggered on node 2 at 08:01. No notification is sent. - At 08:02, a notification is sent to report both instances of the alert.
If an another alert with the same name is triggered, StorageGRID waits 10 minutes before sending a new notification. The new notification reports all active alerts (current alerts that have not been silenced), even if they were reported previously.	- Alert A is triggered on node 1 at 08:00. A notification is sent at 08:02. - Alert A is triggered on node 2 at 08:05. A second notification is sent at 08:15 (10 minutes later). Both nodes are reported.

Behavior	Example
If there are multiple current alerts with the same name and one of those alerts is resolved, a new notification is not sent if the alert reoccurs on the node for which the alert was resolved.	1. Alert A is triggered for node 1. A notification is sent. 2. Alert A is triggered for node 2. A second notification is sent. 3. Alert A is resolved for node 2, but it remains active for node 1. 4. Alert A is triggered again for node 2. No new notification is sent because the alert is still active for node 1.
StorageGRID continues to send email notifications once every 7 days until all instances of the alert are resolved or the alert rule is silenced.	1. Alert A is triggered for node 1 on March 8. A notification is sent. 2. Alert A is not resolved or silenced. Additional notifications are sent on March 15, March 22, March 29, and so on.

Troubleshoot alert email notifications

If the **Email notification failure** alert is triggered or you are unable to receive the test alert email notification, follow these steps to resolve the issue.

Before you begin

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

Steps

1. Verify your settings.
 - a. Select **ALERTS > Email setup**.
 - b. Verify that the Email (SMTP) Server settings are correct.
 - c. Verify that you have specified valid email addresses for the recipients.
2. Check your spam filter, and make sure that the email was not sent to a junk folder.
3. Ask your email administrator to confirm that emails from the sender address aren't being blocked.
4. Collect a log file for the Admin Node, and then contact technical support.

Technical support can use the information in the logs to help determine what went wrong. For example, the `prometheus.log` file might show an error when connecting to the server you specified.

See [Collect log files and system data](#).

Silence alert notifications

Optionally, you can configure silences to temporarily suppress alert notifications.

Before you begin

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

- You have the [Manage alerts or Root access permission](#).

About this task

You can silence alert rules on the entire grid, a single site, or a single node and for one or more severities. Each silence suppresses all notifications for a single alert rule or for all alert rules.

If you have enabled the SNMP agent, silences also suppress SNMP traps and informs.



Be careful when deciding to silence an alert rule. If you silence an alert, you might not detect an underlying problem until it prevents a critical operation from completing.

Steps

1. Select **ALERTS > Silences**.

The Silences page appears.

Silences

You can configure silences to temporarily suppress alert notifications. Each silence suppresses the notifications for an alert rule at one or more severities. You can suppress an alert rule on the entire grid, a single site, or a single node.

+ Create ✎ Edit ✕ Remove				
Alert Rule	Description	Severity	Time Remaining	Nodes
No results found.				

2. Select **Create**.

The Create Silence dialog box appears.

Create Silence

Alert Rule		
Description (optional)		
Duration		Minutes ▼
Severity	☐ Minor only ☐ Minor, major ☐ Minor, major, critical	
Nodes	☐ StorageGRID Deployment ☐ Data Center 1 ☐ DC1-ADM1 ☐ DC1-G1 ☐ DC1-S1 ☐ DC1-S2 ☐ DC1-S3	
Cancel Save		

3. Select or enter the following information:

Field	Description
Alert Rule	The name of the alert rule you want to silence. You can select any default or custom alert rule, even if the alert rule is disabled. Note: Select All rules if you want to silence all alert rules using the criteria specified in this dialog box.
Description	Optionally, a description of the silence. For example, describe the purpose of this silence.
Duration	How long you want this silence to remain in effect, in minutes, hours, or days. A silence can be in effect from 5 minutes to 1,825 days (5 years). Note: You should not silence an alert rule for an extended amount of time. If an alert rule is silenced, you might not detect an underlying problem until it prevents a critical operation from completing. However, you might need to use an extended silence if an alert is triggered by a specific, intentional configuration, such as might be the case for the Services appliance link down alerts and the Storage appliance link down alerts.
Severity	Which alert severity or severities should be silenced. If the alert is triggered at one of the selected severities, no notifications are sent.

Field	Description
Nodes	Which node or nodes you want this silence to apply to. You can suppress an alert rule or all rules on the entire grid, a single site, or a single node. If you select the entire grid, the silence applies to all sites and all nodes. If you select a site, the silence applies only to the nodes at that site. Note: You can't select more than one node or more than one site for each silence. You must create additional silences if you want to suppress the same alert rule on more than one node or more than one site at one time.

4. Select **Save**.

5. If you want to modify or end a silence before it expires, you can edit or remove it.

Option	Description
Edit a silence	- Select ALERTS > Silences. - From the table, select the radio button for the silence you want to edit. - Select Edit. - Change the description, the amount of time remaining, the selected severities, or the affected node. - Select Save.
Remove a silence	- Select ALERTS > Silences. - From the table, select the radio button for the silence you want to remove. - Select Remove. - Select OK to confirm you want to remove this silence. Note: Notifications will now be sent when this alert is triggered (unless suppressed by another silence). If this alert is currently triggered, it might take few minutes for email or SNMP notifications to be sent and for the Alerts page to update.

Related information

[Configure the SNMP agent](#)

Alerts reference

This reference lists the default alerts that appear in the Grid Manager. Recommended actions are in the alert message you receive.

As required, you can create custom alert rules to fit your system management approach.

Some of the default alerts use [Prometheus metrics](#).

Appliance alerts

Alert name	Description
Appliance battery expired	The battery in the appliance's storage controller has expired.
Appliance battery failed	The battery in the appliance's storage controller has failed.
Appliance battery has insufficient learned capacity	The battery in the appliance's storage controller has insufficient learned capacity.
Appliance battery near expiration	The battery in the appliance's storage controller is nearing expiration.
Appliance battery removed	The battery in the appliance's storage controller is missing.
Appliance battery too hot	The battery in the appliance's storage controller is overheated.
Appliance BMC communication error	Communication with the baseboard management controller (BMC) has been lost.
Appliance boot device fault detected	A problem was detected with the boot device in the appliance.
Appliance cache backup device failed	A persistent cache backup device has failed.
Appliance cache backup device insufficient capacity	There is insufficient cache backup device capacity.
Appliance cache backup device write-protected	A cache backup device is write-protected.
Appliance cache memory size mismatch	The two controllers in the appliance have different cache sizes.
Appliance CMOS battery fault	A problem was detected with the CMOS battery in the appliance.
Appliance compute controller chassis temperature too high	The temperature of the compute controller in a StorageGRID appliance has exceeded a nominal threshold.
Appliance compute controller CPU temperature too high	The temperature of the CPU in the compute controller in a StorageGRID appliance has exceeded a nominal threshold.
Appliance compute controller needs attention	A hardware fault has been detected in the compute controller of a StorageGRID appliance.

Alert name	Description
Appliance compute controller power supply A has a problem	Power supply A in the compute controller has a problem.
Appliance compute controller power supply B has a problem	Power supply B in the compute controller has a problem.
Appliance compute hardware monitor service stalled	The service that monitors storage hardware status has stalled.
Appliance DAS drive exceeding limit for data written per day	An excessive amount of data is being written to a drive each day, which might void its warranty.
Appliance DAS drive fault detected	A problem was detected with a direct-attached storage (DAS) drive in the appliance.
Appliance DAS drive locator light on	The drive locator light for one or more direct-attached storage (DAS) drives in an appliance Storage Node is on.
Appliance DAS drive rebuilding	A direct-attached storage (DAS) drive is rebuilding. This is expected if it was recently replaced or removed/reinserted.
Appliance fan fault detected	A problem with a fan unit in the appliance was detected.
Appliance Fibre Channel fault detected	A Fibre Channel link problem has been detected between the appliance storage controller and compute controller
Appliance Fibre Channel HBA port failure	A Fibre Channel HBA port is failing or has failed.
Appliance flash cache drives non-optimal	The drives used for the SSD cache are non-optimal.
Appliance interconnect/battery canister removed	The interconnect/battery canister is missing.
Appliance LACP port missing	A port on a StorageGRID appliance is not participating in the LACP bond.
Appliance NIC fault detected	A problem with a network interface card (NIC) in the appliance was detected.
Appliance overall power supply degraded	The power of a StorageGRID appliance has deviated from the recommended operating voltage.
Appliance SSD critical warning	An appliance SSD is reporting a critical warning.

Alert name	Description
Appliance storage controller A failure	Storage controller A in a StorageGRID appliance has failed.
Appliance storage controller B failure	Storage controller B in a StorageGRID appliance has failed.
Appliance storage controller drive failure	One or more drives in a StorageGRID appliance has failed or is not optimal.
Appliance storage controller hardware issue	SANtricity software is reporting "Needs attention" for a component in a StorageGRID appliance.
Appliance storage controller power supply A failure	Power supply A in a StorageGRID appliance has deviated from the recommended operating voltage.
Appliance storage controller power supply B failure	Power supply B in a StorageGRID appliance has deviated from the recommended operating voltage.
Appliance storage hardware monitor service stalled	The service that monitors storage hardware status has stalled.
Appliance storage shelves degraded	The status of one of the components in the storage shelf for a storage appliance is degraded.
Appliance temperature exceeded	The nominal or maximum temperature for the appliance's storage controller has been exceeded.
Appliance temperature sensor removed	A temperature sensor has been removed.
Appliance UEFI secure boot error	An appliance has not been booted securely.
Disk I/O is very slow	Very slow disk I/O might be impacting grid performance.
Storage appliance fan fault detected	A problem with a fan unit in the storage controller for an appliance was detected.
Storage appliance storage connectivity degraded	There is a problem with one or more connections between the compute controller and storage controller.
Storage device inaccessible	A storage device cannot be accessed.

Audit and syslog alerts

Alert name	Description
Audit logs are being added to the in-memory queue	Node cannot send logs to the local syslog server and the in-memory queue is filling up.
External syslog server forwarding error	Node cannot forward logs to the external syslog server.
Large audit queue	The disk queue for audit messages is full. If this condition is not addressed, S3 or Swift operations might fail.
Logs are being added to the on-disk queue	Node cannot forward logs to the external syslog server and the on-disk queue is filling up.

Bucket alerts

Alert name	Description
FabricPool bucket has unsupported bucket consistency setting	A FabricPool bucket uses the Available or Strong-site consistency level, which is not supported.
FabricPool bucket has unsupported versioning setting	A FabricPool bucket has versioning or S3 Object Lock enabled, which are not supported.

Cassandra alerts

Alert name	Description
Cassandra auto-compactor error	The Cassandra auto-compactor has experienced an error.
Cassandra auto-compactor metrics out of date	The metrics that describe the Cassandra auto-compactor are out of date.
Cassandra communication error	The nodes that run the Cassandra service are having trouble communicating with each other.
Cassandra compactions overloaded	The Cassandra compaction process is overloaded.
Cassandra oversize write error	An internal StorageGRID process sent a write request to Cassandra that was too large.
Cassandra repair metrics out of date	The metrics that describe Cassandra repair jobs are out of date.
Cassandra repair progress slow	The progress of Cassandra database repairs is slow.

Alert name	Description
Cassandra repair service not available	The Cassandra repair service is not available.
Cassandra table corruption	Cassandra has detected table corruption. Cassandra automatically restarts if it detects table corruption.

Cloud Storage Pool alerts

Alert name	Description
Cloud Storage Pool connectivity error	The health check for Cloud Storage Pools detected one or more new errors.
IAM Roles Anywhere end-entity certification expiration	IAM Roles Anywhere end-entity certificate is about to expire.

Cross-grid replication alerts

Alert name	Description
Cross-grid replication permanent failure	A cross-grid replication error occurred that requires user intervention to resolve.
Cross-grid replication resources unavailable	Cross-grid replication requests are pending because a resource is unavailable.

DHCP alerts

Alert name	Description
DHCP lease expired	The DHCP lease on a network interface has expired.
DHCP lease expiring soon	The DHCP lease on a network interface is expiring soon.
DHCP server unavailable	The DHCP server is unavailable.

Debug and trace alerts

Alert name	Description
Debug performance impact	When debug mode is enabled, system performance might be negatively impacted.
Trace configuration enabled	When trace configuration is enabled, system performance might be negatively impacted.

Email and AutoSupport alerts

Alert name	Description
AutoSupport message failed to send	The most recent AutoSupport message failed to send.
Domain name resolution failure	The StorageGRID node has been unable to resolve domain names.
Email notification failure	The email notification for an alert could not be sent.
SNMP inform errors	Errors sending SNMP inform notifications to a trap destination.
SSH or console login detected	In the past 24 hours, a user has logged in with Web Console or SSH.

Erasure coding (EC) alerts

Alert name	Description
EC rebalance failure	The EC rebalance procedure has failed or has been stopped.
EC repair failure	A repair job for EC data has failed or has been stopped.
EC repair stalled	A repair job for EC data has stalled.
Erasure-coded fragment verification error	Erasure-coded fragments can no longer be verified. Corrupt fragments might not be repaired.

Expiration of certificates alerts

Alert name	Description
Admin Proxy CA certificate expiration	One or more certificates in the admin proxy server CA bundle is about to expire.
Expiration of client certificate	One or more client certificates are about to expire.
Expiration of global server certificate for S3 and Swift	The global server certificate for S3 and Swift is about to expire.
Expiration of load balancer endpoint certificate	One or more load balancer endpoint certificates are about to expire.
Expiration of server certificate for Management interface	The server certificate used for the management interface is about to expire.

Alert name	Description
External syslog CA certificate expiration	The certificate authority (CA) certificate used to sign the external syslog server certificate is about to expire.
External syslog client certificate expiration	The client certificate for an external syslog server is about to expire.
External syslog server certificate expiration	The server certificate presented by the external syslog server is about to expire.

Grid Network alerts

Alert name	Description
Grid Network MTU mismatch	The MTU setting for the Grid Network interface (eth0) differs significantly across nodes in the grid.

Grid federation alerts

Alert name	Description
Expiration of grid federation certificate	One or more grid federation certificates are about to expire.
Grid federation connection failure	The grid federation connection between the local and remote grid is not working.

High usage or high latency alerts

Alert name	Description
High Java heap use	A high percentage of Java heap space is being used.
High latency for metadata queries	The average time for Cassandra metadata queries is too long.

Identity federation alerts

Alert name	Description
Identity federation synchronization failure	Unable to synchronize federated groups and users from the identity source.
Identity federation synchronization failure for a tenant	Unable to synchronize federated groups and users from the identity source configured by a tenant.

Information lifecycle management (ILM) alerts

Alert name	Description
ILM placement unachievable	A placement instruction in an ILM rule cannot be achieved for certain objects.
ILM scan rate low	The ILM scan rate is set to less than 100 objects/second.

Key management server (KMS) alerts

Alert name	Description
KMS CA certificate expiration	The certificate authority (CA) certificate used to sign the key management server (KMS) certificate is about to expire.
KMS client certificate expiration	The client certificate for a key management server is about to expire
KMS configuration failed to load	The configuration for the key management server exists but failed to load.
KMS connectivity error	An appliance node could not connect to the key management server for its site.
KMS encryption key name not found	The configured key management server does not have an encryption key that matches the name provided.
KMS encryption key rotation failed	All appliance volumes were successfully decrypted, but one or more volumes could not rotate to the latest key.
KMS is not configured	No key management server exists for this site.
KMS key failed to decrypt an appliance volume	One or more volumes on an appliance with node encryption enabled could not be decrypted with the current KMS key.
KMS server certificate expiration	The server certificate used by the key management server (KMS) is about to expire.
KMS server connectivity failure	An appliance node could not connect to one or more servers in the key management server cluster for its site.

Load balancer alerts

Alert name	Description
Elevated zero-request load balancer connections	An elevated percentage of connections to load balancer endpoints disconnected without performing requests.

Local clock offset alerts

Alert name	Description
Local clock large time offset	The offset between local clock and Network Time Protocol (NTP) time is too large.

Low memory or low space alerts

Alert name	Description
Low audit log disk capacity	The space available for audit logs is low. If this condition is not addressed, S3 or Swift operations might fail.
Low available node memory	The amount of RAM available on a node is low.
Low free space for storage pool	The space available for storing object data in the Storage Node is low.
Low installed node memory	The amount of installed memory on a node is low.
Low metadata storage	The space available for storing object metadata is low.
Low metrics disk capacity	The space available for the metrics database is low.
Low object data storage	The space available for storing object data is low.
Low read-only watermark override	The storage volume soft read-only watermark override is less than the minimum optimized watermark for a Storage Node.
Low root disk capacity	The space available on the root disk is low.
Low system data capacity	The space available for /var/local is low. If this condition is not addressed, S3 or Swift operations might fail.
Low tmp directory free space	The space available in the /tmp directory is low.

Node or node network alerts

Alert name	Description
Admin Network receive usage	The receive usage on the Admin Network is high.
Admin Network transmit usage	The transmit usage on the Admin Network is high.
Firewall configuration failure	Failed to apply firewall configuration.

Alert name	Description
Management interface endpoints in fallback mode	All management interface endpoints have been falling back to the default ports for too long.
Node network connectivity error	Errors have occurred while transferring data between nodes.
Node network reception frame error	A high percentage of the network frames received by a node had errors.
Node not in sync with NTP server	The node is not in sync with the network time protocol (NTP) server.
Node not locked with NTP server	The node is not locked to a network time protocol (NTP) server.
Non-appliance node network down	One or more network devices are down or disconnected.
Services appliance link down on Admin Network	The appliance interface to the Admin Network (eth1) is down or disconnected.
Services appliance link down on Admin Network port 1	The Admin Network port 1 on the appliance is down or disconnected.
Services appliance link down on Client Network	The appliance interface to the Client Network (eth2) is down or disconnected.
Services appliance link down on network port 1	Network port 1 on the appliance is down or disconnected.
Services appliance link down on network port 2	Network port 2 on the appliance is down or disconnected.
Services appliance link down on network port 3	Network port 3 on the appliance is down or disconnected.
Services appliance link down on network port 4	Network port 4 on the appliance is down or disconnected.
Storage appliance link down on Admin Network	The appliance interface to the Admin Network (eth1) is down or disconnected.
Storage appliance link down on Admin Network port 1	The Admin Network port 1 on the appliance is down or disconnected.
Storage appliance link down on Client Network	The appliance interface to the Client Network (eth2) is down or disconnected.
Storage appliance link down on network port 1	Network port 1 on the appliance is down or disconnected.

Alert name	Description
Storage appliance link down on network port 2	Network port 2 on the appliance is down or disconnected.
Storage appliance link down on network port 3	Network port 3 on the appliance is down or disconnected.
Storage appliance link down on network port 4	Network port 4 on the appliance is down or disconnected.
Storage Node not in desired storage state	The LDR service on a Storage Node cannot transition to the desired state because of an internal error or volume related issue
TCP connection usage	The number of TCP connections on this node is approaching the maximum number that can be tracked.
Unable to communicate with node	One or more services are unresponsive, or the node cannot be reached.
Unexpected node reboot	A node rebooted unexpectedly within the last 24 hours.

Object alerts

Alert name	Description
Object existence check failed	The object existence check job has failed.
Object existence check stalled	The object existence check job has stalled.
Objects lost	One or more objects have been lost from the grid.
S3 PUT object size too large	A client is attempting a PUT Object operation that exceeds S3 size limits.
Unidentified corrupt object detected	A file was found in replicated object storage that could not be identified as a replicated object.

Platform services alerts

Alert name	Description
Platform Services pending request capacity low	The number of Platform Services pending requests is approaching capacity.
Platform services unavailable	Too few Storage Nodes with the RSM service are running or available at a site.

Storage volume alerts

Alert name	Description
Storage volume needs attention	A storage volume is offline and needs attention.
Storage volume needs to be restored	A storage volume has been recovered and needs to be restored.
Storage volume offline	A storage volume has been offline for more than 5 minutes.
Storage volume remount attempted	A storage volume was offline and triggered an automatic remount. This could indicate a drive issue or filesystem errors.
Volume Restoration failed to start replicated data repair	Replicated data repair for a repaired volume couldn't be started automatically.

StorageGRID services alerts

Alert name	Description
nginx service using backup configuration	The configuration of the nginx service is invalid. The previous configuration is now being used.
nginx-gw service using backup configuration	The configuration of the nginx-gw service is invalid. The previous configuration is now being used.
Reboot required to disable FIPS	The security policy does not require FIPS mode, but the NetApp Cryptographic Security Module is enabled.
Reboot required to enable FIPS	The security policy requires FIPS mode, but the NetApp Cryptographic Security Module is disabled.
SSH service using backup configuration	The configuration of the SSH service is invalid. The previous configuration is now being used.

Tenant alerts

Alert name	Description
Tenant quota usage high	A high percentage of quota space is being used. This rule is disabled by default because it might cause too many notifications.

Commonly used Prometheus metrics

Refer to this list of commonly used Prometheus metrics to better understand conditions in the default alert rules or to construct the conditions for custom alert rules.

You can also [obtain a complete list of all metrics](#).

For details on the syntax of Prometheus queries, see [Querying Prometheus](#).

What are Prometheus metrics?

Prometheus metrics are time series measurements. The Prometheus service on Admin Nodes collects these metrics from the services on all nodes. Metrics are stored on each Admin Node until the space reserved for Prometheus data is full. When the `/var/local/mysql_ibdata/` volume reaches capacity, the oldest metrics are deleted first.

Where are Prometheus metrics used?

The metrics collected by Prometheus are used in several places in the Grid Manager:

- **Nodes page:** The graphs and charts on the tabs available from the Nodes page use the Grafana visualization tool to display the time-series metrics collected by Prometheus. Grafana displays time-series data in graph and chart formats, while Prometheus serves as the backend data source.



- **Alerts:** Alerts are triggered at specific severity levels when alert rule conditions that use Prometheus metrics evaluate as true.
- **Grid Management API:** You can use Prometheus metrics in custom alert rules or with external automation tools to monitor your StorageGRID system. A complete list of Prometheus metrics is available from the Grid Management API. (From the top of the Grid Manager, select the help icon and select **API documentation > metrics**.) While more than a thousand metrics are available, only a relatively small number are required to monitor the most critical StorageGRID operations.



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

- The **SUPPORT > Tools > Diagnostics** page and the **SUPPORT > Tools > Metrics** page: These pages, which are primarily intended for use by technical support, provide several tools and charts that use the values of Prometheus metrics.



Some features and menu items within the Metrics page are intentionally non-functional and are subject to change.

List of most common metrics

The following list contains the most commonly used Prometheus metrics.



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

alertmanager_notifications_failed_total

The total number of failed alert notifications.

node_filesystem_avail_bytes

The amount of file system space available to non-root users in bytes.

node_memory_MemAvailable_bytes

Memory information field MemAvailable_bytes.

node_network_carrier

Carrier value of `/sys/class/net/iface`.

node_network_receive_errs_total

Network device statistic `receive_errs`.

node_network_transmit_errs_total

Network device statistic `transmit_errs`.

storagegrid_administratively_down

The node is not connected to the grid for an expected reason. For example, the node, or services on the node, has been gracefully shut down, the node is rebooting, or the software is being upgraded.

storagegrid_appliance_compute_controller_hardware_status

The status of the compute controller hardware in an appliance.

storagegrid_appliance_failed_disks

For the storage controller in an appliance, the number of drives that aren't optimal.

storagegrid_appliance_storage_controller_hardware_status

The overall status of the storage controller hardware in an appliance.

storagegrid_content_buckets_and_containers

The total number of S3 buckets and Swift containers known by this Storage Node.

storagegrid_content_objects

The total number of S3 and Swift data objects known by this Storage Node. Count is valid only for data objects created by client applications that interface with the system through S3.

storagegrid_content_objects_lost

The total number of objects this service detects as missing from the StorageGRID system. Action should be taken to determine the cause of the loss and if recovery is possible.

[Troubleshoot lost and missing object data](#)

storagegrid_http_sessions_incoming_attempted

The total number of HTTP sessions that have been attempted to a Storage Node.

storagegrid_http_sessions_incoming_currently_established

The number of HTTP sessions that are currently active (open) on the Storage Node.

storagegrid_http_sessions_incoming_failed

The total number of HTTP sessions that failed to complete successfully, either due to a malformed HTTP request or a failure while processing an operation.

storagegrid_http_sessions_incoming_successful

The total number of HTTP sessions that have completed successfully.

storagegrid_ilm_awaiting_background_objects

The total number of objects on this node awaiting ILM evaluation from the scan.

storagegrid_ilm_awaiting_client_evaluation_objects_per_second

The current rate at which objects are evaluated against the ILM policy on this node.

storagegrid_ilm_awaiting_client_objects

The total number of objects on this node awaiting ILM evaluation from client operations (for example, ingest).

storagegrid_ilm_awaiting_total_objects

The total number of objects awaiting ILM evaluation.

storagegrid_ilm_scan_objects_per_second

The rate at which objects owned by this node are scanned and queued for ILM.

storagegrid_ilm_scan_period_estimated_minutes

The estimated time to complete a full ILM scan on this node.

Note: A full scan does not guarantee that ILM has been applied to all objects owned by this node.

storagegrid_load_balancer_endpoint_cert_expiry_time

The expiration time of the load balancer endpoint certificate in seconds since the epoch.

storagegrid_metadata_queries_average_latency_milliseconds

The average time required to run a query against the metadata store through this service.

storagegrid_network_received_bytes

The total amount of data received since installation.

storagegrid_network_transmitted_bytes

The total amount of data sent since installation.

storagegrid_node_cpu_utilization_percentage

The percentage of available CPU time currently being used by this service. Indicates how busy the service is. The amount of available CPU time depends on the number of CPUs for the server.

storagegrid_ntp_chosen_time_source_offset_milliseconds

Systematic offset of time provided by a chosen time source. Offset is introduced when the delay to reach a time source is not equal to the time required for the time source to reach the NTP client.

storagegrid_ntp_locked

The node is not locked to a Network Time Protocol (NTP) server.

storagegrid_s3_data_transfers_bytes_ingested

The total amount of data ingested from S3 clients to this Storage Node since the attribute was last reset.

storagegrid_s3_data_transfers_bytes_retrieved

The total amount of data retrieved by S3 clients from this Storage Node since the attribute was last reset.

storagegrid_s3_operations_failed

The total number of failed S3 operations (HTTP status codes 4xx and 5xx), excluding those caused by S3 authorization failure.

storagegrid_s3_operations_successful

The total number of successful S3 operations (HTTP status code 2xx).

storagegrid_s3_operations_unauthorized

The total number of failed S3 operations that are the result of an authorization failure.

storagegrid_servercertificate_management_interface_cert_expiry_days

The number of days before the Management Interface certificate expires.

storagegrid_servercertificate_storage_api_endpoints_cert_expiry_days

The number of days before the Object Storage API certificate expires.

storagegrid_service_cpu_seconds

The cumulative amount of time that the CPU has been used by this service since installation.

storagegrid_service_memory_usage_bytes

The amount of memory (RAM) currently in use by this service. This value is identical to that displayed by the Linux top utility as RES.

storagegrid_service_network_received_bytes

The total amount of data received by this service since installation.

storagegrid_service_network_transmitted_bytes

The total amount of data sent by this service.

storagegrid_service_restarts

The total number of times the service has been restarted.

storagegrid_service_runtime_seconds

The total amount of time that the service has been running since installation.

storagegrid_service_uptime_seconds

The total amount of time the service has been running since it was last restarted.

storagegrid_storage_state_current

The current state of the storage services. Attribute values are:

- 10 = Offline
- 15 = Maintenance
- 20 = Read-only
- 30 = Online

storagegrid_storage_status

The current status of the storage services. Attribute values are:

- 0 = No Errors
- 10 = In Transition
- 20 = Insufficient Free Space
- 30 = Volume(s) Unavailable
- 40 = Error

storagegrid_storage_utilization_data_bytes

An estimate of the total size of replicated and erasure-coded object data on the Storage Node.

storagegrid_storage_utilization_metadata_allowed_bytes

The total space on volume 0 of each Storage Node that is allowed for object metadata. This value is always less than the actual space reserved for metadata on a node, because a portion of the reserved space is required for essential database operations (such as compaction and repair) and future hardware and software upgrades. The allowed space for object metadata controls overall object capacity.

storagegrid_storage_utilization_metadata_bytes

The amount of object metadata on storage volume 0, in bytes.

storagegrid_storage_utilization_total_space_bytes

The total amount of storage space allocated to all object stores.

storagegrid_storage_utilization_usable_space_bytes

The total amount of object storage space remaining. Calculated by adding together the amount of available space for all object stores on the Storage Node.

storagegrid_swift_data_transfers_bytes_ingested

The total amount of data ingested from Swift clients to this Storage Node since the attribute was last reset.

storagegrid_swift_data_transfers_bytes_retrieved

The total amount of data retrieved by Swift clients from this Storage Node since the attribute was last reset.

storagegrid_swift_operations_failed

The total number of failed Swift operations (HTTP status codes 4xx and 5xx), excluding those caused by Swift authorization failure.

storagegrid_swift_operations_successful

The total number of successful Swift operations (HTTP status code 2xx).

storagegrid_swift_operations_unauthorized

The total number of failed Swift operations that are the result of an authorization failure (HTTP status codes 401, 403, 405).

storagegrid_tenant_usage_data_bytes

The logical size of all objects for the tenant.

storagegrid_tenant_usage_object_count

The number of objects for the tenant.

storagegrid_tenant_usage_quota_bytes

The maximum amount of logical space available for the tenant's objects. If a quota metric is not provided, an unlimited amount of space is available.

Get a list of all metrics

To obtain the complete list of metrics, use the Grid Management API.

1. From the top of the Grid Manager, select the help icon and select **API documentation**.
2. Locate the **metrics** operations.
3. Execute the `GET /grid/metric-names` operation.
4. Download the results.

Log files reference

Log files reference

StorageGRID provides logs that are used to capture events, diagnostic messages, and error conditions. You might be asked to collect log files and forward them to technical support to assist with troubleshooting.

The logs are categorized as follows:

- [StorageGRID software logs](#)
- [Deployment and maintenance logs](#)
- [About the bycast.log](#)



The details provided for each log type are for reference only. The logs are intended for advanced troubleshooting by technical support. Advanced techniques that involve reconstructing the problem history using the audit logs and the application log files are beyond the scope of these instructions.

Access the logs

To access the logs, you can [collect log files and system data](#) from one or more nodes as a single log file archive. Or, if the primary Admin Node is unavailable or unable to reach a specific node, you can access individual log files for each grid node as follows:

1. Enter the following command: `ssh admin@grid_node_IP`

2. Enter the password listed in the `Passwords.txt` file.
3. Enter the following command to switch to root: `su -`
4. Enter the password listed in the `Passwords.txt` file.

Export logs to the syslog server

Exporting the logs to the syslog server provides these capabilities:

- Receive a list of all Grid Manager and Tenant Manager requests, in addition to S3 and Swift requests.
- Better visibility into S3 requests that return errors, without the performance impact caused by audit logging methods.
- Access to HTTP-layer requests and error codes that are easy to parse.
- Better visibility into requests that were blocked by traffic classifiers at the load balancer.

To export the logs, refer to [Configure audit messages and log destinations](#).

Log file categories

The StorageGRID log file archive contains the logs described for each category and additional files that contain metrics and debug command output.

Archive location	Description
audit	Audit messages generated during normal system operation.
base-os-logs	Base operating system information, including StorageGRID image versions.
bundles	Global configuration information (bundles).
cassandra	Cassandra database information and Reaper repair logs.
ec	VCSs information about the current node and EC group information by profile ID.
grid	General grid logs including debug (<code>broadcast.log</code>) and <code>servermanager</code> logs.
grid.json	Grid configuration file shared across all nodes. Additionally, <code>node.json</code> is specific to the current node.
hagroups	High availability groups metrics and logs.
install	<code>Gdu-server</code> and install logs.
Lambda-arbitrator	Logs related to the S3 Select proxy request.
lumberjack.log	Debug messages related to log collection.

Archive location	Description
Metrics	Service logs for Grafana, Jaeger, node exporter, and Prometheus.
miscd	Miscd access and error logs.
mysql	The mariaDB database configuration and related logs.
net	Logs generated by networking-related scripts and the Dynip service.
nginx	Load balancer and grid federation configuration files and logs. Also includes Grid Manager and Tenant Manager traffic logs.
nginx-gw	• <code>access.log</code>: Grid Manager and Tenant manager request log messages. ◦ These messages are prefixed with <code>mgmt</code>: when exported using syslog. ◦ The format of these log messages is <code>[\$time_iso8601] \$remote_addr \$status \$bytes_sent \$request_length \$request_time "\$endpointId" "\$request" "\$http_host" "\$http_user_agent" "\$http_referer"</code> • <code>cgr-access.log.gz</code>: Inbound cross-grid replication requests. ◦ These messages are prefixed with <code>cgr</code>: when exported using syslog. ◦ The format of these log messages is <code>[\$time_iso8601] \$remote_addr \$status \$bytes_sent \$request_length \$request_time "\$endpointId" "\$upstream_addr" "\$request" "\$http_host"</code> • <code>endpoint-access.log.gz</code>: S3 and Swift requests to load balancer endpoints. ◦ These messages are prefixed with <code>endpoint</code>: when exported using syslog. ◦ The format of these log messages is <code>[\$time_iso8601] \$remote_addr \$status \$bytes_sent \$request_length \$request_time "\$endpointId" "\$upstream_addr" "\$request" "\$http_host"</code> • <code>nginx-gw-dns-check.log</code>: Related to the new DNS check alert.
ntp	NTP configuration file and logs.
Orphaned objects	Logs pertaining to orphaned objects.
os	Node and grid state file, including services <code>pid</code> .
other	Log files under <code>/var/local/log</code> that aren't collected in other folders.

Archive location	Description
perf	Performance information for CPU, networking, and disk I/O.
prometheus-data	Current Prometheus metrics, if the log collection includes Prometheus data.
provisioning	Logs related to grid provisioning process.
raft	Logs from Raft cluster used in platform services.
ssh	Logs related to SSH configuration and service.
snmp	SNMP agent configuration used for sending SNMP notifications.
sockets-data	Sockets data for network debug.
system-commands.txt	Output of StorageGRID container commands. Contains system information, such as networking and disk usage.
synchronize-recovery-package	Related to maintaining consistency of the latest Recovery Package across all Admin Nodes and Storage Nodes that host the ADC service.

StorageGRID software logs

You can use StorageGRID logs to troubleshoot issues.



If you want to send your logs to an external syslog server or change the destination of audit information such as the `bycast.log` and `nms.log`, see [Configure audit messages and log destinations](#).

General StorageGRID logs

File name	Notes	Found on
<code>/var/local/log/bycast.log</code>	The primary StorageGRID troubleshooting file. Select SUPPORT > Tools > Grid topology . Then select Site > Node > SSM > Events .	All nodes
<code>/var/local/log/bycast-err.log</code>	Contains a subset of <code>bycast.log</code> (messages with severity ERROR and CRITICAL). CRITICAL messages are also displayed in the system. Select SUPPORT > Tools > Grid topology . Then select Site > Node > SSM > Events .	All nodes

File name	Notes	Found on
/var/local/core/	Contains any core dump files created if the program terminates abnormally. Possible causes include assertion failures, violations, or thread timeouts. Note: The file <code>`/var/local/core/kexec_cmd`</code> usually exists on appliance nodes and does not indicate an error.	All nodes

Cipher-related logs

File name	Notes	Found on
/var/local/log/ssh-config-generation.log	Contains logs related to generating SSH configurations and reloading SSH services.	All nodes
/var/local/log/nginx/config-generation.log	Contains logs related to generating nginx configurations and reloading nginx services.	All nodes
/var/local/log/nginx-gw/config-generation.log	Contains logs related to generating nginx-gw configurations (and reloading nginx-gw services).	Admin and Gateway Nodes
/var/local/log/update-cipher-configurations.log	Contains logs related to configuring TLS and SSH policies.	All nodes

Grid federation logs

File name	Notes	Found on
/var/local/log/update_grid_federation_config.log	Contains logs related to generating nginx and nginx-gw configurations for grid federation connections.	All nodes

NMS logs

File name	Notes	Found on
/var/local/log/nms.log	• Captures notifications from the Grid Manager and the Tenant Manager. • Captures events related to the operation of the NMS service. For example, email notifications and configuration changes. • Contains XML bundle updates resulting from configuration changes made in the system. • Contains error messages related to the attribute downsampling done once a day. • Contains Java web server error messages, for example, page generation errors and HTTP Status 500 errors.	Admin Nodes
/var/local/log/nms.errlog	Contains error messages related to MySQL database upgrades. Contains the Standard Error (stderr) stream of the corresponding services. There is one log file per service. These files are generally empty unless there are problems with the service.	Admin Nodes
/var/local/log/nms.requestlog	Contains information about outgoing connections from the Management API to internal StorageGRID services.	Admin Nodes

Server Manager logs

File name	Notes	Found on
/var/local/log/servermanager.log	Log file for the Server Manager application running on the server.	All nodes
/var/local/log/GridstatBackend.errlog	Log file for the Server Manager GUI backend application.	All nodes
/var/local/log/gridstat.errlog	Log file for the Server Manager GUI.	All nodes

StorageGRID services logs

File name	Notes	Found on
/var/local/log/acct.errlog		Storage Nodes running the ADC service
/var/local/log/adc.errlog	Contains the Standard Error (stderr) stream of the corresponding services. There is one log file per service. These files are generally empty unless there are problems with the service.	Storage Nodes running the ADC service
/var/local/log/ams.errlog		Admin Nodes
/var/local/log/cassandra/system.log	Information for the metadata store (Cassandra database) that can be used if problems occur when adding new Storage Nodes, or if the nodetool repair task stalls.	Storage Nodes
/var/local/log/cassandra-reaper.log	Information for the Cassandra Reaper service, which performs repairs of the data in the Cassandra database.	Storage Nodes
/var/local/log/cassandra-reaper.errlog	Error information for the Cassandra Reaper service.	Storage Nodes
/var/local/log/chunk.errlog		Storage Nodes
/var/local/log/cmn.errlog		Admin Nodes
/var/local/log/cms.errlog	This log file might be present on systems that have been upgraded from an older version of StorageGRID. It contains legacy information.	Storage Nodes
/var/local/log/dds.errlog		Storage Nodes
/var/local/log/dmv.errlog		Storage Nodes
/var/local/log/dynip*	Contains logs related to the dynip service, which monitors the grid for dynamic IP changes and updates local configuration.	All nodes
/var/local/log/grafana.log	The log associated with the Grafana service, which is used for metrics visualization in the Grid Manager.	Admin Nodes

File name	Notes	Found on
/var/local/log/hagroups.log	The log associated with high availability groups.	Admin Nodes and Gateway Nodes
/var/local/log/hagroups_events.log	Tracks state changes, such as transition from BACKUP to MASTER or FAULT.	Admin Nodes and Gateway Nodes
/var/local/log/idnt.errlog		Storage Nodes running the ADC service
/var/local/log/jaeger.log	The log associated with the jaeger service, which is used for trace collection.	All nodes
/var/local/log/kstn.errlog		Storage Nodes running the ADC service
/var/local/log/lambda*	Contains logs for the S3 Select service.	Admin and Gateway Nodes Only certain Admin and Gateway Nodes contain this log. See the S3 Select requirements and limitations for Admin and Gateway Nodes .
/var/local/log/ldr.errlog		Storage Nodes
/var/local/log/miscd/*.log	Contains logs for the MISCD service (Information Service Control Daemon), which provides an interface for querying and managing services on other nodes and for managing environmental configurations on the node such as querying the state of services running on other nodes.	All nodes
/var/local/log/nginx/*.log	Contains logs for the nginx service, which acts as an authentication and secure communication mechanism for various grid services (such as Prometheus and Dynip) to be able to talk to services on other nodes over HTTPS APIs.	All nodes

File name	Notes	Found on
/var/local/log/nginx-gw/*.log	Contains general logs related to the nginx-gw service, including error logs, and logs for the restricted admin ports on Admin Nodes.	Admin Nodes and Gateway Nodes
/var/local/log/nginx-gw/cgr-access.log.gz	Contains access logs related to cross-grid replication traffic.	Admin Nodes, Gateway Nodes, or both, based on the grid federation configuration. Only found on the destination grid for cross-grid replication.
/var/local/log/nginx-gw/endpoint-access.log.gz	Contains access logs for the Load Balancer service, which provides load balancing of S3 traffic from clients to Storage Nodes.	Admin Nodes and Gateway Nodes
/var/local/log/persistence*	Contains logs for the Persistence service, which manages files on the root disk that need to persist across a reboot.	All nodes
/var/local/log/prometheus.log	For all nodes, contains the node exporter service log and the ade-exporter metrics service log. For Admin Nodes, also contains logs for the Prometheus and Alert Manager services.	All nodes
/var/local/log/raft.log	Contains the output of the library used by the RSM service for the Raft protocol.	Storage Nodes with RSM service
/var/local/log/rms.errlog	Contains logs for the Replicated State Machine Service (RSM) service, which is used for S3 platform services.	Storage Nodes with RSM service
/var/local/log/ssm.errlog		All nodes
/var/local/log/update-s3vs-domains.log	Contains logs related to processing updates for the S3 virtual hosted domain names configuration. See the instructions for implementing S3 client applications.	Admin and Gateway Nodes
/var/local/log/update-snmp-firewall.*	Contain logs related to the firewall ports being managed for SNMP.	All nodes

File name	Notes	Found on
/var/local/log/update-sysl.log	Contains logs related to changes made to the system syslog configuration.	All nodes
/var/local/log/update-traffic-classes.log	Contains logs related to changes to the traffic classifiers configuration.	Admin and Gateway Nodes
/var/local/log/update-utcn.log	Contains logs related to Untrusted Client Network mode on this node.	All nodes

Related information

- [About the bycast.log](#)
- [Use S3 REST API](#)

Deployment and maintenance logs

You can use the deployment and maintenance logs to troubleshoot issues.

File name	Notes	Found on
/var/local/log/install.log	Created during software installation. Contains a record of the installation events.	All nodes
/var/local/log/expansion-progress.log	Created during expansion operations. Contains a record of the expansion events.	Storage Nodes
/var/local/log/pa-move.log	Created while running the <code>pa-move.sh</code> script.	Primary Admin Node
/var/local/log/pa-move-new_pa.log	Created while running the <code>pa-move.sh</code> script.	Primary Admin Node
/var/local/log/pa-move-old_pa.log	Created while running the <code>pa-move.sh</code> script.	Primary Admin Node
/var/local/log/gdu-server.log	Created by the GDU service. Contains events related to provisioning and maintenance procedures managed by the primary Admin Node.	Primary Admin Node
/var/local/log/send_admin_hw.log	Created during installation. Contains debugging information related to a node's communications with the primary Admin Node.	All nodes
/var/local/log/upgrade.log	Created during software upgrade. Contains a record of the software update events.	All nodes

About the bycast.log

The file `/var/local/log/bycast.log` is the primary troubleshooting file for the StorageGRID software. There is a `bycast.log` file for every grid node. The file contains messages specific to that grid node.

The file `/var/local/log/bycast-err.log` is a subset of `bycast.log`. It contains messages of severity ERROR and CRITICAL.

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

File rotation for bycast.log

When the `bycast.log` file reaches 1 GB, the existing file is saved, and a new log file is started.

The saved file is renamed `bycast.log.1`, and the new file is named `bycast.log`. When the new `bycast.log` reaches 1 GB, `bycast.log.1` is renamed and compressed to become `bycast.log.2.gz`, and `bycast.log` is renamed `bycast.log.1`.

The rotation limit for `bycast.log` is 21 files. When the 22nd version of the `bycast.log` file is created, the oldest file is deleted.

The rotation limit for `bycast-err.log` is seven files.



If a log file has been compressed, you must not uncompress it to the same location in which it was written. Uncompressing the file to the same location can interfere with the log rotation scripts.

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

Related information

[Collect log files and system data](#)

Messages in bycast.log

Messages in `bycast.log` are written by the ADE (Asynchronous Distributed Environment). ADE is the runtime environment used by each grid node's services.

Example ADE message:

```
May 15 14:07:11 um-sec-rg1-agn3 ADE: |12455685      0357819531
SVMR EVHR 2019-05-05T27T17:10:29.784677| ERROR 0906 SVMR: Health
check on volume 3 has failed with reason 'TOUT'
```

ADE messages contain the following information:

Message segment	Value in example
Node ID	12455685
ADE process ID	0357819531
Module name	SVMR
Message identifier	EVHR
UTC system time	2019-05-05T27T17:10:29.784677 (YYYY-MM-DDTHH:MM:SS.uuuuuu)
Severity level	ERROR
Internal tracking number	0906
Message	SVMR: Health check on volume 3 has failed with reason 'TOUT'

Message severities in bycast.log

The messages in `bycast.log` are assigned severity levels.

For example:

- **NOTICE** — An event that should be recorded has occurred. Most log messages are at this level.
- **WARNING** — An unexpected condition has occurred.
- **ERROR** — A major error has occurred that will impact operations.
- **CRITICAL** — An abnormal condition has occurred that has stopped normal operations. You should address the underlying condition immediately.

Error codes in bycast.log

Most of the error messages in `bycast.log` contain error codes.

The following table lists common non-numerical codes in `bycast.log`. The exact meaning of a non-numerical code depends on the context in which it is reported.

Error code	Meaning
SUCS	No error
GERR	Unknown
CANC	Canceled
ABRT	Aborted

Error code	Meaning
TOUT	Timeout
INVL	Invalid
NFND	Not found
VERS	Version
CONF	Configuration
FAIL	Failed
ICPL	Incomplete
DONE	Done
SUNV	Service unavailable

The following table lists the numerical error codes in `broadcast.log`.

Error number	Error code	Meaning
001	EPERM	Operation not permitted
002	ENOENT	No such file or directory
003	ESRCH	No such process
004	EINTR	Interrupted system call
005	EIO	I/O error
006	ENXIO	No such device or address
007	E2BIG	Argument list too long
008	ENOEXEC	Exec format error
009	EBADF	Bad file number
010	ECHILD	No child processes
011	EAGAIN	Try again

Error number	Error code	Meaning
012	ENOMEM	Out of memory
013	EACCES	Permission denied
014	EFAULT	Bad address
015	ENOTBLK	Block device required
016	EBUSY	Device or resource busy
017	EEXIST	File exists
018	EXDEV	Cross-device link
019	ENODEV	No such device
020	ENOTDIR	Not a directory
021	EISDIR	Is a directory
022	EINVAL	Invalid argument
023	ENFILE	File table overflow
024	EMFILE	Too many open files
025	ENOTTY	Not a typewriter
026	ETXTBSY	Text file busy
027	EFBIG	File too large
028	ENOSPC	No space left on device
029	ESPIPE	Illegal seek
030	EROFS	Read-only file system
031	EMLINK	Too many links
032	EPIPE	Broken pipe
033	EDOM	Math argument out of domain of func

Error number	Error code	Meaning
034	ERANGE	Math result not representable
035	EDEADLK	Resource deadlock would occur
036	ENAMETOOLONG	File name too long
037	ENOLCK	No record locks available
038	ENOSYS	Function not implemented
039	ENOTEMPTY	Directory not empty
040	ELOOP	Too many symbolic links encountered
041		
042	ENOMSG	No message of desired type
043	EIDRM	Identifier removed
044	ECHRNG	Channel number out of range
045	EL2NSYNC	Level 2 not synchronized
046	EL3HLT	Level 3 halted
047	EL3RST	Level 3 reset
048	ELNRNG	Link number out of range
049	EUNATCH	Protocol driver not attached
050	ENOCSI	No CSI structure available
051	EL2HLT	Level 2 halted
052	EBADE	Invalid exchange
053	EBADR	Invalid request descriptor
054	EXFULL	Exchange full
055	ENOANO	No anode

Error number	Error code	Meaning
056	EBADRQC	Invalid request code
057	EBADSLT	Invalid slot
058		
059	EBFONT	Bad font file format
060	ENOSTR	Device not a stream
061	ENODATA	No data available
062	ETIME	Timer expired
063	ENOSR	Out of streams resources
064	ENONET	Machine is not on the network
065	ENOPKG	Package not installed
066	EREMOTE	Object is remote
067	ENOLINK	Link has been severed
068	EADV	Advertise error
069	ESRMNT	Srmount error
070	ECOMM	Communication error on send
071	EPROTO	Protocol error
072	EMULTIHOP	Multihop attempted
073	EDOTDOT	RFS specific error
074	EBADMSG	Not a data message
075	EOVERFLOW	Value too large for defined data type
076	ENOTUNIQ	Name not unique on network
077	EBADFD	File descriptor in bad state

Error number	Error code	Meaning
078	EREMCHG	Remote address changed
079	ELIBACC	Can't access a needed shared library
080	ELIBBAD	Accessing a corrupted shared library
081	ELIBSCN	
082	ELIBMAX	Attempting to link in too many shared libraries
083	ELIBEXEC	Can't exec a shared library directly
084	EILSEQ	Illegal byte sequence
085	ERESTART	Interrupted system call should be restarted
086	ESTRPIPE	Streams pipe error
087	EUSERS	Too many users
088	ENOTSOCK	Socket operation on non-socket
089	EDESTADDRREQ	Destination address required
090	EMSGSIZE	Message too long
091	EPROTOTYPE	Protocol wrong type for socket
092	ENOPROTOOPT	Protocol not available
093	EPROTONOSUPPORT	Protocol not supported
094	ESOCKTNOSUPPORT	Socket type not supported
095	EOPNOTSUPP	Operation not supported on transport endpoint
096	EPFNOSUPPORT	Protocol family not supported
097	EAFNOSUPPORT	Address family not supported by protocol
098	EADDRINUSE	Address already in use
099	EADDRNOTAVAIL	Can't assign requested address

Error number	Error code	Meaning
100	ENETDOWN	Network is down
101	ENETUNREACH	Network is unreachable
102	ENETRESET	Network dropped connection because of reset
103	ECONNABORTED	Software caused connection to terminate
104	ECONNRESET	Connection reset by peer
105	ENOBUFS	No buffer space available
106	EISCONN	Transport endpoint is already connected
107	ENOTCONN	Transport endpoint is not connected
108	ESHUTDOWN	Can't send after transport endpoint shutdown
109	ETOOMANYREFS	Too many references: can't splice
110	ETIMEDOUT	Connection timed out
111	ECONNREFUSED	Connection refused
112	EHOSTDOWN	Host is down
113	EHOSTUNREACH	No route to host
114	EALREADY	Operation already in progress
115	EINPROGRESS	Operation now in progress
116		
117	EUCLEAN	Structure needs cleaning
118	ENOTNAM	Not a XENIX named type file
119	ENAVAIL	No XENIX semaphores available
120	EISNAM	Is a named type file
121	EREMOTEIO	Remote I/O error

Error number	Error code	Meaning
122	EDQUOT	Quota exceeded
123	ENOMEDIUM	No medium found
124	EMEDIUMTYPE	Wrong medium type
125	ECANCELED	Operation Canceled
126	ENOKEY	Required key not available
127	EKEYEXPIRED	Key has expired
128	EKEYREVOKED	Key has been revoked
129	EKEYREJECTED	Key was rejected by service
130	EOWNERDEAD	For robust mutexes: Owner died
131	ENOTRECOVERABLE	For robust mutexes: State not recoverable

Configure audit message and log destinations

Considerations for using an external syslog server

An external syslog server is a server outside of StorageGRID you can use to collect system audit information in a single location. Using an external syslog server enables you to reduce network traffic on your Admin Nodes and manage the information more efficiently. For StorageGRID, the outbound syslog message packet format is compliant with RFC 3164.

The types of audit information you can send to the external syslog server include:

- Audit logs containing the audit messages generated during normal system operation
- Security-related events such as logins and escalations to root
- Application logs that might be requested if it is necessary to open a support case to troubleshoot an issue you have encountered

When to use an external syslog server

An external syslog server is especially useful if you have a large grid, use multiple types of S3 applications, or want to retain all audit data. Sending audit information to an external syslog server enables you to:

- Collect and manage audit information such as audit messages, application logs, and security events more efficiently.
- Reduce network traffic on your Admin Nodes because audit information is transferred directly from the

various Storage Nodes to the external syslog server, without having to go through an Admin Node.



When logs are sent to an external syslog server, single logs greater than 8,192 bytes are truncated at the end of the message to conform with common limitations in external syslog server implementations.



To maximize the options for full data recovery in the event of a failure of the external syslog server, up to 20 GB of local logs of audit records (`localaudit.log`) are maintained on each node.

How to configure an external syslog server

To learn how to configure an external syslog server, see [Configure audit messages and external syslog server](#).

If you plan to configure use the TLS or RELP/TLS protocol, you must have the following certificates:

- **Server CA certificates:** One or more trusted CA certificates for verifying the external syslog server in PEM encoding. If omitted, the default Grid CA certificate will be used.
- **Client certificate:** The client certificate for authentication to the external syslog server in PEM encoding.
- **Client private key:** Private key for the client certificate in PEM encoding.



If you use a client certificate you must also use a client private key. If you provide an encrypted private key, you must also provide the passphrase. There is no significant security benefit from using an encrypted private key because the key and passphrase must be stored; using an unencrypted private key, if available, is recommended for simplicity.

How to estimate the size of the external syslog server

Normally, your grid is sized to achieve a required throughput, defined in terms of S3 operations per second or bytes per second. For example, you might have a requirement that your grid handle 1,000 S3 operations per second, or 2,000 MB per second, of object ingests and retrievals. You should size your external syslog server according to your grid's data requirements.

This section provides some heuristic formulas that help you estimate the rate and average size of log messages of various types that your external syslog server needs to be capable of handling, expressed in terms of the known or desired performance characteristics of the grid (S3 operations per second).

Use S3 operations per second in estimation formulas

If your grid was sized for a throughput expressed in bytes per second, you must convert this sizing into S3 operations per second to use the estimation formulas. To convert grid throughput, you must first determine your average object size, which you can do using the information in existing audit logs and metrics (if any), or by using your knowledge of the applications that will use StorageGRID. For example, if your grid was sized to achieve a throughput of 2,000 MB/second, and your average object size is 2 MB, then your grid was sized to be able to handle 1,000 S3 operations per second (2,000 MB / 2 MB).



The formulas for external syslog server sizing in the following sections provide common-case estimates (rather than worst-case estimates). Depending on your configuration and workload, you might see a higher or lower rate of syslog messages or volume of syslog data than the formulas predict. The formulas are meant to be used as guidelines only.

Estimation formulas for audit logs

If you have no information about your S3 workload other than number of S3 operations per second your grid is expected to support, you can estimate the volume of audit logs your external syslog server will need to handle using the following formulas, under the assumption that you leave the Audit Levels set to the default values (all categories set to Normal, except Storage, which is set to Error):

```
Audit Log Rate = 2 x S3 Operations Rate
Audit Log Average Size = 800 bytes
```

For example, if your grid is sized for 1,000 S3 operations per second, your external syslog server should be sized to support 2,000 syslog messages per second and should be able to receive (and typically store) audit log data at a rate of 1.6 MB per second.

If you know more about your workload, more accurate estimations are possible. For audit logs, the most important additional variables are the percentage of S3 operations that are PUTs (vs. GETS), and the average size, in bytes, of the following S3 fields (4-character abbreviations used in the table are audit log field names):

Code	Field	Description
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
S3BK	S3 bucket	The S3 bucket name.
S3KY	S3 key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.

Let's use P to represent the percentage of S3 operations that are PUTs, where $0 \leq P \leq 1$ (so for a 100% PUT workload, $P = 1$, and for a 100% GET workload, $P = 0$).

Let's use K to represent the average size of the sum of the S3 account names, S3 bucket, and S3 key. Suppose the S3 account name is always my-s3-account (13 bytes), buckets have fixed-length names like /my/application/bucket-12345 (28 bytes), and objects have fixed-length keys like 5733a5d7-f069-41ef-8fbd-13247494c69c (36 bytes). Then the value of K is 90 (13+13+28+36).

If you can determine values for P and K , you can estimate the volume of audit logs your external syslog server will need to handle using the following formulas, under the assumption that you leave the Audit Levels set to the defaults (all categories set to Normal, except Storage, which is set to Error):

```
Audit Log Rate = ((2 x P) + (1 - P)) x S3 Operations Rate
Audit Log Average Size = (570 + K) bytes
```

For example, if your grid is sized for 1,000 S3 operations per second, your workload is 50% PUTs, and your S3 account names, bucket names, and object names average 90 bytes, your external syslog server should be sized to support 1,500 syslog messages per second and should be able to receive (and typically store) audit log data at a rate of approximately 1 MB per second.

Estimation formulas for non-default audit levels

The formulas provided for audit logs assume the use of default audit level settings (all categories set to Normal, except Storage, which is set to Error). Detailed formulas for estimating the rate and average size of audit messages for non-default audit level settings aren't available. However, the following table can be used to make a rough estimate of the rate; you can use the average size formula provided for audit logs, but be aware that it is likely to result in an over-estimate because the "extra" audit messages are, on average, smaller than the default audit messages.

Condition	Formula
Replication: Audit levels all set to Debug or Normal	Audit log rate = 8 x S3 Operations Rate
Erasure coding: audit levels all set to Debug or Normal	Use same formula as for default settings

Estimation formulas for security events

Security events aren't correlated with S3 operations and typically produce a negligible volume of logs and data. For these reasons, no estimation formulas are provided.

Estimation formulas for application logs

If you have no information about your S3 workload other than the number of S3 operations per second your grid is expected to support, you can estimate the volume of applications logs your external syslog server will need to handle using the following formulas:

```
Application Log Rate = 3.3 x S3 Operations Rate
Application Log Average Size = 350 bytes
```

So, for example, if your grid is sized for 1,000 S3 operations per second, your external syslog server should be sized to support 3,300 application logs per second and be able to receive (and store) application log data at a rate of about 1.2 MB per second.

If you know more about your workload, more accurate estimations are possible. For application logs, the most important additional variables are the data protection strategy (replication vs. erasure coding), the percentage of S3 operations that are PUTs (vs. GETs/other), and the average size, in bytes, of the following S3 fields (4-character abbreviations used in table are audit log field names):

Code	Field	Description
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.

Code	Field	Description
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
S3BK	S3 bucket	The S3 bucket name.
S3KY	S3 key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.

Example sizing estimations

This section explains example cases of how to use the estimation formulas for grids with the following methods of data protection:

- Replication
- Erasure coding

If you use replication for data protection

Let P represent the percentage of S3 operations that are PUTs, where $0 \leq P \leq 1$ (so for a 100% PUT workload, $P = 1$, and for a 100% GET workload, $P = 0$).

Let K represent the average size of the sum of the S3 account names, S3 bucket, and S3 key. Suppose the S3 account name is always my-s3-account (13 bytes), buckets have fixed-length names like /my/application/bucket-12345 (28 bytes), and objects have fixed-length keys like 5733a5d7-f069-41ef-8fbd-13247494c69c (36 bytes). Then K has a value of 90 (13+13+28+36).

If you can determine values for P and K , you can estimate the volume of application logs your external syslog server will have to be able to handle using the following formulas.

```
Application Log Rate = ((1.1 x P) + (2.5 x (1 - P))) x S3 Operations Rate
Application Log Average Size = (P x (220 + K)) + ((1 - P) x (240 + (0.2 x K))) Bytes
```

So, for example, if your grid is sized for 1,000 S3 operations per second, your workload is 50% PUTs, and your S3 account names, bucket names, and object names average 90 bytes, your external syslog server should be sized to support 1800 application logs per second, and will be receiving (and typically storing) application data at a rate of 0.5 MB per second.

If you use erasure coding for data protection

Let P represent the percentage of S3 operations that are PUTs, where $0 \leq P \leq 1$ (so for a 100% PUT workload, $P = 1$, and for a 100% GET workload, $P = 0$).

Let K represent the average size of the sum of the S3 account names, S3 bucket, and S3 key. Suppose the S3 account name is always my-s3-account (13 bytes), buckets have fixed-length names like

/my/application/bucket-12345 (28 bytes), and objects have fixed-length keys like 5733a5d7-f069-41ef-8fbd-13247494c69c (36 bytes). Then K has a value of 90 (13+13+28+36).

If you can determine values for P and K, you can estimate the volume of application logs your external syslog server will have to be able to handle using the following formulas.

$$\begin{aligned} \text{Application Log Rate} &= ((3.2 \times P) + (1.3 \times (1 - P))) \times \text{S3 Operations Rate} \\ \text{Application Log Average Size} &= (P \times (240 + (0.4 \times K))) + ((1 - P) \times (185 + (0.9 \times K))) \text{ Bytes} \end{aligned}$$

So, for example, if your grid is sized for 1,000 S3 operations per second, your workload is 50% PUTs, and your S3 account names, bucket names, and object names average 90 bytes, your external syslog server should be sized to support 2,250 application logs per second and should be able to receive (and typically store) application data at a rate of 0.6 MB per second.

Configure audit messages and external syslog server

You can configure a number of settings related to audit messages. You can adjust the number of audit messages recorded; define any HTTP request headers you want to include in client read and write audit messages; configure an external syslog server; and specify where audit logs, security event logs, and StorageGRID software logs are sent.

Audit messages and logs record system activities and security events, and are essential tools for monitoring and troubleshooting. All StorageGRID nodes generate audit messages and logs to track system activity and events.

Optionally, you can configure an external syslog server to save audit information remotely. Using an external server minimizes the performance impact of audit message logging without reducing the completeness of audit data. An external syslog server is especially useful if you have a large grid, use multiple types of S3 applications, or want to retain all audit data. See [Configure audit messages and external syslog server](#) for details.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Maintenance or Root access permission](#).
- If you plan to configure an external syslog server, you have reviewed the [considerations for using an external syslog server](#) and ensured that the server has enough capacity to receive and store the log files.
- If you plan to configure an external syslog server using TLS or RELP/TLS protocol, you have the required server CA and client certificates and the client private key.

Change audit message levels

You can set a different audit level for each of the following categories of messages in the audit log:

Audit category	Default setting	More information
System	Normal	System audit messages
Storage	Error	Object storage audit messages

Audit category	Default setting	More information
Management	Normal	Management audit message
Client reads	Normal	Client read audit messages
Client writes	Normal	Client write audit messages
ILM	Normal	ILM audit messages
Cross-grid replication	Error	CGRR: Cross-Grid Replication Request



These defaults apply if you initially installed StorageGRID using version 10.3 or later. If you initially used an earlier version of StorageGRID, the default for all categories is set to Normal.



During upgrades, audit level configurations will not be effective immediately.

Steps

1. Select **CONFIGURATION > Monitoring > Audit and syslog server**.
2. For each category of audit message, select an audit level from the drop-down list:

Audit level	Description
Off	No audit messages from the category are logged.
Error	Only error messages are logged—audit messages for which the result code was not "successful" (SUCS).
Normal	Standard transactional messages are logged—the messages listed in these instructions for the category.
Debug	Deprecated. This level behaves the same as the Normal audit level.

The messages included for any particular level include those that would be logged at the higher levels. For example, the Normal level includes all of the Error messages.



If you don't require a detailed record of client read operations for your S3 applications, optionally change the **Client Reads** setting to **Error** to decrease the number of audit messages recorded in the audit log.

3. Select **Save**.

A green banner indicates your configuration has been saved.

Define HTTP request headers

You can optionally define any HTTP request headers you want to include in client read and write audit messages. These protocol headers apply to S3 requests only.

Steps

1. In the **Audit protocol headers** section, define the HTTP request headers you want to include in client read and write audit messages.

Use an asterisk (*) as a wildcard to match zero or more characters. Use the escape sequence (*) to match a literal asterisk.

2. Select **Add another header** to create additional headers, if needed.

When HTTP headers are found in a request, they are included in the audit message under the field HTRH.



Audit protocol request headers are logged only if the audit level for **Client Reads** or **Client Writes** is not **Off**.

3. Select **Save**

A green banner indicates your configuration has been saved.

Use an external syslog server

You can optionally configure an external syslog server to save audit logs, application logs, and security event logs to a location outside of your grid.



If you don't want to use an external syslog server, skip this step and go to [Select audit information destinations](#).



If the configuration options available in this procedure aren't flexible enough to meet your requirements, additional configuration options can be applied using the `audit-destinations` endpoints, which are in the private API section of the [Grid Management API](#). For example, you can use the API if you want to use different syslog servers for different groups of nodes.

Enter syslog information

Access the Configure external syslog server wizard and provide the information StorageGRID needs to access the external syslog server.

Steps

1. From the Audit and syslog server page, select **Configure external syslog server**. Or, if you have previously configured an external syslog server, select **Edit external syslog server**.

The Configure external syslog server wizard appears.

2. For the **Enter syslog info** step of the wizard, enter a valid fully qualified domain name or an IPv4 or IPv6 address for the external syslog server in the **Host** field.
3. Enter the destination port on the external syslog server (must be an integer between 1 and 65535). The default port is 514.
4. Select the protocol used to send audit information to the external syslog server.

Using **TLS** or **RELPL/TLS** is recommended. You must upload a server certificate to use either of these options. Using certificates helps secure the connections between your grid and the external syslog server. For more information, see [Manage security certificates](#).

All protocol options require support by, and configuration of, the external syslog server. You must choose an option that is compatible with the external syslog server.



Reliable Event Logging Protocol (RELPL) extends the functionality of the syslog protocol to provide reliable delivery of event messages. Using RELPL can help prevent the loss of audit information if your external syslog server has to restart.

5. Select **Continue**.

6. If you selected **TLS** or **RELPL/TLS**, upload the server CA certificates, client certificate, and client private key.
 - a. Select **Browse** for the certificate or key you want to use.
 - b. Select the certificate or key file.
 - c. Select **Open** to upload the file.

A green check appears next to the certificate or key file name, notifying you that it has been uploaded successfully.

7. Select **Continue**.

Manage syslog content

You can select which information to send to the external syslog server.

Steps

1. For the **Manage syslog content** step of the wizard, select each type of audit information you want to send to the external syslog server.
 - **Send audit logs:** Sends StorageGRID events and system activities
 - **Send security events:** Sends security events such as when an unauthorized user attempts to sign in or a user signs in as root
 - **Send application logs:** Sends [StorageGRID software log files](#) useful for troubleshooting, including:
 - `broadcast-err.log`
 - `broadcast.log`
 - `jaeger.log`
 - `nms.log` (Admin Nodes only)
 - `prometheus.log`
 - `raft.log`
 - `hagroups.log`
 - **Send access logs:** Sends HTTP access logs for external requests to Grid Manager, Tenant Manager, configured load balancer endpoints, and grid federation requests from remote systems.
2. Use the drop-down menus to select the severity and facility (type of message) for each category of audit information you want to send.

Setting severity and facility values can help you aggregate the logs in customizable ways for easier analysis.

- a. For **Severity**, select **Passthrough**, or select a severity value between 0 and 7.

If you select a value, the selected value will be applied to all messages of this type. Information about different severities will be lost if you override severity with a fixed value.

Severity	Description
Passthrough	Each message sent to the external syslog to have the same severity value as when it was logged locally onto the node: • For audit logs, the severity is "info." • For security events, the severity values are generated by the Linux distribution on the nodes. • For application logs, the severities vary between "info" and "notice," depending on what the issue is. For example, adding an NTP server and configuring an HA group gives a value of "info," while intentionally stopping the SSM or RSM service gives a value of "notice." • For access logs, the severity is "info."
0	Emergency: System is unusable
1	Alert: Action must be taken immediately
2	Critical: Critical conditions
3	Error: Error conditions
4	Warning: Warning conditions
5	Notice: Normal but significant condition
6	Informational: Informational messages
7	Debug: Debug-level messages

- b. For **Facility**, select **Passthrough**, or select a facility value between 0 and 23.

If you select a value, it will be applied to all messages of this type. Information about different facilities will be lost if you override facility with a fixed value.

Facility	Description
Passthrough	Each message sent to the external syslog to have the same facility value as when it was logged locally onto the node: • For audit logs, the facility sent to the external syslog server is "local7." • For security events, the facility values are generated by the linux distribution on the nodes. • For application logs, the application logs sent to the external syslog server have the following facility values: ◦ <code>broadcast.log</code>: user or daemon ◦ <code>broadcast-err.log</code>: user, daemon, local3, or local4 ◦ <code>jaeger.log</code>: local2 ◦ <code>nms.log</code>: local3 ◦ <code>prometheus.log</code>: local4 ◦ <code>raft.log</code>: local5 ◦ <code>hagroups.log</code>: local6 • For access logs, the facility sent to the external syslog server is "local0."
0	kern (kernel messages)
1	user (user-level messages)
2	mail
3	daemon (system daemons)
4	auth (security/authorization messages)
5	syslog (messages generated internally by syslogd)
6	lpr (line printer subsystem)
7	news (network news subsystem)
8	UUCP
9	cron (clock daemon)
10	security (security/authorization messages)
11	FTP

Facility	Description
12	NTP
13	logaudit (log audit)
14	logalert (log alert)
15	clock (clock daemon)
16	local0
17	local1
18	local2
19	local3
20	local4
21	local5
22	local6
23	local7

3. Select **Continue**.

Send test messages

Before starting to use an external syslog server, you should request that all nodes in your grid send test messages to the external syslog server. You should use these test messages to help you validate your entire log collection infrastructure before you commit to sending data to the external syslog server.



Don't use the external syslog server configuration until you confirm that the external syslog server received a test message from each node in your grid and that the message was processed as expected.

Steps

1. If you don't want to send test messages because you are certain your external syslog server is configured properly and can receive audit information from all the nodes in your grid, select **Skip and finish**.

A green banner indicates that the configuration has been saved.

2. Otherwise, select **Send test messages** (recommended).

Test results continuously appear on the page until you stop the test. While the test is in progress, your audit messages continue to be sent to your previously configured destinations.

3. If you receive any errors, correct them and select **Send test messages** again.

See [Troubleshoot an external syslog server](#) to help you resolve any errors.

4. Wait until you see a green banner indicating all nodes have passed testing.

5. Check your syslog server to determine if test messages are being received and processed as expected.



If you are using UDP, check your entire log collection infrastructure. The UDP protocol does not allow for as rigorous error detection as the other protocols.

6. Select **Stop and finish**.

You are returned to the **Audit and syslog server** page. A green banner indicates that the syslog server configuration has been saved.



StorageGRID audit information is not sent to the external syslog server until you select a destination that includes the external syslog server.

Select audit information destinations

You can specify where audit logs, security event logs, and [StorageGRID software logs](#) are sent.

StorageGRID defaults to local node audit destinations and stores the audit information in `/var/local/log/localaudit.log`.



When using `/var/local/log/localaudit.log`, the Grid Manager and Tenant Manager audit log entries might be sent to a Storage Node. You can find which node has the most recent entries by using the `run-each-node --parallel "zgrep MGAU /var/local/log/localaudit.log | tail"` command.

Some destinations are available only if you have configured an external syslog server.

Steps

1. On the Audit and syslog server page, select the destination for audit information.



Local nodes only and **External syslog server** typically provide better performance.

Option	Description
Local nodes only (default)	Audit messages, security event logs, and application logs are not sent to Admin Nodes. Instead, they are saved only on the nodes that generated them ("the local node"). The audit information generated on every local node is stored in <code>/var/local/log/localaudit.log</code>. Note: StorageGRID periodically removes local logs in a rotation to free up space. When the log file for a node reaches 1 GB, the existing file is saved, and a new log file is started. The rotation limit for the log is 21 files. When the 22nd version of the log file is created, the oldest log file is deleted. On average about 20 GB of log data is stored on each node.
Admin Nodes/local nodes	Audit messages are sent to the audit log on Admin Nodes, and security event logs and application logs are stored on the nodes that generated them. The audit information is stored in the following files: • 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. It might contain secondary information, such as an additional copy of some messages.
External syslog server	Audit information is sent to an external syslog server and saved on the local nodes (<code>/var/local/log/localaudit.log</code>). The type of information sent depends upon how you configured the external syslog server. This option is enabled only after you have configured an external syslog server.
Admin Node and external syslog server	Audit messages are sent to the audit log (<code>/var/local/audit/export/audit.log</code>) on Admin Nodes, and audit information is sent to the external syslog server and saved on the local node (<code>/var/local/log/localaudit.log</code>). The type of information sent depends upon how you configured the external syslog server. This option is enabled only after you have configured an external syslog server.

2. Select **Save**.

A warning message appears.

3. Select **OK** to confirm that you want to change the destination for audit information.

A green banner indicates that the audit configuration has been saved.

New logs are sent to the destinations you selected. Existing logs remain in their current location.

Use SNMP monitoring

Use SNMP monitoring

If you want to monitor StorageGRID using the Simple Network Management Protocol (SNMP), you must configure the SNMP agent that is included with StorageGRID.

- [Configure the SNMP agent](#)
- [Update the SNMP agent](#)

Capabilities

Each StorageGRID node runs an SNMP agent, or daemon, that provides a MIB. The StorageGRID MIB contains table and notification definitions for alerts. The MIB also contains system description information such as platform and model number for each node. Each StorageGRID node also supports a subset of MIB-II objects.



See [Access MIB files](#) if you want to download the MIB files on your grid nodes.

Initially, SNMP is disabled on all nodes. When you configure the SNMP agent, all StorageGRID nodes receive the same configuration.

The StorageGRID SNMP agent supports all three versions of the SNMP protocol. It provides read-only MIB access for queries, and it can send two types of event-driven notifications to a management system:

Traps

Traps are notifications sent by the SNMP agent that don't require acknowledgment by the management system. Traps serve to notify the management system that something has happened within StorageGRID, such as an alert being triggered.

Traps are supported in all three versions of SNMP.

Informs

Informs are similar to traps, but they require acknowledgment by the management system. If the SNMP agent doesn't receive an acknowledgment within a certain amount of time, it resends the inform until an acknowledgment is received or the maximum retry value has been reached.

Informs are supported in SNMPv2c and SNMPv3.

Trap and inform notifications are sent in the following cases:

- A default or custom alert is triggered at any severity level. To suppress SNMP notifications for an alert, you must [configure a silence](#) for the alert. Alert notifications are sent by the [preferred sender Admin Node](#).

Each alert is mapped to one of three trap types based on the severity level of the alert: activeMinorAlert, activeMajorAlert, and activeCriticalAlert. For a list of the alerts that can trigger these traps, see the [Alerts reference](#).

SNMP version support

The table provides a high-level summary of what is supported for each SNMP version.

	SNMPv1	SNMPv2c	SNMPv3
Queries (GET and GETNEXT)	Read-only MIB queries	Read-only MIB queries	Read-only MIB queries
Query authentication	Community string	Community string	User-based Security Model (USM) user
Notifications (TRAP and INFORM)	Traps only	Traps and informs	Traps and informs
Notification authentication	Default trap community or a custom community string for each trap destination	Default trap community or a custom community string for each trap destination	USM user for each trap destination

Limitations

- StorageGRID supports read-only MIB access. Read-write access is not supported.
- All nodes in the grid receive the same configuration.
- SNMPv3: StorageGRID does not support the Transport Support Mode (TSM).
- SNMPv3: The only authentication protocol supported is SHA (HMAC-SHA-96).
- SNMPv3: The only privacy protocol supported is AES.

Configure the SNMP agent

You can configure the StorageGRID SNMP agent to use a third-party SNMP management system for read-only MIB access and notifications.

Before you begin

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

About this task

The StorageGRID SNMP agent supports SNMPv1, SNMPv2c, and SNMPv3. You can configure the agent for one or more versions.

For SNMPv3, only User Security Model (USM) authentication is supported.

All nodes in the grid use the same SNMP configuration.

Specify basic configuration

As a first step, enable the StorageGRID SMNP agent and provide basic information.

Steps

1. Select **CONFIGURATION > Monitoring > SNMP agent**.

The SNMP agent page appears.

2. To enable the SNMP agent on all grid nodes, select the **Enable SNMP** checkbox.
3. Enter the following information in the Basic configuration section.

Field	Description
System contact	Optional. The primary contact for the StorageGRID system, which is returned in SNMP messages as sysContact. The System contact is typically an email address. This value applies to all nodes in the StorageGRID system. System contact can be a maximum of 255 characters.
System location	Optional. The location of the StorageGRID system, which is returned in SNMP messages as sysLocation. The System location can be any information that is useful for identifying where your StorageGRID system is located. For example, you might use the street address of a facility. This value applies to all nodes in the StorageGRID system. System location can be a maximum of 255 characters.
Enable SNMP agent notifications	• If selected, the StorageGRID SNMP agent sends trap and inform notifications. • If not selected, the SNMP agent supports read-only MIB access, but it doesn't send any SNMP notifications.
Enable authentication traps	If selected, the StorageGRID SNMP agent sends authentication traps if it receives improperly authenticated protocol messages.

Enter community strings

If you use SNMPv1 or SNMPv2c, complete the Community strings section.

When the management system queries the StorageGRID MIB, it sends a community string. If the community string matches one of the values specified here, the SNMP agent sends a response to the management system.

Steps

1. For **Read-only community**, optionally enter a community string to allow read-only MIB access on IPv4 and IPv6 agent addresses.



To ensure the security of your StorageGRID system, don't use "public" as the community string. If you leave this field blank, the SNMP agent uses the grid ID of your StorageGRID system as the community string.

Each community string can be a maximum of 32 characters and can't contain whitespace characters.

2. Select **Add another community string** to add additional strings.

Up to five strings are allowed.

Create trap destinations

Use the Trap destinations tab in the Other configurations section to define one or more destinations for StorageGRID trap or inform notifications. When you enable the SNMP agent and select **Save**, StorageGRID sends notifications to each defined destination when alerts are triggered. Standard notifications are also sent for the supported MIB-II entities (for example, ifDown and coldStart).

Steps

1. For the **Default trap community** field, optionally enter the default community string you want to use for SNMPv1 or SNMPv2 trap destinations.

As required, you can provide a different ("custom") community string when you define a specific trap destination.

Default trap community can be a maximum of 32 characters and can't contain whitespace characters.

2. To add a trap destination, select **Create**.
3. Select which SNMP version will be used for this trap destination.
4. Complete the Create trap destination form for the version you selected.

SNMPv1

If you selected SNMPv1 as the version, complete these fields.

Field	Description
Type	Must be Trap for SNMPv1.
Host	An IPv4 or IPv6 address or a fully-qualified domain name (FQDN) to receive the trap.
Port	Use 162, which is the standard port for SNMP traps unless you must use another value.
Protocol	Use UDP, which is the standard SNMP trap protocol unless you need to use TCP.
Community string	Use the default trap community, if one was specified, or enter a custom community string for this trap destination. The custom community string can be a maximum of 32 characters and can't contain whitespace.

SNMPv2c

If you selected SNMPv2c as the version, complete these fields.

Field	Description
Type	Whether the destination will be used for traps or informs.
Host	An IPv4 or IPv6 address or FQDN to receive the trap.
Port	Use 162, which is the standard port for SNMP traps unless you must use another value.
Protocol	Use UDP, which is the standard SNMP trap protocol unless you need to use TCP.
Community string	Use the default trap community, if one was specified, or enter a custom community string for this trap destination. The custom community string can be a maximum of 32 characters and can't contain whitespace.

SNMPv3

If you selected SNMPv3 as the version, complete these fields.

Field	Description
Type	Whether the destination will be used for traps or informs.
Host	An IPv4 or IPv6 address or FQDN to receive the trap.
Port	Use 162, which is the standard port for SNMP traps unless you must use another value.
Protocol	Use UDP, which is the standard SNMP trap protocol unless you need to use TCP.
USM user	The USM user that will be used for authentication. • If you selected Trap, only USM users without authoritative engine IDs are shown. • If you selected Inform, only USM users with authoritative engine IDs are shown. • If no users are shown: 1. Create and save the trap destination. 2. Go to Create USM users and create the user. 3. Return to the Trap destinations tab, select the saved destination from the table, and select Edit. 4. Select the user.

5. Select **Create**.

The trap destination is created and added to the table.

Create agent addresses

Optionally, use the Agent addresses tab in the Other configurations section to specify one or more "listening addresses." These are the StorageGRID addresses on which the SNMP agent can receive queries.

If you don't configure an agent address, the default listening address is UDP port 161 on all StorageGRID networks.

Steps

1. Select **Create**.
2. Enter the following information.

Field	Description
Internet protocol	Whether this address will use IPv4 or IPv6. By default, SNMP uses IPv4.

Field	Description
Transport protocol	Whether this address will use UDP or TCP. By default, SNMP uses UDP.
StorageGRID network	Which StorageGRID network the agent will listen on. • Grid, Admin, and Client Networks: The SNMP agent will listen for queries on all three networks. • Grid Network • Admin Network • Client Network Note: If you use the Client Network for insecure data and you create an agent address for the Client Network, be aware that SNMP traffic will also be insecure.
Port	Optionally, the port number that the SNMP agent should listen on. The default UDP port for an SNMP agent is 161, but you can enter any unused port number. Note: When you save the SNMP agent, StorageGRID automatically opens the agent address ports on the internal firewall. You must ensure that any external firewalls allow access to these ports.

3. Select **Create**.

The agent address is created and added to the table.

Create USM users

If you are using SNMPv3, use the USM users tab in the Other configurations section to define the USM users who are authorized to query the MIB or to receive traps and informs.



SNMPv3 *inform* destinations must have users with engine IDs. SNMPv3 *trap* destination can't have users with engine IDs.

These steps don't apply if you are only using SNMPv1 or SNMPv2c.

Steps

1. Select **Create**.
2. Enter the following information.

Field	Description
Username	A unique name for this USM user. Username can have a maximum of 32 characters and can't contain whitespace characters. The username can't be changed after the user is created.
Read-only MIB access	If selected, this user should have read-only access to the MIB.
Authoritative engine ID	If this user will be used in an inform destination, the authoritative engine ID for this user. Enter 10 to 64 hex characters (5 to 32 bytes) with no spaces. This value is required for USM users that will be selected in trap destinations for informs. This value is not allowed for USM users that will be selected in trap destinations for traps. Note: This field is not shown if you selected Read-only MIB access because USM users who have read-only MIB access can't have engine IDs.
Security level	The security level for the USM user: • authPriv: This user communicates with authentication and privacy (encryption). You must specify an authentication protocol and password and a privacy protocol and password. • authNoPriv: This user communicates with authentication and without privacy (no encryption). You must specify an authentication protocol and password.
Authentication protocol	Always set to SHA, which is the only protocol supported (HMAC-SHA-96).
Password	The password this user will use for authentication.
Privacy protocol	Shown only if you selected authPriv and always set to AES, which is the only privacy protocol supported.
Password	Shown only if you selected authPriv . The password this user will use for privacy.

3. Select **Create**.

The USM user is created and added to the table.

4. When you have completed the SNMP agent configuration, select **Save**.

The new SNMP agent configuration becomes active.

Update the SNMP agent

You can disable SNMP notifications, update community strings, or add or remove agent addresses, USM users, and trap destinations.

Before you begin

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

About this task

See [Configure the SNMP agent](#) for details about each field on the SNMP agent page. You must select **Save** at the bottom of the page to commit any changes you make on each tab.

Steps

1. Select **CONFIGURATION > Monitoring > SNMP agent**.

The SNMP agent page appears.

2. To disable the SNMP agent on all grid nodes, clear the **Enable SNMP** checkbox, and select **Save**.

If you re-enable the SNMP agent, any previous SNMP configuration settings are retained.

3. Optionally, update the information in the Basic configuration section:
 - a. As required, update the **System contact** and **System location**.
 - b. Optionally, select or clear the **Enable SNMP agent notifications** checkbox to control whether the StorageGRID SNMP agent sends trap and inform notifications.

When this checkbox is cleared, the SNMP agent supports read-only MIB access, but it doesn't send SNMP notifications.

- c. Optionally, select or clear the **Enable authentication traps** checkbox to control whether the StorageGRID SNMP agent sends authentication traps when it receives improperly authenticated protocol messages.
4. If you use SNMPv1 or SNMPv2c, optionally update or add a **Read-only community** in the Community strings section.
 5. To update trap destinations, select the Trap destinations tab in the Other configurations section.

Use this tab to define one or more destinations for StorageGRID trap or inform notifications. When you enable the SNMP agent and select **Save**, StorageGRID sends notifications to each defined destination when alerts are triggered. Standard notifications are also sent for the supported MIB-II entities (for example, ifDown and coldStart).

For details about what to enter, see [Create trap destinations](#).

- Optionally, update or remove the default trap community.

If you remove the default trap community, you must first ensure that any existing trap destinations use a custom community string.

- To add a trap destination, select **Create**.
- To edit a trap destination, select the radio button, and select **Edit**.

- To remove a trap destination, select the radio button, and select **Remove**.
- To commit your changes, select **Save** at the bottom of the page.

6. To update agent addresses, select the Agent addresses tab in the Other configurations section.

Use this tab to specify one or more "listening addresses." These are the StorageGRID addresses on which the SNMP agent can receive queries.

For details about what to enter, see [Create agent addresses](#).

- To add an agent address, select **Create**.
- To edit an agent address, select the radio button, and select **Edit**.
- To remove an agent address, select the radio button, and select **Remove**.
- To commit your changes, select **Save** at the bottom of the page.

7. To update USM users, select the USM users tab in the Other configurations section.

Use this tab to define the USM users who are authorized to query the MIB or to receive traps and informs.

For details about what to enter, see [Create USM users](#).

- To add a USM user, select **Create**.
- To edit a USM user, select the radio button, and select **Edit**.

The username for an existing USM user can't be changed. If you need to change a username, you must remove the user and create a new one.



If you add or remove a user's authoritative engine ID and that user is currently selected for a destination, you must edit or remove the destination. Otherwise, a validation error occurs when you save the SNMP agent configuration.

- To remove a USM user, select the radio button, and select **Remove**.



If the user you removed is currently selected for a trap destination, you must edit or remove the destination. Otherwise, a validation error occurs when you save the SNMP agent configuration.

- To commit your changes, select **Save** at the bottom of the page.

8. When you have updated the SNMP agent configuration, select **Save**.

Access MIB files

MIB files contain definitions and information about the properties of managed resources and services for the nodes in your grid. You can access MIB files that define the objects and notifications for StorageGRID. These files can be useful for monitoring your grid.

See [Use SNMP monitoring](#) for more information about SNMP and MIB files.

Access MIB files

Follow these steps to access the MIB files.

Steps

1. Select **CONFIGURATION > Monitoring > SNMP agent**.
2. On the SNMP agent page, select the file you want to download:
 - **NETAPP-STORAGEGRID-MIB.txt**: Defines the alert table and notifications (traps) accessible on all Admin Nodes.
 - **ES-NETAPP-06-MIB.mib**: Defines objects and notifications for E-Series-based appliances.
 - **MIB_1_10.zip**: Defines objects and notifications for appliances with a BMC interface.



You can also access MIB files at the following location on any StorageGRID node:
`/usr/share/snmp/mibs`

3. To extract the StorageGRID OIDs from the MIB file:
 - a. Get the OID of the root of the StorageGRID MIB:

```
root@user-adm1:~ # snmptranslate -On -IR storagegrid
```

Result: `.1.3.6.1.4.1.789.28669` (28669 is always the OID for StorageGRID)

- b. Grep for the StorageGRID OID in the entire tree (using `paste` to join lines):

```
root@user-adm1:~ # snmptranslate -Tso | paste -d " " - - | grep 28669
```



The `snmptranslate` command has many options that are useful for exploring the MIB. This command is available on any StorageGRID node.

MIB file contents

All objects are under the StorageGRID OID.

Object name	Object ID (OID)	Description
<code>.iso.org.dod.internet. private.enterprises .netapp.storagegrid</code>	<code>.1.3.6.1.4.1.789.28669</code>	The MIB module for NetApp StorageGRID entities.

MIB objects

Object name	Object ID (OID)	Description
<code>activeAlertCount</code>	<code>.1.3.6.1.4.1.789.28669.1.3</code>	The number of active alerts in the <code>activeAlertTable</code> .
<code>activeAlertTable</code>	<code>.1.3.6.1.4.1.789.28669.1.4</code>	A table of active alerts in StorageGRID.

Object name	Object ID (OID)	Description
activeAlertId	.1.3.6.1.4.1. 789.28669.1.4.1.1	The ID of the alert. Only unique in the current set of active alerts.
activeAlertName	.1.3.6.1.4.1. 789.28669.1.4.1.2	The name of the alert.
activeAlertInstance	.1.3.6.1.4.1. 789.28669.1.4.1.3	The name of the entity that generated the alert, typically the node name.
activeAlertSeverity	.1.3.6.1.4.1. 789.28669.1.4.1.4	The severity of the alert.
activeAlertStartTime	.1.3.6.1.4.1. 789.28669.1.4.1.5	The date and time the alert was triggered.

Notification types (Traps)

All notifications include the following variables as varbinds:

- activeAlertId
- activeAlertName
- activeAlertInstance
- activeAlertSeverity
- activeAlertStartTime

Notification type	Object ID (OID)	Description
activeMinorAlert	.1.3.6.1.4.1. 789.28669.0.6	An alert with minor severity
activeMajorAlert	.1.3.6.1.4.1. 789.28669.0.7	An alert with major severity
activeCriticalAlert	.1.3.6.1.4.1. 789.28669.0.8	An alert with critical severity

Collect additional StorageGRID data

Use charts and graphs

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




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

Types of charts

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

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



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

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

Dashboard

16 Buckets
[View buckets](#)

2 Platform services
endpoints
[View endpoints](#)

0 Groups
[View groups](#)

1 User
[View users](#)

Storage usage [?](#)

6.5 TB of 7.2 TB used

0.7 TB (10.1%) remaining



Bucket name	Space used	Number of objects
Bucket-15	969.2 GB	913,425
Bucket-04	937.2 GB	576,806
Bucket-13	815.2 GB	957,389
Bucket-06	812.5 GB	193,843
Bucket-10	473.9 GB	583,245
Bucket-03	403.2 GB	981,226
Bucket-07	362.5 GB	420,726
Bucket-05	294.4 GB	785,190
8 other buckets	1.4 TB	3,007,036

Top buckets by capacity limit usage [?](#)

Bucket name	Usage
Bucket-10	82%
Bucket-03	57%
Bucket-15	20%

Tenant details [?](#)

Name: Tenant02

ID: 3341 1240 0546 8283 2208


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

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

There are four types of graphs:

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

DC1-S2 (Storage Node)

[Overview](#)[Hardware](#)[Network](#)[Storage](#)[Objects](#)[ILM](#)[Tasks](#)[1 hour](#)[1 day](#)[1 week](#)[1 month](#)[Custom](#)

Network interfaces

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

Network communication

Receive

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


Transmit

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




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


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




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



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



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



Related information

- [View the Nodes page](#)
- [View the Grid Topology tree](#)
- [Review support metrics](#)

Chart legend

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

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

Display charts and graphs

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


Before you begin

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


Steps

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

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



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



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



Charts aren't available for all metrics and attributes.

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




Reports (Charts): DDS (DC1-S1) - Data Store




Attribute: Total Queries - Successful ▼
Quick Query: Last Hour ▼ Update

Vertical Scaling: ☒
Raw Data: ☐

YYYY/MM/DD HH:MM:SS
Start Date: 2021/09/24 13:24:01
End Date: 2021/09/24 14:24:01




Close



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

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



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

Overview

Alarms

Reports

Configuration

Main



Overview: SSM (DC1-ADM1) - Resources

Updated: 2018-05-07 16:29:52 MDT

Computational Resources

Service Restarts:	1	
Service Runtime:	6 days	
Service Uptime:	6 days	
Service CPU Seconds:	10666 s	
Service Load:	0.266 %	

Memory

Installed Memory:	8.38 GB	
Available Memory:	2.9 GB	

Processors

Processor Number	Vendor	Type	Cache
1	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
2	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
3	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
4	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
5	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
6	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
7	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB
8	GenuineIntel	Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz	15 MiB

- Select the chart icon  next to the attribute.

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

Generate charts

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

Before you begin

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

Steps

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

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

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

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

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

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

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

8. Select **Update**.

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

Use text reports

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

Raw text reports

A raw text report displays details about the selected attribute:

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

Text Results for Services: Load - System Logging

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

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

Aggregate text reports

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

Each entry displays the following information:

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

Text Results for Attribute Send to Relay Rate

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

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

Generate text reports

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

Before you begin

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

About this task

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

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

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

Steps

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

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

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

7. If you selected Custom Query, you need to customize the time period to report on by entering the **Start Date** and **End Date**.

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

8. Click **Update**.

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

Export text reports

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

About this task

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

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Create a text report.
3. Click *Export* .

Overview

Alarms

Reports

Configuration

Charts

Text

**Reports (Text): SSM (170-176) - Events**

Attribute:

Attribute Send to Relay Rate

Quick Query:

Custom Query

Update

Results Per Page:

5

Raw Data: ☒

YYYY/MM/DD HH:MM:SS

Start Date:

2010/07/19 08:42:09

End Date:

2010/07/20 08:42:09

Text Results for Attribute Send to Relay Rate
2010-07-19 08:42:09 PDT To 2010-07-20 08:42:09 PDT1 - 5 of 254 

Previous « 1 2 3 4 5 » Next

The Export Text Report window opens displaying the report.

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

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

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

Monitor PUT and GET performance

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

About this task

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

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

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

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

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

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

Monitor object verification operations

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

Before you begin

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

About this task

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

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

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



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

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

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

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

Steps

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

Verification		
Status: ?	No errors	
Percent complete: ?	0.00%	
Average stat time: ?	0.00 microseconds	
Objects verified: ?	0	
Object verification rate: ?	0.00 objects / second	
Data verified: ?	0 bytes	
Data verification rate: ?	0.00 bytes / second	
Missing objects: ?	0	
Corrupt objects: ?	0	
Corrupt objects unidentified: ?	0	
Quarantined objects: ?	0	

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

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

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

2. Review the results from object existence check jobs:

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

Object existence check

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

Active job

Job history

Delete

Search...

☐	Job ID ?	Status	Nodes (volumes) ?	Missing object copies detected ?
☐	15816859223101303015	Completed	DC2-S1 (3 volumes)	0
☐	12538643155010477372	Completed	DC1-S3 (1 volume)	0
☐	5490044849774982476	Completed	DC1-S2 (1 volume)	0
☐	3395284277055907678	Completed	DC1-S1 (3 volumes) DC1-S2 (3 volumes) DC1-S3 (3 volumes) and 7 more	0

Monitor events

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

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

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

Steps

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

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

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

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

Reports (Text): SSM (170-41) - Events

Attribute: Results Per Page: Start Date: End Date:

Quick Query: Raw Data: ☒

Text Results for Last Event
2009-04-15 15:19:53 PDT To 2009-04-15 15:24:53 PDT

1 - 2 of 2

Time Received	Sample Time	Value
2009-04-15 15:24:22	2009-04-15 15:24:22	hdc: task_no_data_intr: status=0x51 { DriveReady SeekComplete Error }
2009-04-15 15:24:11	2009-04-15 15:23:39	hdc: task_no_data_intr: status=0x51 { DriveReady SeekComplete Error }

Create custom syslog events

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

About this task

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

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

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

Steps

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

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

Events
Updated: 2021-10-22 11:15:34 MDT

Custom Events (1 - 1 of 1)

Event	Actions
shutdown	

Show 10 Records Per Page Refresh Previous 1 Next

Apply Changes

4. Select **Apply Changes**.
5. Select **SUPPORT > Tools > Grid topology**.
6. Select **grid node > SSM > Events**.
7. Locate the entry for Custom Events in the Events table, and monitor the value for **Count**.

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

Overview
Alarms
Reports
Configuration

Main



Overview: SSM (DC1-ADM1) - Events
Updated: 2021-10-22 11:19:18 MDT

System Events

Log Monitor State: Connected
Total Events: 0
Last Event: No Events

Description	Count
Abnormal Software Events	0
Account Service Events	0
Cassandra Errors	0
Cassandra Heap Out Of Memory Errors	0
Chunk Service Events	0
Custom Events	0
Data-Mover Service Events	0
File System Errors	0
Forced Termination Events	0
Grid Node Errors	0
Hotfix Installation Failure Events	0
I/O Errors	0
IDE Errors	0
Identity Service Events	0
Kernel Errors	0
Kernel Memory Allocation Failure	0
Keystone Service Events	0
Network Receive Errors	0
Network Transmit Errors	0
Out Of Memory Errors	0
Replicated State Machine Service Events	0
SCSI Errors	0

Reset the count of custom events to zero

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

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

Steps

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

Overview

Alarms

Reports

Configuration

Main

Alarms



Configuration: SSM (DC2-ADM1) - Events

Updated: 2018-04-11 10:35:44 MDT

Description	Count	Reset
Abnormal Software Events	0	☐
Account Service Events	0	☐
Cassandra Errors	0	☐
Cassandra Heap Out Of Memory Errors	0	☐
Custom Events	0	☒
File System Errors	0	☐
Forced Termination Events	0	☐

4. Select **Apply Changes**.

Review audit messages

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

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

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

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

You can access audit log files directly from the command line of the Admin Node.

StorageGRID can send audit information by default, or you can change the destination:

- StorageGRID defaults to local node audit destinations.
- Grid Manager and Tenant Manager audit log entries might be sent to a Storage Node.
- Optionally, you can change the destination of audit logs and send audit information to an external syslog server. Local logs of audit records continue to be generated and stored when an external syslog server is

configured.

- [Learn about configuring audit messages and log destinations.](#)

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

Collect log files and system data

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

Before you begin

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

About this task

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

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

Steps

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

The screenshot shows the StorageGRID log collection configuration window. On the left, a tree view shows the hierarchy: StorageGRID (expanded) -> DC1 (expanded) -> DC1-ADM1, DC1-G1, DC1-S1 (selected), DC1-S2, DC1-S3, DC1-S4; and DC2 (expanded) -> DC2-ADM1, DC2-G1, DC2-S1 (selected), DC2-S2, DC2-S3, DC2-S4. On the right, the 'Log Start Time' is set to 2021-12-03 06:31 AM MST and the 'Log End Time' is set to 2021-12-03 10:31 AM MST. Under 'Log Types', 'Application Logs' is checked, while 'Audit Logs', 'Network Trace', and 'Prometheus Database' are unchecked. There is a large empty 'Notes' text box and a 'Provisioning Passphrase' field with 10 masked characters. A blue 'Collect Logs' button is located at the bottom right.

2. Select the grid nodes for which you want to collect log files.

As required, you can collect log files for the entire grid or an entire data center site.

3. Select a **Start Time** and **End Time** to set the time range of the data to be included in the log files.

If you select a very long time period or collect logs from all nodes in a large grid, the log archive could become too large to be stored on a node, or too large to be collected to the primary Admin Node for download. If this occurs, you must restart log collection with a smaller set of data.

4. Select the types of logs you want to collect.

- **Application Logs:** Application-specific logs that technical support uses most frequently for troubleshooting. The logs collected are a subset of the available application logs.
- **Audit Logs:** Logs containing the audit messages generated during normal system operation.
- **Network Trace:** Logs used for network debugging.
- **Prometheus Database:** Time series metrics from the services on all nodes.

5. Optionally, enter notes about the log files you are gathering in the **Notes** text box.

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

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

7. Select **Collect Logs**.

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

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

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

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

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

After you finish

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

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

Manually trigger an AutoSupport package

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

Before you begin

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

Steps

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

StorageGRID attempts to send an AutoSupport package to the NetApp Support Site. If the attempt is successful, the **Most Recent Result** and **Last Successful Time** values on the **Results** tab are updated. If there is a problem, the **Most Recent Result** value updates to "Failed," and StorageGRID does not try to send the AutoSupport package again.




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

View the Grid Topology tree

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

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



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

StorageGRID attributes

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

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

- **Nodes page:** Many of the values shown on the Nodes page are StorageGRID attributes. (Prometheus metrics are also shown on the Nodes pages.)
- **Grid Topology tree:** Attribute values are shown in the Grid Topology tree (**SUPPORT > Tools > Grid topology**).
- **Events:** System events occur when certain attributes record an error or fault condition for a node, including errors such as network errors.

Attribute values

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

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

Review support metrics

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

Before you begin

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

About this task

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



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

Steps

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

An example of the Metrics page is shown here:

Metrics

Access charts and metrics to help troubleshoot issues.

The tools available on this page are intended for use by technical support. Some features and menu items within these tools are intentionally non-functional.

Prometheus

Prometheus is an open-source toolkit for collecting metrics. The Prometheus interface allows you to query the current values of metrics and to view charts of the values over time.

Access the Prometheus UI using the link below. You must be signed in to the Grid Manager.

- [https://](#)

Grafana

Grafana is open-source software for metrics visualization. The Grafana interface provides pre-constructed dashboards that contain graphs of important metric values over time.

Access the Grafana dashboards using the links below. You must be signed in to the Grid Manager.

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

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


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



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

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

The Grafana interface for the link you selected appears.



Run diagnostics

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

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

Before you begin

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

About this task

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

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

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

Steps

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

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

In this example, all diagnostics have a Normal status.

Diagnostics

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

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

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

Run Diagnostics

✓ Cassandra automatic restarts



✓ Cassandra blocked task queue too large



✓ Cassandra commit log latency



✓ Cassandra commit log queue depth



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

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

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



You can't change these thresholds.

- **Status values:** A table showing the status and the value of the diagnostic throughout the StorageGRID system.

In this example, the current CPU utilization for every node in a StorageGRID system is shown. All node values are below the Attention and Caution thresholds, so the overall status of the diagnostic is Normal.

CPU utilization

Checks the current CPU utilization on each node.

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

Status Normal

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

Thresholds
 Attention $\geq 75\%$
 Caution $\geq 95\%$

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

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

This link is not displayed for all diagnostics.

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



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



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

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

☐ Enable query history

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


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

Execute

- insert metric at cursor - ▾

Graph Console

1h ⏪ Until ⏩ Res. (s) ☐ stacked



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

Remove Graph

Add Graph

Create custom monitoring applications

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

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

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

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



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

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

Troubleshoot StorageGRID system

Troubleshoot a StorageGRID system

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

Often, you can resolve problems on your own; however, you might need to escalate some issues to technical support.

Define the problem

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

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

Question	Example response
What is the StorageGRID system doing or not doing? What are its symptoms?	Client applications are reporting that objects can't be ingested into StorageGRID.
When did the problem start?	Object ingest was first denied at about 14:50 on January 8, 2020.
How did you first notice the problem?	Notified by client application. Also received alert email notifications.
Does the problem happen consistently, or only sometimes?	Problem is ongoing.
If the problem happens regularly, what steps cause it to occur	Problem happens every time a client tries to ingest an object.

Question	Example response
If the problem happens intermittently, when does it occur? Record the times of each incident that you are aware of.	Problem is not intermittent.
Have you seen this problem before? How often have you had this problem in the past?	This is the first time I have seen this issue.

Assess the risk and impact on the system

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

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

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

Collect data

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

Type of data to collect	Why collect this data	Instructions
Create timeline of recent changes	Changes to your StorageGRID system, its configuration, or its environment can cause new behavior.	• Create a timeline of recent changes
Review alerts	Alerts can help you quickly determine the root cause of a problem by providing important clues as to the underlying issues that might be causing it. Review the list of current alerts to see if StorageGRID has identified the root cause of a problem for you. Review alerts triggered in the past for additional insights.	• View current and resolved alerts

Type of data to collect	Why collect this data	Instructions
Monitor events	Events include any system error or fault events for a node, including errors such as network errors. Monitor events to learn more about issues or to help with troubleshooting.	• Monitor events
Identify trends using charts and text reports	Trends can provide valuable clues about when issues first appeared, and can help you understand how quickly things are changing.	• Use charts and graphs • Use text reports
Establish baselines	Collect information about the normal levels of various operational values. These baseline values, and deviations from these baselines, can provide valuable clues.	• Establish baselines
Perform ingest and retrieval tests	To troubleshoot performance issues with ingest and retrieval, use a workstation to store and retrieve objects. Compare results against those seen when using the client application.	• Monitor PUT and GET performance
Review audit messages	Review audit messages to follow StorageGRID operations in detail. The details in audit messages can be useful for troubleshooting many types of issues, including performance issues.	• Review audit messages
Check object locations and storage integrity	If you are having storage problems, verify that objects are being placed where you expect. Check the integrity of object data on a Storage Node.	• Monitor object verification operations • Confirm object data locations • Verify object integrity
Collect data for technical support	Technical support might ask you to collect data or review specific information to help troubleshoot issues.	• Collect log files and system data • Manually trigger an AutoSupport package • Review support metrics

Create a timeline of recent changes

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

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

Create a table of recent changes to your system that includes information about when each change occurred

and any relevant details about the change, such information about what else was happening while the change was in progress:

Time of change	Type of change	Details
For example: • When did you start the node recovery? • When did the software upgrade complete? • Did you interrupt the process?	What happened? What did you do?	Document any relevant details about the change. For example: • Details of the network changes. • Which hotfix was installed. • How client workloads changed. Make sure to note if more than one change was happening at the same time. For example, was this change made while an upgrade was in progress?

Examples of significant recent changes

Here are some examples of potentially significant changes:

- Was the StorageGRID system recently installed, expanded, or recovered?
- Has the system been upgraded recently? Was a hotfix applied?
- Has any hardware been repaired or changed recently?
- Has the ILM policy been updated?
- Has the client workload changed?
- Has the client application or its behavior changed?
- Have you changed load balancers, or added or removed a high availability group of Admin Nodes or Gateway Nodes?
- Have any tasks been started that might take a long time to complete? Examples include:
 - Recovery of a failed Storage Node
 - Storage Node decommissioning
- Have any changes been made to user authentication, such as adding a tenant or changing LDAP configuration?
- Is data migration taking place?
- Were platform services recently enabled or changed?
- Was compliance enabled recently?
- Have Cloud Storage Pools been added or removed?
- Have any changes been made to storage compression or encryption?
- Have there been any changes to the network infrastructure? For example, VLANs, routers, or DNS.
- Have any changes been made to NTP sources?
- Have any changes been made to the Grid, Admin, or Client Network interfaces?
- Have any other changes been made to the StorageGRID system or its environment?

Establish baselines

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

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

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

Analyze data

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

The analysis is problem-dependent, but in general:

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

Escalation information checklist

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

	Item	Notes
	Problem statement	What are the problem symptoms? When did the problem start? Does it happen consistently or intermittently? If intermittently, what times has it occurred? Define the problem
	Impact assessment	What is the severity of the problem? What is the impact to the client application? • Has the client connected successfully before? • Can the client ingest, retrieve, and delete data?

	Item	Notes
	StorageGRID System ID	Select MAINTENANCE > System > License . The StorageGRID System ID is shown as part of the current license.
	Software version	From the top of the Grid Manager, select the help icon and select About to see the StorageGRID version.
	Customization	Summarize how your StorageGRID system is configured. For example, list the following: • Does the grid use storage compression, storage encryption, or compliance? • Does ILM make replicated or erasure-coded objects? Does ILM ensure site redundancy? Do ILM rules use the Balanced, Strict, or Dual Commit ingest behaviors?
	Log files and system data	Collect log files and system data for your system. Select SUPPORT > Tools > Logs . You can collect logs for the entire grid, or for selected nodes. If you are collecting logs only for selected nodes, be sure to include at least one Storage Node that has the ADC service. (The first three Storage Nodes at a site include the ADC service.) Collect log files and system data
	Baseline information	Collect baseline information regarding ingest operations, retrieval operations, and storage consumption. Establish baselines
	Timeline of recent changes	Create a timeline that summarizes any recent changes to the system or its environment. Create a timeline of recent changes
	History of efforts to diagnose the issue	If you have taken steps to diagnose or troubleshoot the issue yourself, make sure to record the steps you took and the outcome.

Troubleshoot object and storage issues

Confirm object data locations

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

object data is being stored where intended.

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: "Identifier" with the text "source/testobject" and "Version ID (optional)" with the text "MEJGMkMyQzgtNEY5OC0xMUU3LTkzMEYlRDkyNTAwQKY5NOMx". A blue "Look Up" button is at the bottom.

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": {
      "PAIRS": "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.

- Select **MAINTENANCE > Object existence check > Job history**.
- Determine which jobs are ready to be rerun:
 - Look at the **End time** column to determine which jobs were run more than three weeks ago.

- ii. For those jobs, scan the Consistency control column for strong-site or strong-global.
- c. Select the 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.



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.

Troubleshoot metadata issues

If metadata issues occur, alerts will inform you of the source of the issues and recommended actions to take. In particular, you must add new Storage Nodes if the Low metadata storage alert is triggered.


Before you begin

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

About this task

Follow the recommended actions for each metadata-related alert that is triggered. If the **Low metadata storage** alert is triggered, you must add new Storage Nodes.

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



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

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

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

```
storagegrid_storage_utilization_metadata_bytes/storagegrid_storage_utilization_metadata_allowed_bytes
```

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

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



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

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



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

Steps

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



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

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

Troubleshoot certificate errors

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

About this task


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

If you are accessing the Grid Manager or Tenant Manager using a domain name instead of an IP address, the

browser shows a certificate error without an option to bypass if either of the following occurs:

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

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



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

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

Steps

If you received an alert notification about an expired certificate, access the certificate details:

. Select **CONFIGURATION > Security > Certificates** and then [select the appropriate certificate tab](#).

1. Check the validity period of the certificate.
Some web browsers and S3 clients don't accept certificates with a validity period greater than 398 days.
2. If the certificate has expired or will expire soon, upload or generate a new certificate.
 - For a server certificate, see the steps for [configuring a custom server certificate for the Grid Manager](#)

and the Tenant Manager.

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

Troubleshoot Admin Node and user interface issues

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

Admin Node sign-in errors

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

Before you begin

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

About this task

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

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

Steps

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

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

2. If your StorageGRID system has more than one Admin Node, try signing in to the Grid Manager from another Admin Node to check the status of an unavailable Admin Node.

- If you are able to sign in, you can use the **Dashboard**, **NODES**, **Alerts**, and **SUPPORT** options to help determine the cause of the error.
- If you have only one Admin Node or you still can't sign in, go to the next step.

3. Determine if the node's hardware is offline.

4. If single sign-on (SSO) is enabled for your StorageGRID system, refer to the steps for [configuring single sign-on](#).

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



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

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

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

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

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

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

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

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

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

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

```

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

```

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

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

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

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

10. Review the following logs:

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

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

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

```
# gpt-list-services adc
```

Admin Nodes use the ADC service during the authentication process.

12. From the Admin Node, use ssh to log in to each of the ADC Storage Nodes, using the IP addresses you identified.
13. View the status of all services running on the grid node: `storagegrid-status`

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

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

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

User interface issues

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

Steps

1. Make sure you're using a [supported web browser](#).
2. Clear your web browser cache.

Clearing the cache removes outdated resources used by the previous version of StorageGRID software, and permits the user interface to operate correctly again. For instructions, see the documentation for your web browser.

Troubleshoot network, hardware, and platform issues

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

"422: Unprocessable Entity" errors

The error 422: Unprocessable Entity can occur for different reasons. Check the error message to determine what caused your issue.

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

Error message	Root cause and corrective action
422: Unprocessable Entity Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration. Unable to authenticate, please verify your username and password: LDAP Result Code 8 "Strong Auth Required": 00002028: LdapErr: DSID-0C090256, comment: The server requires binds to turn on integrity checking if SSL\TLS are not already active on the connection, data 0, v3839	This message might occur if you select the Do not use TLS option for Transport Layer Security (TLS) when configuring identity federation using Windows Active Directory (AD). Using the Do not use TLS option is not supported for use with AD servers that enforce LDAP signing. You must select either the Use STARTTLS option or the Use LDAPS option for TLS.
422: Unprocessable Entity Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration. Unable to begin TLS, verify your certificate and TLS configuration: LDAP Result Code 200 "Network Error": TLS handshake failed (EOF)	This message appears if you try to use an unsupported cipher to make a Transport Layer Security (TLS) connection from StorageGRID to an external system used for identify federation or Cloud Storage Pools. Check the ciphers that are offered by the external system. The system must use one of the ciphers supported by StorageGRID for outgoing TLS connections, as shown in the instructions for administering StorageGRID.

Grid Network MTU mismatch alert

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

About this task

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

Steps

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

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

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

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

Positional arguments	Description
mtu	The MTU to set. Must be in the range 1280 to 9216.
network	The networks to apply the MTU to. Include one or more of the following network types: • grid • admin • client

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

Node network reception frame error alert

Node network reception frame error alerts can be caused by connectivity issues between StorageGRID and your network hardware. This alert clears on its own after the underlying problem is addressed.

About this task

Node network reception frame error alerts can be caused by the following issues with networking hardware

that connects to StorageGRID:

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

Steps

1. Follow the troubleshooting steps for all potential causes of this alert given your network configuration.
2. Perform the following steps depending on the cause of the error:

FEC mismatch



These steps are applicable only for **Node network reception frame error** alerts caused by FEC mismatch on StorageGRID appliances.

- a. Check the FEC status of the port in the switch attached to your StorageGRID appliance.
- b. Check the physical integrity of the cables from the appliance to the switch.
- c. If you want to change FEC settings to try to resolve the alert, first ensure that the appliance is configured for **Auto** mode on the Link Configuration page of the StorageGRID Appliance Installer (see the instructions for your appliance):
 - [SG6160](#)
 - [SGF6112](#)
 - [SG6000](#)
 - [SG5800](#)
 - [SG5700](#)
 - [SG110 and SG1100](#)
 - [SG100 and SG1000](#)
- d. Change the FEC settings on the switch ports. The StorageGRID appliance ports will adjust their FEC settings to match, if possible.

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



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

Switch port and NIC MTU mismatch

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

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



For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The **Grid Network MTU mismatch** alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values don't have to be the same for all network types. See [Troubleshoot the Grid Network MTU mismatch alert](#) for more information.



Also see [Change MTU setting](#).

High link error rates

- a. Enable FEC, if not already enabled.
- b. Verify that your network cabling is of good quality and is not damaged or improperly connected.
- c. If the cables don't appear to be the problem, contact technical support.



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

NIC ring buffer overrun

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

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

3. Monitor the problem and contact technical support if the alert doesn't resolve.

Time synchronization errors

You might see issues with time synchronization in your grid.

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



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

Linux: Network connectivity issues

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

MAC address cloning

In some cases, network issues can be resolved by using MAC address cloning. If you are using virtual hosts, set the value of the MAC address cloning key for each of your networks to "true" in your node configuration file. This setting causes the MAC address of the StorageGRID container to use the MAC address of the host. To create node configuration files, see the instructions for [Red Hat Enterprise Linux](#) or [Ubuntu or Debian](#).



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

For more information about enabling MAC cloning, see the instructions for [Red Hat Enterprise Linux](#) or [Ubuntu or Debian](#).

Promiscuous mode

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

For more information about using Promiscuous Mode, see the instructions for [Red Hat Enterprise Linux](#) or [Ubuntu or Debian](#).

Linux: Node status is "orphaned"

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

About this task

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

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

Steps

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

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

If the node is orphaned, the response is

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

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

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

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

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

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

Linux: Troubleshoot IPv6 support

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

About this task

To see the IPv6 address that has been assigned to a grid node:

1. Select **NODES** and select the node.
2. Select **Show additional IP addresses** next to **IP Addresses** on the Overview tab.

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

Steps

1. Log in to the host as root or using an account with sudo permission.
2. Run the following command: `sysctl net.ipv6.conf.all.disable_ipv6`

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

The result should be 0.

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



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

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

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

The result should be 1.

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



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

5. Exit the container: `exit`

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

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

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

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

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

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

8. Run these commands to restart the StorageGRID container:

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

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

Troubleshoot an external syslog server

The following table describes the error messages that might be related using to an external syslog server and lists corrective actions.

For more information about sending audit information to an external syslog server, see:

- [Considerations for using an external syslog server](#)
- [Configure audit messages and external syslog server](#)

Error message	Description and recommended actions
Cannot resolve hostname	The FQDN you entered for the syslog server could not be resolved to an IP address. 1. Check the hostname you entered. If you entered an IP address, make sure it is a valid IP address in W.X.Y.Z ("dotted decimal") notation. 2. Check that the DNS servers are configured correctly. 3. Confirm that each node can access the IP addresses for the DNS server.

Error message	Description and recommended actions
Connection refused	A TCP or TLS connection to the syslog server was refused. There might be no service listening on the TCP or TLS port for the host, or a firewall might be blocking access. 1. Check that you entered the correct FQDN or IP address, port, and protocol for the syslog server. 2. Confirm that the host for the syslog service is running a syslog daemon that is listening on the specified port. 3. Confirm that a firewall is not blocking access to TCP/TLS connections from the nodes to the IP and port of the syslog server.
Network unreachable	The syslog server is not on a directly attached subnet. A router returned an ICMP failure message to indicate it could not forward the test messages from the listed nodes to the syslog server. 1. Check that you entered the correct FQDN or IP address for the syslog server. 2. For each node listed, check the Grid Network Subnet List, the Admin Networks Subnet Lists, and the Client Network gateways. Confirm these are configured to route traffic to the syslog server over the expected network interface and gateway (Grid, Admin, or Client).
Host unreachable	The syslog server is on a directly attached subnet (subnet used by the listed nodes for their Grid, Admin, or Client IP addresses). The nodes attempted to send test messages, but did not receive responses to ARP requests for the syslog server's MAC address. 1. Check that you entered the correct FQDN or IP address for the syslog server. 2. Check that the host running the syslog service is up.
Connection timed out	A TCP/TLS connection attempt was made, but no response was received from the syslog server for a long time. There might be a routing misconfiguration or a firewall might be dropping traffic without sending any response (a common configuration). 1. Check that you entered the correct FQDN or IP address for the syslog server. 2. For each node listed, check the Grid Network Subnet List, the Admin Networks Subnet Lists, and the Client Network gateways. Confirm these are configured to route traffic to the syslog server using the network interface and gateway (Grid, Admin, or Client) over which you expect the syslog server to be reached. 3. Confirm that a firewall is not blocking access to TCP/TLS connections from the nodes listed to the IP and port of the syslog server.

Error message	Description and recommended actions
Connection closed by partner	A TCP connection to the syslog server was successfully established but was later closed. Reasons for this might include: • The syslog server might have been restarted or rebooted. • The node and the syslog server might have different TCP/TLS settings. • An intermediate firewall might be closing idle TCP connections. • A non-syslog server listening on the syslog server port might have closed the connection. To resolve this issue: 1. Check that you entered the correct FQDN or IP address, port, and protocol for the syslog server. 2. If you are using TLS, confirm the syslog server is also using TLS. If you are using TCP, confirm the syslog server is also using TCP. 3. Check that an intermediate firewall is not configured to close idle TCP connections.
TLS certificate error	The server certificate received from the syslog server was not compatible with the CA certificate bundle and client certificate you provided. 1. Confirm that the CA certificate bundle and client certificate (if any) are compatible with the server certificate on the syslog server. 2. Confirm that the identities in the server certificate from the syslog server include the expected IP or FQDN values.
Forwarding suspended	Syslog records are no longer being forwarded to the syslog server and StorageGRID is unable to detect the reason. Review the debugging logs provided with this error to attempt to determine the root cause.
TLS session terminated	The syslog server terminated the TLS session and StorageGRID is unable to detect the reason. 1. Review the debugging logs provided with this error to attempt to determine the root cause. 2. Check that you entered the correct FQDN or IP address, port, and protocol for the syslog server. 3. If you are using TLS, confirm the syslog server is also using TLS. If you are using TCP, confirm the syslog server is also using TCP. 4. Confirm that the CA certificate bundle and client certificate (if any) are compatible with the server certificate from the syslog server. 5. Confirm that the identities in the server certificate from the syslog server include the expected IP or FQDN values.

Error message	Description and recommended actions
Results query failed	The Admin Node used for syslog server configuration and testing is unable to request test results from the nodes listed. One or more nodes might be down. 1. Follow standard troubleshooting steps to ensure that the nodes are online and all expected services are running. 2. Restart the miscd service on the nodes listed.

Review audit logs

Audit messages and logs

These instructions contain information about the structure and content of StorageGRID audit messages and audit logs. You can use this information to read and analyze the audit trail of system activity.

These instructions are for administrators responsible for producing reports of system activity and usage that require analysis of the StorageGRID system's audit messages.

To use the text log file, you must have access to the configured audit share on the Admin Node.

For information about configuring audit message levels and using an external syslog server, see [Configure audit messages and log destinations](#).

Audit message flow and retention

All StorageGRID services generate audit messages during normal system operation. You should understand how these audit messages move through the StorageGRID system to the `audit.log` file.


Audit message flow

Audit messages are processed by Admin Nodes and by those Storage Nodes that have an Administrative Domain Controller (ADC) service.

As shown in the audit message flow diagram, each StorageGRID node sends its audit messages to one of the ADC services at the data center site. The ADC service is automatically enabled for the first three Storage Nodes installed at each site.

In turn, each ADC service acts as a relay and sends its collection of audit messages to every Admin Node in the StorageGRID system, which gives each Admin Node a complete record of system activity.


Each Admin Node stores audit messages in text log files; the active log file is named `audit.log`.



Audit message retention

StorageGRID uses a copy-and-delete process to ensure that no audit messages are lost before they can be written to the audit log.

When a node generates or relays an audit message, the message is stored in an audit message queue on the system disk of the grid node. A copy of the message is always held in an audit message queue until the message is written to the audit log file in the Admin Node's `/var/local/log` directory. This helps prevent loss of an audit message during transport.



The audit message queue can temporarily increase due to network connectivity issues or insufficient audit capacity. As the queues increase, they consume more of the available space in each node's `/var/local/` directory. If the issue persists and a node's audit message directory becomes too full, the individual nodes will prioritize processing their backlog and become temporarily unavailable for new messages.

Specifically, you might see the following behaviors:

- If the `/var/local/log` directory used by an Admin Node becomes full, the Admin Node will be flagged as unavailable to new audit messages until the directory is no longer full. S3 client requests aren't affected. The XAMS (Unreachable Audit Repositories) alarm is triggered when an audit repository is unreachable.
- If the `/var/local/` directory used by a Storage Node with the ADC service becomes 92% full, the node will be flagged as unavailable to audit messages until the directory is only 87% full. S3 client requests to other nodes aren't affected. The NRLY (Available Audit Relays) alarm is triggered when audit relays are unreachable.



If there are no available Storage Nodes with the ADC service, the Storage Nodes store the audit messages locally in the `/var/local/log/localaudit.log` file.

- If the `/var/local/` directory used by a Storage Node becomes 85% full, the node will start refusing S3 client requests with `503 Service Unavailable`.

The following types of issues can cause audit message queues to grow very large:

- The outage of an Admin Node or a Storage Node with the ADC service. If one of the system's nodes is

down, the remaining nodes might become backlogged.

- A sustained activity rate that exceeds the audit capacity of the system.
- The `/var/local/` space on an ADC Storage Node becoming full for reasons unrelated to audit messages. When this happens, the node stops accepting new audit messages and prioritizes its current backlog, which can cause backlogs on other nodes.

Large audit queue alert and Audit Messages Queued (AMQS) alarm

To help you monitor the size of audit message queues over time, the **Large audit queue** alert and the legacy AMQS alarm are triggered when the number of messages in a Storage Node queue or Admin Node queue reaches certain thresholds.

If the **Large audit queue** alert or the legacy AMQS alarm is triggered, start by checking the load on the system—if there have been a significant number of recent transactions, the alert and the alarm should resolve over time and can be ignored.

If the alert or alarm persists and increases in severity, view a chart of the queue size. If the number is steadily increasing over hours or days, the audit load has likely exceeded the audit capacity of the system. Reduce the client operation rate or decrease the number of audit messages logged by changing the audit level for Client Writes and Client Reads to Error or Off. See [Configure audit messages and log destinations](#).

Duplicate messages

The StorageGRID system takes a conservative approach if a network or node failure occurs. For this reason, duplicate messages might exist in the audit log.

Access audit log file

The audit share contains the active `audit.log` file and any compressed audit log files. You can access audit log files directly from the command line of the Admin Node.

Before you begin

- You have [specific access permissions](#).
- You must have the `Passwords.txt` file.
- You must know the IP address of an Admin Node.

Steps

1. Log in to an Admin Node:
 - a. Enter the following command: `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 `#`.

2. Go to the directory containing the audit log files:

```
cd /var/local/log
```

3. View the current or a saved audit log file, as required.

Audit log file rotation

Audit logs files are saved to an Admin Node's `/var/local/log` directory. The active audit log files are named `audit.log`.



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

Once a day, the active `audit.log` file is saved, and a new `audit.log` file is started. The name of the saved file indicates when it was saved, in the format `yyyy-mm-dd.txt`. If more than one audit log is created in a single day, the file names use the date the file was saved, appended by a number, in the format `yyyy-mm-dd.txt.n`. For example, `2018-04-15.txt` and `2018-04-15.txt.1` are the first and second log files created and saved on 15 April 2018.

After a day, the saved file is compressed and renamed, in the format `yyyy-mm-dd.txt.gz`, which preserves the original date. Over time, this results in the consumption of storage allocated for audit logs on the Admin Node. A script monitors the audit log space consumption and deletes log files as necessary to free space in the `/var/local/log` directory. Audit logs are deleted based on the date they were created, with the oldest being deleted first. You can monitor the script's actions in the following file: `/var/local/log/manage-audit.log`.

This example shows the active `audit.log` file, the previous day's file (`2018-04-15.txt`), and the compressed file for the prior day (`2018-04-14.txt.gz`).

```
audit.log
2018-04-15.txt
2018-04-14.txt.gz
```

Audit log file format

Audit log file format

The audit log files are found on every Admin Node and contain a collection of individual audit messages.

Each audit message contains the following:

- The Coordinated Universal Time (UTC) of the event that triggered the audit message (ATIM) in ISO 8601 format, followed by a space:

`YYYY-MM-DDTHH:MM:SS.UUUUUU`, where `UUUUUU` are microseconds.

- The audit message itself, enclosed within square brackets and beginning with `AUDT`.

The following example shows three audit messages in an audit log file (line breaks added for readability). These messages were generated when a tenant created an S3 bucket and added two objects to that bucket.

2019-08-07T18:43:30.247711

```
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991681][TIME(UI64):73520][SAI
P(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCkNB8M3MTWNt-
PhoTDwB9JOk7PtyLkQmA=="][SUSR(CSTR):"urn:sgws:identity::175300642415970547
18:root"]
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"buc
ket1"][AVER(UI32):10][ATIM(UI64):1565203410247711]
[ATYP(FC32):SPUT][ANID(UI32):12454421][AMID(FC32):S3RQ][ATID(UI64):7074142
142472611085]]
```

2019-08-07T18:43:30.783597

```
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991696][TIME(UI64):120713][SA
IP(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCkNB8M3MTWNt-
PhoTDwB9JOk7PtyLkQmA=="][SUSR(CSTR):"urn:sgws:identity::175300642415970547
18:root"]
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"buc
ket1"][S3KY(CSTR):"fh-small-0"]
[CBID(UI64):0x779557A069B2C037][UUID(CSTR):"94BA6949-38E1-4B0C-BC80-
EB44FB4FCC7F"][CSIZ(UI64):1024][AVER(UI32):10]
[ATIM(UI64):1565203410783597][ATYP(FC32):SPUT][ANID(UI32):12454421][AMID(F
C32):S3RQ][ATID(UI64):8439606722108456022]]
```

2019-08-07T18:43:30.784558

```
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991693][TIME(UI64):121666][SA
IP(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCkNB8M3MTWNt-
PhoTDwB9JOk7PtyLkQmA=="][SUSR(CSTR):"urn:sgws:identity::175300642415970547
18:root"]
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"buc
ket1"][S3KY(CSTR):"fh-small-2000"]
[CBID(UI64):0x180CBD8E678EED17][UUID(CSTR):"19CE06D0-D2CF-4B03-9C38-
E578D66F7ADD"][CSIZ(UI64):1024][AVER(UI32):10]
[ATIM(UI64):1565203410784558][ATYP(FC32):SPUT][ANID(UI32):12454421][AMID(F
C32):S3RQ][ATID(UI64):13489590586043706682]]
```

In their default format, the audit messages in the audit log files aren't easy to read or interpret. You can use the [audit-explain tool](#) to obtain simplified summaries of the audit messages in the audit log. You can use the [audit-sum tool](#) to summarize how many write, read, and delete operations were logged and how long these operations took.

Use audit-explain tool

You can use the `audit-explain` tool to translate the audit messages in the audit log in to an easy-to-read format.

Before you begin

- You have [specific access permissions](#).
- You must have the `Passwords.txt` file.
- You must know the IP address of the primary Admin Node.

About this task

The `audit-explain` tool, available on the primary Admin Node, provides simplified summaries of the audit messages in an audit log.



The `audit-explain` tool is primarily intended for use by technical support during troubleshooting operations. Processing `audit-explain` queries can consume a large amount of CPU power, which might impact StorageGRID operations.

This example shows typical output from the `audit-explain` tool. These four [SPUT](#) audit messages were generated when the S3 tenant with account ID 92484777680322627870 used S3 PUT requests to create a bucket named "bucket1" and add three objects to that bucket.

```
SPUT S3 PUT bucket bucket1 account:92484777680322627870 usec:124673
SPUT S3 PUT object bucket1/part1.txt tenant:92484777680322627870
cbid:9DCB157394F99FE5 usec:101485
SPUT S3 PUT object bucket1/part2.txt tenant:92484777680322627870
cbid:3CFBB07AB3D32CA9 usec:102804
SPUT S3 PUT object bucket1/part3.txt tenant:92484777680322627870
cbid:5373D73831ECC743 usec:93874
```

The `audit-explain` tool can do the following:

- Process plain or compressed audit logs. For example:

```
audit-explain audit.log
```

```
audit-explain 2019-08-12.txt.gz
```

- Process multiple files simultaneously. For example:

```
audit-explain audit.log 2019-08-12.txt.gz 2019-08-13.txt.gz
```

```
audit-explain /var/local/log/*
```

- Accept input from a pipe, which allows you to filter and preprocess the input using the `grep` command or other means. For example:

```
grep SPUT audit.log | audit-explain
```

```
grep bucket-name audit.log | audit-explain
```

Because audit logs can be very large and slow to parse, you can save time by filtering parts that you want to look at and running `audit-explain` on the parts, instead of the entire file.



The `audit-explain` tool does not accept compressed files as piped input. To process compressed files, provide their file names as command-line arguments, or use the `zcat` tool to decompress the files first. For example:

```
zcat audit.log.gz | audit-explain
```

Use the `help` (`-h`) option to see the available options. For example:

```
$ audit-explain -h
```

Steps

1. Log in to the primary Admin Node:

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

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

2. Enter the following command, where `/var/local/log/audit.log` represents the name and the location of the file or files you want to analyze:

```
$ audit-explain /var/local/log/audit.log
```

The `audit-explain` tool prints human-readable interpretations of all messages in the specified file or files.



To reduce line lengths and to aid readability, timestamps aren't shown by default. If you want to see the timestamps, use the `timestamp` (`-t`) option.

Use audit-sum tool

You can use the `audit-sum` tool to count the write, read, head, and delete audit messages and to see the minimum, maximum, and average time (or size) for each operation type.

Before you begin

- You have [specific access permissions](#).
- You must have the `Passwords.txt` file.
- You must know the IP address of the primary Admin Node.

About this task

The `audit-sum` tool, available on the primary Admin Node, summarizes how many write, read, and delete operations were logged and how long these operations took.



The `audit-sum` tool is primarily intended for use by technical support during troubleshooting operations. Processing `audit-sum` queries can consume a large amount of CPU power, which might impact StorageGRID operations.

This example shows typical output from the `audit-sum` tool. This example shows how long protocol operations took.

```

message group      count      min(sec)      max(sec)
average(sec)
=====
=====
IDEL                274
SDEL               213371      0.004         20.934
0.352
SGET               201906      0.010         1740.290
1.132
SHEA               22716       0.005          2.349
0.272
SPUT              1771398      0.011         1770.563
0.487

```

The `audit-sum` tool provides counts and times for the following S3, Swift, and ILM audit messages in an audit log.



Audit codes are removed from the product and documentation as features are deprecated. If you encounter an audit code that is not listed here, check the previous versions of this topic for older SG releases. For example, [StorageGRID 11.8 Using audit sum tool documentation](#).

Code	Description	Refer to
IDEL	ILM Initiated Delete: Logs when ILM starts the process of deleting an object.	IDEL: ILM Initiated Delete
SDEL	S3 DELETE: Logs a successful transaction to delete an object or bucket.	SDEL: S3 DELETE
SGET	S3 GET: Logs a successful transaction to retrieve an object or list the objects in a bucket.	SGET: S3 GET
SHEA	S3 HEAD: Logs a successful transaction to check for the existence of an object or bucket.	SHEA: S3 HEAD
SPUT	S3 PUT: Logs a successful transaction to create a new object or bucket.	SPUT: S3 PUT
WDEL	Swift DELETE: Logs a successful transaction to delete an object or container.	WDEL: Swift DELETE

Code	Description	Refer to
WGET	Swift GET: Logs a successful transaction to retrieve an object or list the objects in a container.	WGET: Swift GET
WHEA	Swift HEAD: Logs a successful transaction to check for the existence of an object or container.	WHEA: Swift HEAD
WPUT	Swift PUT: Logs a successful transaction to create a new object or container.	WPUT: Swift PUT

The `audit-sum` tool can do the following:

- Process plain or compressed audit logs. For example:

```
audit-sum audit.log

audit-sum 2019-08-12.txt.gz
```

- Process multiple files simultaneously. For example:

```
audit-sum audit.log 2019-08-12.txt.gz 2019-08-13.txt.gz

audit-sum /var/local/log/*
```

- Accept input from a pipe, which allows you to filter and preprocess the input using the `grep` command or other means. For example:

```
grep WGET audit.log | audit-sum

grep bucket1 audit.log | audit-sum

grep SPUT audit.log | grep bucket1 | audit-sum
```



This tool does not accept compressed files as piped input. To process compressed files, provide their file names as command-line arguments, or use the `zcat` tool to decompress the files first. For example:

```
audit-sum audit.log.gz

zcat audit.log.gz | audit-sum
```

You can use command-line options to summarize operations on buckets separately from operations on objects or to group message summaries by bucket name, by time period, or by target type. By default, the summaries show the minimum, maximum, and average operation time, but you can use the `size (-s)` option to look at object size instead.

Use the `help (-h)` option to see the available options. For example:

```
$ audit-sum -h
```


Steps

- 1. Log in to the primary Admin Node:
 - a. Enter the following command: `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 `#`.

- 2. If you want to analyze all messages related to write, read, head, and delete operations, follow these steps:
 - a. Enter the following command, where `/var/local/log/audit.log` represents the name and the location of the file or files you want to analyze:

```
$ audit-sum /var/local/log/audit.log
```

This example shows typical output from the `audit-sum` tool. This example shows how long protocol operations took.

message group	count	min(sec)	max(sec)
average(sec)			
=====	=====	=====	=====
=====			
IDEL	274		
SDEL	213371	0.004	20.934
0.352			
SGET	201906	0.010	1740.290
1.132			
SHEA	22716	0.005	2.349
0.272			
SPUT	1771398	0.011	1770.563
0.487			

In this example, SGET (S3 GET) operations are the slowest on average at 1.13 seconds, but SGET and SPUT (S3 PUT) operations both show long worst-case times of about 1,770 seconds.

- b. To show the slowest 10 retrieval operations, use the `grep` command to select only SGET messages and add the long output option (`-l`) to include object paths:

```
grep SGET audit.log | audit-sum -l
```

The results include the type (object or bucket) and path, which allows you to grep the audit log for other messages relating to these particular objects.

```

Total:          201906 operations
Slowest:        1740.290 sec
Average:         1.132 sec
Fastest:         0.010 sec
Slowest operations:
    time(usec)      source ip      type      size(B) path
    =====
1740289662    10.96.101.125      object    5663711385
backup/r9010aQ8JB-1566861764-4519.iso
1624414429    10.96.101.125      object    5375001556
backup/r9010aQ8JB-1566861764-6618.iso
1533143793    10.96.101.125      object    5183661466
backup/r9010aQ8JB-1566861764-4518.iso
70839         10.96.101.125      object      28338
bucket3/dat.1566861764-6619
68487         10.96.101.125      object      27890
bucket3/dat.1566861764-6615
67798         10.96.101.125      object      27671
bucket5/dat.1566861764-6617
67027         10.96.101.125      object      27230
bucket5/dat.1566861764-4517
60922         10.96.101.125      object      26118
bucket3/dat.1566861764-4520
35588         10.96.101.125      object      11311
bucket3/dat.1566861764-6616
23897         10.96.101.125      object      10692
bucket3/dat.1566861764-4516

```

From this example output, you can see that the three slowest S3 GET requests were for objects about 5 GB in size, which is much larger than the other objects. The large size accounts for the slow worst-case retrieval times.

3. If you want to determine what sizes of objects are being ingested into and retrieved from your grid, use the size option (-s):

```
audit-sum -s audit.log
```

message group average (MB)	count	min (MB)	max (MB)
=====	=====	=====	=====
IDEL 1654.502	274	0.004	5000.000
SDEL 1.695	213371	0.000	10.504
SGET 14.920	201906	0.000	5000.000
SHEA 2.967	22716	0.001	10.504
SPUT 2.495	1771398	0.000	5000.000

In this example, the average object size for SPUT is under 2.5 MB, but the average size for SGET is much larger. The number of SPUT messages is much higher than the number of SGET messages, indicating that most objects are never retrieved.

4. If you want to determine if retrievals were slow yesterday:

- a. Issue the command on the appropriate audit log and use the group-by-time option (`-gt`), followed by the time period (for example, 15M, 1H, 10S):

```
grep SGET audit.log | audit-sum -gt 1H
```

message group average(sec)	count	min(sec)	max(sec)
=====	=====	=====	=====
2019-09-05T00 1.254	7591	0.010	1481.867
2019-09-05T01 1.115	4173	0.011	1740.290
2019-09-05T02 1.562	20142	0.011	1274.961
2019-09-05T03 1.254	57591	0.010	1383.867
2019-09-05T04 1.405	124171	0.013	1740.290
2019-09-05T05 1.562	420182	0.021	1274.511
2019-09-05T06 5.562	1220371	0.015	6274.961
2019-09-05T07 2.002	527142	0.011	1974.228
2019-09-05T08 1.105	384173	0.012	1740.290
2019-09-05T09 1.354	27591	0.010	1481.867

These results show that S3 GET traffic spiked between 06:00 and 07:00. The max and average times are both considerably higher at these times as well, and they did not ramp up gradually as the count increased. This suggests that capacity was exceeded somewhere, perhaps in the network or in the grid's ability to process requests.

- b. To determine what size objects were being retrieved each hour yesterday, add the size option (`-s`) to the command:

```
grep SGET audit.log | audit-sum -gt 1H -s
```

message group average(B)	count	min(B)	max(B)
=====	=====	=====	=====
2019-09-05T00 1.976	7591	0.040	1481.867
2019-09-05T01 2.062	4173	0.043	1740.290
2019-09-05T02 2.303	20142	0.083	1274.961
2019-09-05T03 1.182	57591	0.912	1383.867
2019-09-05T04 1.528	124171	0.730	1740.290
2019-09-05T05 2.398	420182	0.875	4274.511
2019-09-05T06 51.328	1220371	0.691	5663711385.961
2019-09-05T07 2.147	527142	0.130	1974.228
2019-09-05T08 1.878	384173	0.625	1740.290
2019-09-05T09 1.354	27591	0.689	1481.867

These results indicate that some very large retrievals occurred when the overall retrieval traffic was at its maximum.

- c. To see more detail, use the [audit-explain tool](#) to review all the SGET operations during that hour:

```
grep 2019-09-05T06 audit.log | grep SGET | audit-explain | less
```

If the output of the `grep` command is expected to be many lines, add the `less` command to show the contents of the audit log file one page (one screen) at a time.

5. If you want to determine if SPUT operations on buckets are slower than SPUT operations for objects:

- a. Start by using the `-go` option, which groups messages for object and bucket operations separately:

```
grep SPUT sample.log | audit-sum -go
```

message group average(sec)	count	min(sec)	max(sec)
=====	=====	=====	=====
SPUT.bucket 0.125	1	0.125	0.125
SPUT.object 0.236	12	0.025	1.019

The results show that SPUT operations for buckets have different performance characteristics than SPUT operations for objects.

- b. To determine which buckets have the slowest SPUT operations, use the `-gb` option, which groups messages by bucket:

```
grep SPUT audit.log | audit-sum -gb
```

message group average(sec)	count	min(sec)	max(sec)
=====	=====	=====	=====
SPUT.cho-non-versioning 1.571	71943	0.046	1770.563
SPUT.cho-versioning 1.415	54277	0.047	1736.633
SPUT.cho-west-region 1.329	80615	0.040	55.557
SPUT.ltd002 0.361	1564563	0.011	51.569

- c. To determine which buckets have the largest SPUT object size, use both the `-gb` and the `-s` options:

```
grep SPUT audit.log | audit-sum -gb -s
```

message group average (B)	count	min (B)	max (B)
=====	=====	=====	=====
SPUT.cho-non-versioning 21.672	71943	2.097	5000.000
SPUT.cho-versioning 21.120	54277	2.097	5000.000
SPUT.cho-west-region 14.433	80615	2.097	800.000
SPUT.ldt002 0.352	1564563	0.000	999.972

Audit message format

Audit message format

Audit messages exchanged within the StorageGRID system include standard information common to all messages and specific content describing the event or activity being reported.

If the summary information provided by the [audit-explain](#) and [audit-sum](#) tools is insufficient, refer to this section to understand the general format of all audit messages.

The following is an example audit message as it might appear in the audit log file:

```
2014-07-17T03:50:47.484627
[AUDT:[RSLT(FC32):VRGN][AVER(UI32):10][ATIM(UI64):1405569047484627][ATYP(FC32):SYSU][ANID(UI32):11627225][AMID(FC32):ARNI][ATID(UI64):9445736326500603516]]
```

Each audit message contains a string of attribute elements. The entire string is enclosed in brackets ([]), and each attribute element in the string has the following characteristics:

- Enclosed in brackets []
- Introduced by the string AUDT, which indicates an audit message
- Without delimiters (no commas or spaces) before or after
- Terminated by a line feed character \n

Each element includes an attribute code, a data type, and a value that are reported in this format:

```
[ATTR(type):value][ATTR(type):value]...
[ATTR(type):value]\n
```

The number of attribute elements in the message depends on the event type of the message. The attribute elements aren't listed in any particular order.

The following list describes the attribute elements:

- **ATTR** is a four-character code for the attribute being reported. There are some attributes that are common to all audit messages and others that are event-specific.
- **type** is a four-character identifier of the programming data type of the value, such as UI64, FC32, and so on. The type is enclosed in parentheses ().
- **value** is the content of the attribute, typically a numeric or text value. Values always follow a colon (:). Values of data type CSTR are surrounded by double quotes " ".

Data types

Different data types are used to store information in audit messages.

Type	Description
UI32	Unsigned long integer (32 bits); it can store the numbers 0 to 4,294,967,295.
UI64	Unsigned double long integer (64 bits); it can store the numbers 0 to 18,446,744,073,709,551,615.
FC32	Four-character constant; a 32-bit unsigned integer value represented as four ASCII characters such as "ABCD."
IPAD	Used for IP addresses.
CSTR	A variable-length array of UTF-8 characters. Characters can be escaped with the following conventions: • Backslash is \. • Carriage return is \r. • Double quotes is \". • Line feed (new line) is \n. • Characters can be replaced by their hexadecimal equivalents (in the format \xHH, where HH is the hexadecimal value representing the character).

Event-specific data

Each audit message in the audit log records data specific to a system event.

Following the opening [AUDT: container that identifies the message itself, the next set of attributes provide information about the event or action described by the audit message. These attributes are highlighted in the following example:


```

2018-12-05T08:24:45.921845 [AUDT:*[RSLT(FC32):SUCS]*
\[TIME(UI64):11454]\[SAIP(IPAD):"10.224.0.100"]\[S3AI(CSTR):"60025621595611246499"]
\[SACC(CSTR):"account"]\[S3AK(CSTR):"SGKH4_Nc8SO1H6w3w0nCOFCGgk__E6dYzKlumRs
KJA==" ]
\[SUSR(CSTR):"urn:sgws:identity::60025621595611246499:root"]
\[SBAI(CSTR):"60025621595611246499"]\[SBAC(CSTR):"account"]\[S3BK(CSTR):"bucket"]
\[S3KY(CSTR):"object"]\[CBID(UI64):0xCC128B9B9E428347]
\[UUID(CSTR):"B975D2CE-E4DA-4D14-8A23-
1CB4B83F2CD8"]\[CSIZ(UI64):30720][AVER(UI32):10]
\[ATIM(UI64):1543998285921845][ATYP(FC32):SHEA][ANID(UI32):12281045][AMID(FC32):S3RQ]
\[ATID(UI64):15552417629170647261]]

```

The `ATYP` element (underlined in the example) identifies which event generated the message. This example message includes the `SHEA` message code (`[ATYP(FC32):SHEA]`), indicating it was generated by a successful S3 HEAD request.

Common elements in audit messages

All audit messages contain the common elements.

Code	Type	Description
AMID	FC32	Module ID: A four-character identifier of the module ID that generated the message. This indicates the code segment within which the audit message was generated.
ANID	UI32	Node ID: The grid node ID assigned to the service that generated the message. Each service is allocated a unique identifier at the time the StorageGRID system is configured and installed. This ID can't be changed.
ASES	UI64	Audit Session Identifier: In previous releases, this element indicated the time at which the audit system was initialized after the service started up. This time value was measured in microseconds since the operating system epoch (00:00:00 UTC on 1 January, 1970). Note: This element is obsolete and no longer appears in audit messages.
ASQN	UI64	Sequence Count: In previous releases, this counter was incremented for each generated audit message on the grid node (ANID) and reset to zero at service restart. Note: This element is obsolete and no longer appears in audit messages.
ATID	UI64	Trace ID: An identifier that is shared by the set of messages that were triggered by a single event.

Code	Type	Description
ATIM	UI64	Timestamp: The time the event was generated that triggered the audit message, measured in microseconds since the operating system epoch (00:00:00 UTC on 1 January, 1970). Note that most available tools for converting the timestamp to local date and time are based on milliseconds. Rounding or truncation of the logged timestamp might be required. The human-readable time that appears at the beginning of the audit message in the <code>audit.log</code> file is the ATIM attribute in ISO 8601 format. The date and time are represented as <code>YYYY-MMDDTHH:MM:SS.UUUUUU</code>, where the <code>T</code> is a literal string character indicating the beginning of the time segment of the date. <code>UUUUUU</code> are microseconds.
ATYP	FC32	Event Type: A four-character identifier of the event being logged. This governs the "payload" content of the message: the attributes that are included.
AVER	UI32	Version: The version of the audit message. As the StorageGRID software evolves, new versions of services might incorporate new features in audit reporting. This field enables backward compatibility in the AMS service to process messages from older versions of services.
RSLT	FC32	Result: The result of event, process, or transaction. If is not relevant for a message, <code>NONE</code> is used rather than <code>SUCS</code> so that the message is not accidentally filtered.

Audit message examples

You can find detailed information in each audit message. All audit messages use the same format.

The following is an example audit message as it might appear in the `audit.log` file:

```
2014-07-17T21:17:58.959669
[AUDT: [RSLT (FC32) : SUCS] [TIME (UI64) : 246979] [S3AI (CSTR) : "bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"] [
S3AK (CSTR) : "UJXDKKQOXB7YARDS71Q2"] [S3BK (CSTR) : "s3small11"] [S3K
Y (CSTR) : "hello1"] [CBID (UI64) : 0x50C4F7AC2BC8EDF7] [CSIZ (UI64) : 0
] [AVER (UI32) : 10] [ATIM (UI64) : 1405631878959669] [ATYP (FC32) : SPUT
] [ANID (UI32) : 12872812] [AMID (FC32) : S3RQ] [ATID (UI64) : 1579224144
102530435]]
```

The audit message contains information about the event being recorded, as well as information about the audit message itself.

To identify which event is recorded by the audit message, look for the ATYP attribute (highlighted below):

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"] [
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"] [S3BK(CSTR):"s3small11"] [S3K
Y(CSTR):"hello1"] [CBID(UI64):0x50C4F7AC2BC8EDF7] [CSIZ(UI64):0
] [AVER(UI32):10] [ATIM(UI64):1405631878959669] [ATYP(FC32):SP
UT] [ANID(UI32):12872812] [AMID(FC32):S3RQ] [ATID(UI64):1579224
144102530435]]
```

The value of the ATYP attribute is SPUT. **SPUT** represents an S3 PUT transaction, which logs the ingest of an object to a bucket.

The following audit message also shows the bucket to which the object is associated:

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"] [
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"] [S3BK\ (CSTR\): "s3small11"] [S3
KY(CSTR):"hello1"] [CBID(UI64):0x50C4F7AC2BC8EDF7] [CSIZ(UI64):
0] [AVER(UI32):10] [ATIM(UI64):1405631878959669] [ATYP(FC32):SPU
T] [ANID(UI32):12872812] [AMID(FC32):S3RQ] [ATID(UI64):157922414
4102530435]]
```

To discover when the PUT event occurred, note the Universal Coordinated Time (UTC) timestamp at the beginning of the audit message. This value is a human-readable version of the ATIM attribute of the audit message itself:

```
2014-07-17T21:17:58.959669
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):246979][S3AI(CSTR):"bc644d
381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"] [
S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"] [S3BK(CSTR):"s3small11"] [S3K
Y(CSTR):"hello1"] [CBID(UI64):0x50C4F7AC2BC8EDF7] [CSIZ(UI64):0
] [AVER(UI32):10] [ATIM\ (UI64\): 1405631878959669] [ATYP(FC32):SP
UT] [ANID(UI32):12872812] [AMID(FC32):S3RQ] [ATID(UI64):15792241
44102530435]]
```

ATIM records the time, in microseconds, since the beginning of the UNIX epoch. In the example, the value 1405631878959669 translates to Thursday, 17-Jul-2014 21:17:59 UTC.

Audit messages and the object lifecycle

When are audit message generated?

Audit messages are generated each time an object is ingested, retrieved, or deleted. You

can identify these transactions in the audit log by locating S3 API-specific audit messages.

Audit messages are linked through identifiers specific to each protocol.

Protocol	Code
Linking S3 operations	S3BK (bucket), S3KY (key), or both
Linking Swift operations	WCON (container), WOBJ (object), or both
Linking internal operations	CBID (object's internal identifier)

Timing of audit messages

Because of factors such as timing differences between grid nodes, object size, and network delays, the order of audit messages generated by the different services can vary from that shown in the examples in this section.

Object ingest transactions

You can identify client ingest transactions in the audit log by locating S3 API-specific audit messages.

Not all audit messages generated during an ingest transaction are listed in the following tables. Only the messages required to trace the ingest transaction are included.

S3 ingest audit messages

Code	Name	Description	Trace	See
SPUT	S3 PUT transaction	An S3 PUT ingest transaction has completed successfully.	CBID, S3BK, S3KY	SPUT: S3 PUT
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	ORLM: Object Rules Met

Swift ingest audit messages

Code	Name	Description	Trace	See
WPUT	Swift PUT transaction	A Swift PUT ingest transaction has successfully completed.	CBID, WCON, WOBJ	WPUT: Swift PUT
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	ORLM: Object Rules Met

Example: S3 object ingest

The series of audit messages below is an example of the audit messages generated and saved to the audit log when an S3 client ingests an object to a Storage Node (LDR service).

In this example, the active ILM policy includes the Make 2 Copies ILM rule.



Not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 ingest transaction (SPUT) are listed.

This example assumes that an S3 bucket has been previously created.

SPUT: S3 PUT

The SPUT message is generated to indicate that an S3 PUT transaction has been issued to create an object in a specific bucket.

```
2017-07-
17T21:17:58.959669[AUDT:[RSLT(FC32):SUCS][TIME(UI64):25771][SAIP(IPAD):"10
.96.112.29"][S3AI(CSTR):"70899244468554783528"][SACC(CSTR):"test"][S3AK(CS
TR):"SGKHyalRU_5cLflqajtaFmxJn946lAWRJfBF33gAOg=="][SUSR(CSTR):"urn:sgws:i
dentity::70899244468554783528:root"][SBAI(CSTR):"70899244468554783528"][SB
AC(CSTR):"test"][S3BK(CSTR):"example"][S3KY(CSTR):"testobject-0-
3"][CBID\ (UI64\):0x8EF52DF8025E63A8][CSIZ(UI64):30720][AVER(UI32):10][ATIM
(UI64):150032627859669][ATYP\ (FC32\):SPUT][ANID(UI32):12086324][AMID(FC32)
:S3RQ][ATID(UI64):14399932238768197038]]
```

ORLM: Object Rules Met

The ORLM message indicates that the ILM policy has been satisfied for this object. The message includes the object's CBID and the name of the ILM rule that was applied.

For replicated objects, the LOCS field includes the LDR node ID and volume ID of the object locations.

```
2019-07-
17T21:18:31.230669[AUDT:[CBID\ (UI64\):0x50C4F7AC2BC8EDF7][RULE(CSTR):"Make
2 Copies"][STAT(FC32):DONE][CSIZ(UI64):0][UUID(CSTR):"0B344E18-98ED-4F22-
A6C8-A93ED68F8D3F"][LOCS(CSTR):"CLDI 12828634 2148730112, CLDI 12745543
2147552014"][RSLT(FC32):SUCS][AVER(UI32):10][ATYP\ (FC32\):ORLM][ATIM(UI64)
:1563398230669][ATID(UI64):15494889725796157557][ANID(UI32):13100453][AMID
(FC32):BCMS]]
```

For erasure-coded objects, the LOCS field includes the erasure-coding profile ID and the erasure coding group ID

```
2019-02-23T01:52:54.647537
[AUDT:[CBID(UI64):0xFA8ABE5B5001F7E2][RULE(CSTR):"EC_2_plus_1"][STAT(FC32):DONE][CSIZ(UI64):10000][UUID(CSTR):"E291E456-D11A-4701-8F51-D2F7CC9AFECA"][LOCS(CSTR):"CLEC 1 A471E45D-A400-47C7-86AC-12E77F229831"][RSLT(FC32):SUCS][AVER(UI32):10][ATIM(UI64):1550929974537]\[ATYP\ (FC32\):ORLM\][ANID(UI32):12355278][AMID(FC32):ILMX][ATID(UI64):4168559046473725560]]
```

The PATH field includes S3 bucket and key information or Swift container and object information, depending on which API was used.

```
2019-09-15.txt:2018-01-24T13:52:54.131559
[AUDT:[CBID(UI64):0x82704DFA4C9674F4][RULE(CSTR):"Make 2 Copies"][STAT(FC32):DONE][CSIZ(UI64):3145729][UUID(CSTR):"8C1C9CAC-22BB-4880-9115-CE604F8CE687"][PATH(CSTR):"frisbee_Bucket1/GridDataTests151683676324774_1_1vf9d"][LOCS(CSTR):"CLDI 12525468, CLDI 12222978"][RSLT(FC32):SUCS][AVER(UI32):10][ATIM(UI64):1568555574559][ATYP(FC32):ORLM][ANID(UI32):12525468][AMID(FC32):OBDI][ATID(UI64):344833886538369336]]
```

Object delete transactions

You can identify object delete transactions in the audit log by locating S3 API-specific audit messages.

Not all audit messages generated during a delete transaction are listed in the following tables. Only messages required to trace the delete transaction are included.

S3 delete audit messages

Code	Name	Description	Trace	See
SDEL	S3 Delete	Request made to delete the object from a bucket.	CBID, S3KY	SDEL: S3 DELETE

Swift delete audit messages

Code	Name	Description	Trace	See
WDEL	Swift Delete	Request made to delete the object from a container, or the container.	CBID, WOBJ	WDEL: Swift DELETE

Example: S3 object deletion

When an S3 client deletes an object from a Storage Node (LDR service), an audit message is generated and saved to the audit log.



Not all audit messages generated during a delete transaction are listed in the example below. Only those related to the S3 delete transaction (SDEL) are listed.

SDEL: S3 Delete

Object deletion begins when the client sends a DeleteObject request to an LDR service. The message contains the bucket from which to delete the object and the object's S3 Key, which is used to identify the object.

```
2017-07-
17T21:17:58.959669[AUDT:[RSLT(FC32):SUCS][TIME(UI64):14316][SAIP(IPAD):"10
.96.112.29"][S3AI(CSTR):"70899244468554783528"][SACC(CSTR):"test"][S3AK(CS
TR):"SGKHyalRU_5cLflqajtaFmxJn946lAWRJfBF33gAOg=="][SUSR(CSTR):"urn:sgws:i
dentity::70899244468554783528:root"][SBAI(CSTR):"70899244468554783528"][SB
AC(CSTR):"test"]\[S3BK\CSTR\):"example"\]\[S3KY\CSTR\):"testobject-0-
7"\][CBID(UI64):0x339F21C5A6964D89][CSIZ(UI64):30720][AVER(UI32):10][ATI
M(UI64):150032627859669][ATYP\ (FC32\) :SDEL][ANID(UI32):12086324][AMID(FC32
):S3RQ][ATID(UI64):4727861330952970593]]
```

Object retrieve transactions

You can identify object retrieve transactions in the audit log by locating S3 API-specific audit messages.

Not all audit messages generated during a retrieve transaction are listed in the following tables. Only messages required to trace the retrieve transaction are included.

S3 retrieval audit messages

Code	Name	Description	Trace	See
SGET	S3 GET	Request made to retrieve an object from a bucket.	CBID, S3BK, S3KY	SGET: S3 GET

Swift retrieval audit messages

Code	Name	Description	Trace	See
WGET	Swift GET	Request made to retrieve an object from a container.	CBID, WCON, WOBJ	WGET: Swift GET

Example: S3 object retrieval

When an S3 client retrieves an object from a Storage Node (LDR service), an audit message is generated and saved to the audit log.

Note that not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 retrieval transaction (SGET) are listed.

SGET: S3 GET

Object retrieval begins when the client sends a `GetObject` request to an LDR service. The message contains the bucket from which to retrieve the object and the object's S3 Key, which is used to identify the object.

```
2017-09-20T22:53:08.782605
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):47807][SAIP(IPAD):"10.96.112.26"][S3AI(
CSTR):"43979298178977966408"][SACC(CSTR):"s3-account-
a"][S3AK(CSTR):"SGKht7GzEcu0yXhFhT_rL5mep4nJt1w75GBh-
O_FEw=="][SUSR(CSTR):"urn:sgws:identity::43979298178977966408:root"][SBAI(
CSTR):"43979298178977966408"][SBAC(CSTR):"s3-account-
a"]\[S3BK\CSTR\):"bucket-
anonymous"\]\[S3KY\CSTR\):"Hello.txt"\][CBID(UI64):0x83D70C6F1F662B02][CS
IZ(UI64):12][AVER(UI32):10][ATIM(UI64):1505947988782605]\[ATYP\ (FC32\):SGE
T\][ANID(UI32):12272050][AMID(FC32):S3RQ][ATID(UI64):17742374343649889669]
]
```

If the bucket policy allows, a client can anonymously retrieve objects, or can retrieve objects from a bucket that is owned by a different tenant account. The audit message contains information about the bucket owner's tenant account so that you can track these anonymous and cross-account requests.

In the following example message, the client sends a `GetObject` request for an object stored in a bucket that they don't own. The values for `SBAI` and `SBAC` record the bucket owner's tenant account ID and name, which differs from the tenant account ID and name of the client recorded in `S3AI` and `SACC`.

```
2017-09-20T22:53:15.876415
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):53244][SAIP(IPAD):"10.96.112.26"]\[S3AI
\CSTR\):"17915054115450519830"\]\[SACC\CSTR\):"s3-account-
b"\][S3AK(CSTR):"SGKHpoblWlP_kBkqSCbTi754Ls81BUog67I2LlSiUg=="][SUSR(CSTR)
:"urn:sgws:identity::17915054115450519830:root"]\[SBAI\CSTR\):"4397929817
8977966408"\]\[SBAC\CSTR\):"s3-account-a"\][S3BK(CSTR):"bucket-
anonymous"][S3KY(CSTR):"Hello.txt"][CBID(UI64):0x83D70C6F1F662B02][CSIZ(UI
64):12][AVER(UI32):10][ATIM(UI64):1505947995876415][ATYP(FC32):SGET][ANID(
UI32):12272050][AMID(FC32):S3RQ][ATID(UI64):6888780247515624902]]
```

Example: S3 Select on an object

When an S3 client issues an S3 Select query on an object, audit messages are generated and saved to the audit log.

Note that not all audit messages generated during a transaction are listed in the example below. Only those related to the S3 Select transaction (SelectObjectContent) are listed.

Each query results in two audit messages: one that performs the authorization of the S3 Select request (the S3SR field is set to "select") and a subsequent standard GET operation that retrieves the data from storage during processing.

```
2021-11-08T15:35:30.750038
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1636385730715700][TIME(UI64):29173][SAIP(IPAD):"192.168.7.44"][S3AI(CSTR):"63147909414576125820"][SACC(CSTR):"Tenant1636027116"][S3AK(CSTR):"AUFD1XNVZ905F3TW7KSU"][SUSR(CSTR):"urn:sgws:identity::63147909414576125820:root"][SBAI(CSTR):"63147909414576125820"][SBAC(CSTR):"Tenant1636027116"][S3BK(CSTR):"619c0755-9e38-42e0-a614-05064f74126d"][S3KY(CSTR):"SUB-EST2020_ALL.csv"][CBID(UI64):0x0496F0408A721171][UUID(CSTR):"D64B1A4A-9F01-4EE7-B133-08842A099628"][CSIZ(UI64):0][S3SR(CSTR):"select"][AVER(UI32):10][ATIM(UI64):1636385730750038][ATYP(FC32):SPOS][ANID(UI32):12601166][AMID(FC32):S3RQ][ATID(UI64):1363009709396895985]]
```

```
2021-11-08T15:35:32.604886
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1636383069486504][TIME(UI64):430690][SAIP(IPAD):"192.168.7.44"][HTRH(CSTR):"{\"x-forwarded-for\": \"unix:\"}"][S3AI(CSTR):"63147909414576125820"][SACC(CSTR):"Tenant1636027116"][S3AK(CSTR):"AUFD1XNVZ905F3TW7KSU"][SUSR(CSTR):"urn:sgws:identity::63147909414576125820:root"][SBAI(CSTR):"63147909414576125820"][SBAC(CSTR):"Tenant1636027116"][S3BK(CSTR):"619c0755-9e38-42e0-a614-05064f74126d"][S3KY(CSTR):"SUB-EST2020_ALL.csv"][CBID(UI64):0x0496F0408A721171][UUID(CSTR):"D64B1A4A-9F01-4EE7-B133-08842A099628"][CSIZ(UI64):10185581][MTME(UI64):1636380348695262][AVER(UI32):10][ATIM(UI64):1636385732604886][ATYP(FC32):SGET][ANID(UI32):12733063][AMID(FC32):S3RQ][ATID(UI64):16562288121152341130]]
```

Metadata update messages

Audit messages are generated when an S3 client updates an object's metadata.

S3 metadata update audit messages

Code	Name	Description	Trace	See
SUPD	S3 Metadata Updated	Generated when an S3 client updates the metadata for an ingested object.	CBID, S3KY, HTRH	SUPD: S3 Metadata Updated

Example: S3 metadata update

The example shows a successful transaction to update the metadata for an existing S3 object.

SUPD: S3 Metadata Update

The S3 client makes a request (SUPD) to update the specified metadata (`x-amz-meta-*`) for the S3 object (S3KY). In this example, request headers are included in the field HTRH because it has been configured as an audit protocol header (**CONFIGURATION > Monitoring > Audit and syslog server**). See [Configure audit messages and log destinations](#).

```
2017-07-11T21:54:03.157462
[AUDT:[RSLT(FC32):SUCS][TIME(UI64):17631][SAIP(IPAD):"10.96.100.254"]
[HTRH(CSTR):"{\"accept-encoding\": \"identity\", \"authorization\": \"AWS
LIUF17FGJARQHPY2E761:jul/hnZs/uNY+aVvV0lTSYhEGts=\",
\"content-length\": \"0\", \"date\": \"Tue, 11 Jul 2017 21:54:03
GMT\", \"host\": \"10.96.99.163:18082\",
\"user-agent\": \"aws-cli/1.9.20 Python/2.7.6 Linux/3.13.0-119-generic
botocore/1.3.20\",
\"x-amz-copy-source\": \"/testbkt1/testobj1\", \"x-amz-metadata-
directive\": \"REPLACE\", \"x-amz-meta-city\": \"Vancouver\"}"]
[S3AI(CSTR):"20956855414285633225"][SACC(CSTR):"acct1"][S3AK(CSTR):"SGKHyy
v9ZQqWRbJSQc5vI7mgioJwrdplShE02AUaww=="]
[SUSR(CSTR):"urn:sgws:identity::20956855414285633225:root"]
[SBAI(CSTR):"20956855414285633225"][SBAC(CSTR):"acct1"][S3BK(CSTR):"testbk
t1"]
[S3KY(CSTR):"testobj1"][CBID(UI64):0xCB1D5C213434DD48][CSIZ(UI64):10][AVER
(UI32):10]
[ATIM(UI64):1499810043157462][ATYP(FC32):SUPD][ANID(UI32):12258396][AMID(F
C32):S3RQ]
[ATID(UI64):8987436599021955788]]
```

Audit messages

Audit message descriptions

Detailed descriptions of audit messages returned by the system are listed in the following sections. Each audit message is first listed in a table that groups related messages by the class of activity that the message represents. These groupings are useful both for understanding the types of activities that are audited, and for selecting the desired type of audit message filtering.

The audit messages are also listed alphabetically by their four-character codes. This alphabetic list enables you to find information about specific messages.

The four-character codes used throughout this chapter are the ATYP values found in the audit messages as shown in the following example message:

```
2014-07-17T03:50:47.484627
```

```
\[AUDT:[RSLT(FC32):VRGN][AVER(UI32):10][ATIM(UI64):1405569047484627][ATYP\  
(FC32\):SYSU][ANID(UI32):11627225][AMID(FC32):ARNI][ATID(UI64):94457363265  
00603516]]
```

For information about setting audit message levels, changing log destinations, and using an external syslog server for your audit information, see [Configure audit messages and log destinations](#)

Audit message categories

System audit messages

The audit messages belonging to the system audit category are used for events related to the auditing system itself, grid node states, system-wide task activity (grid tasks), and service backup operations.

Code	Message title and description	See
ECMC	Missing Erasure-Coded Data Fragment: Indicates that a missing erasure-coded data fragment has been detected.	ECMC: Missing Erasure-Coded Data Fragment
ECOC	Corrupt Erasure-Coded Data Fragment: Indicates that a corrupt erasure-coded data fragment has been detected.	ECOC: Corrupt Erasure-Coded Data Fragment
ETAF	Security Authentication Failed: A connection attempt using Transport Layer Security (TLS) failed.	ETAF: Security Authentication Failed
GNRG	GNDS Registration: A service updated or registered information about itself in the StorageGRID system.	GNRG: GNDS Registration
GNUR	GNDS Unregistration: A service has unregistered itself from the StorageGRID system.	GNUR: GNDS Unregistration
GTED	Grid Task Ended: The CMN service finished processing the grid task.	GTED: Grid Task Ended
GTST	Grid Task Started: The CMN service started to process the grid task.	GTST: Grid Task Started
GTSU	Grid Task Submitted: A grid task was submitted to the CMN service.	GTSU: Grid Task Submitted
LLST	Location Lost: This audit message is generated when a location is lost.	LLST: Location Lost

Code	Message title and description	See
OLST	Object Lost: A requested object cannot be located within the StorageGRID system.	OLST: System Detected Lost Object
SADD	Security Audit Disable: Audit message logging was turned off.	SADD: Security Audit Disable
SADE	Security Audit Enable: Audit message logging has been restored.	SADE: Security Audit Enable
SVRF	Object Store Verify Fail: A content block failed verification checks.	SVRF: Object Store Verify Fail
SVRU	Object Store Verify Unknown: Unexpected object data detected in the object store.	SVRU: Object Store Verify Unknown
SYSD	Node Stop: A shutdown was requested.	SYSD: Node Stop
SYST	Node Stopping: A service initiated a graceful stop.	SYST: Node Stopping
SYSU	Node Start: A service started; the nature of the previous shutdown is indicated in the message.	SYSU: Node Start

Object storage audit messages

The audit messages belonging to the object storage audit category are used for events related to the storage and management of objects within the StorageGRID system. These include object storage and retrievals, grid-node to grid-node transfers, and verifications.



Audit codes are removed from the product and documentation as features are deprecated. If you encounter an audit code that is not listed here, check the previous versions of this topic for older SG releases. For example, [StorageGRID 11.8 object storage audit messages](#).

Code	Description	See
BROR	Bucket Read Only Request: A bucket entered or exited read-only mode.	BROR: Bucket Read Only Request
CBSE	Object Send End: The source entity completed a grid-node to grid-node data transfer operation.	CBSE: Object Send End
CBRE	Object Receive End: The destination entity completed a grid-node to grid-node data transfer operation.	CBRE: Object Receive End

Code	Description	See
CGRR	Cross-Grid Replication Request: StorageGRID attempted a cross-grid replication operation to replicate objects between buckets in a grid federation connection.	CGRR: Cross-Grid Replication Request
EBDL	Empty Bucket Delete: The ILM scanner deleted an object in a bucket that is deleting all objects (performing an empty bucket operation).	EBDL: Empty Bucket Delete
EBKR	Empty Bucket Request: A user sent a request to turn empty bucket on or off (that is, to delete bucket objects or to stop deleting objects).	EBKR: Empty Bucket Request
SCMT	Object Store Commit: A content block was completely stored and verified, and can now be requested.	SCMT: Object Store Commit Request
SREM	Object Store Remove: A content block was deleted from a grid node, and can no longer be requested directly.	SREM: Object Store Remove

Client read audit messages

Client read audit messages are logged when an S3 client application makes a request to retrieve an object.

Code	Description	Used by	See
S3SL	S3 Select request: Logs a completion after an S3 Select request has been returned to the client. The S3SL message can include error message and error code details. The request might not have been successful.	S3 client	S3SL: S3 Select request
SGET	S3 GET: Logs a successful transaction to retrieve an object or list the objects in a bucket. Note: If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	SGET: S3 GET
SHEA	S3 HEAD: Logs a successful transaction to check for the existence of an object or bucket.	S3 client	SHEA: S3 HEAD
WGET	Swift GET: Logs a successful transaction to retrieve an object or list the objects in a container.	Swift client	WGET: Swift GET
WHEA	Swift HEAD: Logs a successful transaction to check for the existence of an object or container.	Swift client	WHEA: Swift HEAD

Client write audit messages

Client write audit messages are logged when an S3 client application makes a request to create or modify an object.

Code	Description	Used by	See
OVWR	Object Overwrite: Logs a transaction to overwrite one object with another object.	S3 and Swift clients	OVWR: Object Overwrite
SDEL	S3 DELETE: Logs a successful transaction to delete an object or bucket. Note: If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	SDEL: S3 DELETE
SPOS	S3 POST: Logs a successful transaction to restore an object from AWS Glacier storage to a Cloud Storage Pool.	S3 client	SPOS: S3 POST
SPUT	S3 PUT: Logs a successful transaction to create a new object or bucket. Note: If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	SPUT: S3 PUT
SUPD	S3 Metadata Updated: Logs a successful transaction to update the metadata for an existing object or bucket.	S3 client	SUPD: S3 Metadata Updated
WDEL	Swift DELETE: Logs a successful transaction to delete an object or container.	Swift client	WDEL: Swift DELETE
WPUT	Swift PUT: Logs a successful transaction to create a new object or container.	Swift client	WPUT: Swift PUT

Management audit message

The Management category logs user requests to the Management API.

Code	Message title and description	See
MGAU	Management API audit message: A log of user requests.	MGAU: Management audit message

ILM audit messages

The audit messages belonging to the ILM audit category are used for events related to information lifecycle management (ILM) operations.

Code	Message title and description	See
IDEL	ILM Initiated Delete: This audit message is generated when ILM starts the process of deleting an object.	IDEL: ILM Initiated Delete
LKCU	Overwritten Object Cleanup. This audit message is generated when an overwritten object is automatically removed to free up storage space.	LKCU: Overwritten Object Cleanup
ORLM	Object Rules Met: This audit message is generated when object data is stored as specified by the ILM rules.	ORLM: Object Rules Met

Audit message reference

BROR: Bucket Read Only Request

The LDR service generates this audit message when a bucket enters or exits read-only mode. For example, a bucket enters read-only mode while all objects are being deleted.

Code	Field	Description
BKHD	Bucket UUID	The bucket ID.
BROV	Bucket read-only request value	Whether the bucket is being made read-only or is leaving the read-only state (1 = read-only, 0 = not-read-only).
BROS	Bucket read-only reason	The reason the bucket is being made read-only or leaving the read-only state. For example, emptyBucket.
S3AI	S3 tenant account ID	The ID of the tenant account that sent the request. An empty value indicates anonymous access.
S3BK	S3 bucket	The S3 bucket name.

CBRB: Object Receive Begin

During normal system operations, content blocks are continuously transferred between different nodes as data is accessed, replicated and retained. When transfer of a content block from one node to another is initiated, this message is issued by the destination entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the first sequence count requested. If successful, the transfer begins from this sequence count.
CTES	Expected End Sequence Count	Indicates the last sequence count requested. If successful, the transfer is considered complete when this sequence count has been received.
RSLT	Transfer Start Status	Status at the time the transfer was started: SUCS: Transfer started successfully.

This audit message means a node-to-node data transfer operation was initiated on a single piece of content, as identified by its Content Block Identifier. The operation requests data from "Start Sequence Count" to "Expected End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow, and when combined with storage audit messages, to verify replica counts.

CBRE: Object Receive End

When transfer of a content block from one node to another is completed, this message is issued by the destination entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.

Code	Field	Description
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.
RSLT	Transfer Result	The result of the transfer operation (from the perspective of the sending entity): SUCS: transfer successfully completed; all requested sequence counts were sent. CONL: connection lost during transfer CTMO: connection timed-out during establishment or transfer UNRE: destination node ID unreachable CRPT: transfer ended due to reception of corrupt or invalid data

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from "Start Sequence Count" to "Actual End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

CBSB: Object Send Begin

During normal system operations, content blocks are continuously transferred between different nodes as data is accessed, replicated and retained. When transfer of a content block from one node to another is initiated, this message is issued by the source entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the first sequence count requested. If successful, the transfer begins from this sequence count.
CTES	Expected End Sequence Count	Indicates the last sequence count requested. If successful, the transfer is considered complete when this sequence count has been received.
RSLT	Transfer Start Status	Status at the time the transfer was started: SUCS: transfer started successfully.

This audit message means a node-to-node data transfer operation was initiated on a single piece of content, as identified by its Content Block Identifier. The operation requests data from "Start Sequence Count" to "Expected End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow, and when combined with storage audit messages, to verify replica counts.

CBSE: Object Send End

When transfer of a content block from one node to another is completed, this message is issued by the source entity.

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node session/connection.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.

Code	Field	Description
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH: The transfer operation was requested by the sending entity. PULL: The transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.
RSLT	Transfer Result	The result of the transfer operation (from the perspective of the sending entity): SUCS: Transfer successfully completed; all requested sequence counts were sent. CONL: connection lost during transfer CTMO: connection timed-out during establishment or transfer UNRE: destination node ID unreachable CRPT: transfer ended due to reception of corrupt or invalid data

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from "Start Sequence Count" to "Actual End Sequence Count". Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

CGRR: Cross-Grid Replication Request

This message is generated when StorageGRID attempts a cross-grid replication operation to replicate objects between buckets in a grid federation connection.

Code	Field	Description
CSIZ	Object Size	The size of the object in bytes. The CSIZ attribute was introduced in StorageGRID 11.8. As a result, cross-grid replication requests spanning a StorageGRID 11.7 to 11.8 upgrade might have an inaccurate total object size.
S3AI	S3 tenant account ID	The ID of the tenant account that owns the bucket from which the object is being replicated.
GFID	Grid federation connection ID	The ID of the grid federation connection being used for cross-grid replication.
OPER	CGR operation	The type of cross-grid replication operation that was attempted: • 0 = Replicate object • 1 = Replicate multipart object • 2 = Replicate delete marker
S3BK	S3 bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name.
VSID	Version ID	The version ID of the specific version of an object that was being replicated.
RSLT	Result Code	Returns successful (SUCS) or general error (GERR).

EBDL: Empty Bucket Delete

The ILM scanner deleted an object in a bucket that is deleting all objects (performing an empty bucket operation).

Code	Field	Description
CSIZ	Object Size	The size of the object in bytes.
PATH	S3 Bucket/Key	The S3 bucket name and S3 key name.
SEGC	Container UUID	UUID of the container for the segmented object. This value is available only if the object is segmented.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.

Code	Field	Description
RSLT	Result of the delete operation	The result of event, process, or transaction. If is not relevant for a message, NONE is used rather than SUCS so that the message is not accidentally filtered.

EBKR: Empty Bucket Request

This message indicates a user sent a request to turn empty bucket on or off (that is, to delete bucket objects or to stop deleting objects).

Code	Field	Description
BUID	Bucket UUID	The bucket ID.
EBJS	Empty Bucket JSON Configuration	Contains the JSON representing the current Empty Bucket configuration.
S3AI	S3 tenant account ID	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.

ECMC: Missing Erasure-Coded Data Fragment

This audit message indicates that the system has detected a missing erasure-coded data fragment.

Code	Field	Description
VCMC	VCS ID	The name of the VCS that contains the missing chunk.
MCID	Chunk ID	The identifier of the missing erasure-coded fragment.
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this particular message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

ECOC: Corrupt Erasure-Coded Data Fragment

This audit message indicates that the system has detected a corrupt erasure-coded data fragment.

Code	Field	Description
VCCO	VCS ID	The name of the VCS that contains the corrupt chunk.

Code	Field	Description
VLID	Volume ID	The RangeDB Volume that contains the corrupt erasure-coded fragment.
CCID	Chunk ID	The identifier of the corrupt erasure-coded fragment.
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this particular message. 'NONE' is used rather than 'SUCS' so that this message is not filtered.

ETAF: Security Authentication Failed

This message is generated when a connection attempt using Transport Layer Security (TLS) has failed.

Code	Field	Description
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection over which the authentication failed.
RUID	User Identity	A service dependent identifier representing the identity of the remote user.
RSLT	Reason Code	The reason for the failure: SCNI: Secure connection establishment failed. CERM: Certificate was missing. CERT: Certificate was invalid. CERE: Certificate was expired. CERR: Certificate was revoked. CSGN: Certificate signature was invalid. CSGU: Certificate signer was unknown. UCRM: User credentials were missing. UCRI: User credentials were invalid. UCRU: User credentials were disallowed. TOUT: Authentication timed out.

When a connection is established to a secure service that uses TLS, the credentials of the remote entity are verified using the TLS profile and additional logic built into the service. If this authentication fails due to invalid, unexpected, or disallowed certificates or credentials, an audit message is logged. This enables queries for

unauthorized access attempts and other security-related connection problems.

The message could result from a remote entity having an incorrect configuration, or from attempts to present invalid or disallowed credentials to the system. This audit message should be monitored to detect attempts to gain unauthorized access to the system.

GNRG: GNDS Registration

The CMN service generates this audit message when a service has updated or registered information about itself in the StorageGRID system.

Code	Field	Description
RSLT	Result	The result of the update request: • SUCS: Successful • SUNV: Service Unavailable • GERR: Other failure
GNID	Node ID	The node ID of the service that initiated the update request.
GNTTP	Device Type	The grid node's device type (for example, BLDR for an LDR service).
GNDV	Device Model version	The string identifying the grid node's device model version in the DMDL bundle.
GNGP	Group	The group to which the grid node belongs (in the context of link costs and service-query ranking).
GNIA	IP Address	The grid node's IP address.

This message is generated whenever a grid node updates its entry in the Grid Nodes Bundle.

GNUR: GNDS Unregistration

The CMN service generates this audit message when a service has unregistered information about itself from the StorageGRID system.

Code	Field	Description
RSLT	Result	The result of the update request: • SUCS: Successful • SUNV: Service Unavailable • GERR: Other failure
GNID	Node ID	The node ID of the service that initiated the update request.

GTED: Grid Task Ended

This audit message indicates that the CMN service has finished processing the specified grid task and has moved the task to the Historical table. If the result is SUCS, ABRT, or ROLF, there will be a corresponding Grid Task Started audit message. The other results indicate that processing of this grid task never started.

Code	Field	Description
TSID	Task ID	This field uniquely identifies a generated grid task and allows the grid task to be managed over its lifecycle. Note: The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.
RSLT	Result	The final status result of the grid task: • SUCS: The grid task completed successfully. • ABRT: The grid task was terminated without a rollback error. • ROLF: The grid task was terminated and was unable to complete the rollback process. • CANC: The grid task was canceled by the user before it was started. • EXPR: The grid task expired before it was started. • IVLD: The grid task was invalid. • AUTH: The grid task was unauthorized. • DUPL: The grid task was rejected as a duplicate.

GTST: Grid Task Started

This audit message indicates that the CMN service has started to process the specified grid task. The audit message immediately follows the Grid Task Submitted message for grid tasks initiated by the internal Grid Task Submission service and selected for automatic activation. For grid tasks submitted into the Pending table, this message is generated when the user starts the grid task.

Code	Field	Description
TSID	Task ID	This field uniquely identifies a generated grid task and allows the task to be managed over its lifecycle. Note: The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.
RSLT	Result	The result. This field has only one value: • SUCS: The grid task was started successfully.

GTSU: Grid Task Submitted

This audit message indicates that a grid task has been submitted to the CMN service.

Code	Field	Description
TSID	Task ID	Uniquely identifies a generated grid task and allows the task to be managed over its lifecycle. Note: The Task ID is assigned at the time that a grid task is generated, not the time that it is submitted. It is possible for a given grid task to be submitted multiple times, and in this case the Task ID field is not sufficient to uniquely link the Submitted, Started, and Ended audit messages.
TTYP	Task Type	The type of grid task.
TVER	Task Version	A number indicating the version of the grid task.
TDSC	Task Description	A human-readable description of the grid task.
VATS	Valid After Timestamp	The earliest time (UINT64 microseconds from January 1, 1970 - UNIX time) at which the grid task is valid.
VBTS	Valid Before Timestamp	The latest time (UINT64 microseconds from January 1, 1970 - UNIX time) at which the grid task is valid.
TSRC	Source	The source of the task: • TXTB: The grid task was submitted through the StorageGRID system as a signed text block. • GRID: The grid task was submitted through the internal Grid Task Submission Service.

Code	Field	Description
ACTV	Activation Type	The type of activation: • AUTO: The grid task was submitted for automatic activation. • PEND: The grid task was submitted into the pending table. This is the only possibility for the TXTB source.
RSLT	Result	The result of the submission: • SUCS: The grid task was submitted successfully. • FAIL: The task has been moved directly to the historical table.

IDEL: ILM Initiated Delete

This message is generated when ILM starts the process of deleting an object.

The IDEL message is generated in either of these situations:

- **For objects in compliant S3 buckets:** This message is generated when ILM starts the process of auto-deleting an object because its retention period has expired (assuming the auto-delete setting is enabled and legal hold is off).
- **For objects in non-compliant S3 buckets.** This message is generated when ILM starts the process of deleting an object because no placement instructions in the active ILM policies currently apply to the object.

Code	Field	Description
CBID	Content Block Identifier	The CBID of the object.
CMPA	Compliance: Auto delete	For objects in compliant S3 buckets only. 0 (false) or 1 (true), indicating whether a compliant object should be deleted automatically when its retention period ends, unless the bucket is under a legal hold.
CMPL	Compliance: Legal hold	For objects in compliant S3 buckets only. 0 (false) or 1 (true), indicating whether the bucket is currently under a legal hold.
CMPR	Compliance: Retention period	For objects in compliant S3 buckets only. The length of the object's retention period in minutes.
CTME	Compliance: Ingest time	For objects in compliant S3 buckets only. The object's ingest time. You can add the retention period in minutes to this value to determine when the object can be deleted from the bucket.
DMRK	Delete Marker Version ID	The version ID of the delete marker created when deleting an object from a versioned bucket. Operations on buckets don't include this field.

Code	Field	Description
CSIZ	Content size	The size of the object in bytes.
LOCS	Locations	The storage location of object data within the StorageGRID system. The value for LOCS is "" if the object has no locations (for example, it has been deleted). CLEC: for erasure-coded objects, the erasure-coding profile ID and the erasure coding group ID that is applied to the object's data. CLDI: for replicated objects, the LDR node ID and the volume ID of the object's location. CLNL: ARC node ID of the object's location if the object data is archived.
PATH	S3 Bucket/Key	The S3 bucket name and S3 key name.
RSLT	Result	The result of the ILM operation. SUCS: The ILM operation was successful.
RULE	Rules Label	• If an object in a compliant S3 bucket is being deleted automatically because its retention period has expired, this field is blank. • If the object is being deleted because there are no more placement instructions that currently apply to the object, this field shows the human-readable label of the last ILM rule that applied to the object.
SGRP	Site (Group)	If present, the object was deleted at the site specified, which is not the site where the object was ingested.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was deleted. Operations on buckets and objects in unversioned buckets don't include this field.

LKCU: Overwritten Object Cleanup

This message is generated when StorageGRID removes an overwritten object that previously required cleanup to free up storage space. An object is overwritten when an S3 client writes an object to a path already containing a object. The removal process occurs automatically and in the background.

Code	Field	Description
CSIZ	Content size	The size of the object in bytes.

Code	Field	Description
LTyp	Type of cleanup	<i>Internal use only.</i>
LUID	Removed Object UUID	The identifier of the object that was removed.
PATH	S3 Bucket/Key	The S3 bucket name and S3 key name.
SEGC	Container UUID	UUID of the container for the segmented object. This value is available only if the object is segmented.
UUID	Universally Unique Identifier	The identifier of the object that still exists. This value is available only if the object has not been deleted.

LKDM: Leaked Object Cleanup

This message is generated when a leaked chunk has been cleaned or deleted. A chunk can be part of a replicated object or an erasure-encoded object.

Code	Field	Description
CLOC	Chunk location	The file path of the leaked chunk that got deleted.
CTYP	Chunk type	Type of chunk: ec: Erasure-coded object chunk repl: Replicated object chunk
LTyp	Leak type	The five types of leaks that can be detected: object_leaked: Object doesn't exist in the grid location_leaked: Object exists in the grid, but found location doesn't belong to object mup_seg_leaked: Multipart upload was stopped or not completed, and the segment/part was left out segment_leaked: Parent UUID/CBID (associated container object) is valid but doesn't contain this segment no_parent: Container object is deleted, but object segment was left out and not deleted

Code	Field	Description
CTIM	Chunk create time	Time the leaked chunk was created.
UUID	Universally Unique Identifier	The identifier of the object the chunk belongs to.
CBID	Content Block Identifier	CBID of the object the leaked chunk belongs to.
CSIZ	Content size	The size of the chunk in bytes.

LLST: Location Lost

This message is generated whenever a location for an object copy (replicated or erasure-coded) can't be found.

Code	Field	Description
CBIL	CBID	The affected CBID.
ECPR	Erasure-Coding Profile	For erasure-coded object data. The ID of the erasure-coding profile used.
LTYP	Location Type	CLDI (Online): For replicated object data CLEC (Online): For erasure-coded object data CLNL (Nearline): For archived replicated object data
NOID	Source Node ID	The node ID on which the locations were lost.
PCLD	Path to replicated object	The complete path to the disk location of the lost object data. Only returned when LTYP has a value of CLDI (that is, for replicated objects). Takes the form <code>/var/local/rangedb/2/p/13/13/00oJs6X%{h{U)SeUFxE@</code>
RSLT	Result	Always NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.
TSRC	Triggering Source	USER: User triggered SYST: System triggered

Code	Field	Description
UUID	Universally Unique ID	The identifier of the affected object in the StorageGRID system.

MGAU: Management audit message

The Management category logs user requests to the Management API. Every HTTP request that is not a GET or HEAD request to a valid API URI logs a response containing the username, IP, and type of request to the API. Invalid API URIs (such as /api/v3-authorize) and invalid requests to valid API URIs are not logged.

Code	Field	Description
MDIP	Destination IP Address	The server (destination) IP address.
MDNA	Domain name	The host domain name.
MPAT	Request PATH	The request path.
MPQP	Request query parameters	The query parameters for the request.
MRBD	Request body	The content of the request body. While the response body is logged by default, the request body is logged in certain cases when the response body is empty. Because the following information is not available in the response body, it is taken from the request body for the following POST methods: • Username and account ID in POST authorize • New subnets configuration in POST /grid/grid-networks/update • New NTP servers in POST /grid/ntp-servers/update • Decommissioned server IDs in POST /grid/servers/decommission Note: Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).
MRMD	Request method	The HTTP request method: • POST • PUT • DELETE • PATCH
MRSC	Response code	The response code.

Code	Field	Description
MRSP	Response body	The content of the response (the response body) is logged by default. Note: Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).
MSIP	Source IP address	The client (source) IP address.
MUUN	User URN	The URN (uniform resource name) of the user who sent the request.
RSLT	Result	Returns successful (SUCS) or the error reported by the backend.

OLST: System Detected Lost Object

This message is generated when the DDS service can't locate any copies of an object within the StorageGRID system.

Code	Field	Description
CBID	Content Block Identifier	The CBID of the lost object.
NOID	Node ID	If available, the last known direct or near-line location of the lost object. It is possible to have just the Node ID without a Volume ID if the volume information is not available.
PATH	S3 Bucket/Key	If available, the S3 bucket name and S3 key name.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.
UUID	Universally Unique ID	The identifier of the lost object within the StorageGRID system.
VOLI	Volume ID	If available, the Volume ID of the Storage Node for the last known location of the lost object.

ORLM: Object Rules Met

This message is generated when the object is successfully stored and copied as specified by the ILM rules.



The ORLM message is not generated when an object is successfully stored by the default Make 2 Copies rule if another rule in the policy uses the Object Size advanced filter.

Code	Field	Description
BUID	Bucket Header	Bucket ID field. Used for internal operations. Appears only if STAT is PRGD.
CBID	Content Block Identifier	The CBID of the object.
CSIZ	Content size	The size of the object in bytes.
LOCS	Locations	The storage location of object data within the StorageGRID system. The value for LOCS is "" if the object has no locations (for example, it has been deleted). CLEC: for erasure-coded objects, the erasure-coding profile ID and the erasure coding group ID that is applied to the object's data. CLDI: for replicated objects, the LDR node ID and the volume ID of the object's location. CLNL: ARC node ID of the object's location if the object data is archived.
PATH	S3 Bucket/Key	The S3 bucket name and S3 key name.
RSLT	Result	The result of the ILM operation. SUCS: The ILM operation was successful.
RULE	Rules Label	The human-readable label given to the ILM rule applied to this object.
SEGC	Container UUID	UUID of the container for the segmented object. This value is available only if the object is segmented.
SGCB	Container CBID	CBID of the container for the segmented object. This value is available only for segmented and multipart objects.
STAT	Status	The status of ILM operation. DONE: ILM operations against the object have completed. DFER: The object has been marked for future ILM re-evaluation. PRGD: The object has been deleted from the StorageGRID system. NLOC: The object data can no longer be found in the StorageGRID system. This status might indicate that all copies of object data are missing or damaged.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.

Code	Field	Description
VSID	Version ID	The version ID of a new object created in a versioned bucket. Operations on buckets and objects in unversioned buckets don't include this field.

The ORLM audit message can be issued more than once for a single object. For instance, it is issued whenever one of the following events occur:

- ILM rules for the object are satisfied forever.
- ILM rules for the object are satisfied for this epoch.
- ILM rules have deleted the object.
- The background verification process detects that a copy of replicated object data is corrupt. The StorageGRID system performs an ILM evaluation to replace the corrupt object.

Related information

- [Object ingest transactions](#)
- [Object delete transactions](#)

OVWR: Object Overwrite

This message is generated when an external (client-requested) operation causes one object to be overwritten by another object.

Code	Field	Description
CBID	Content Block Identifier (new)	The CBID for the new object.
CSIZ	Previous Object Size	The size, in bytes, of the object being overwritten.
OCBD	Content Block Identifier (previous)	The CBID for the previous object.
UUID	Universally Unique ID (new)	The identifier of the new object within the StorageGRID system.
OUID	Universally Unique ID (previous)	The identifier for the previous object within the StorageGRID system.
PATH	S3 Object Path	The S3 object path used for both the previous and new object
RSLT	Result Code	Result of the Object Overwrite transaction. Result is always: SUCS: Successful

Code	Field	Description
SGRP	Site (Group)	If present, the overwritten object was deleted at the site specified, which is not the site where the overwritten object was ingested.

S3SL: S3 Select request

This message logs a completion after an S3 Select request has been returned to the client. The S3SL message can include error message and error code details. The request might not have been successful.

Code	Field	Description
BYSC	Bytes Scanned	Number of bytes scanned (received) from Storage Nodes. BYSC and BYPR are likely to be different if the object is compressed. If the object is compressed BYSC would have the compressed byte count and BYPR would be the bytes after decompression.
BYPR	Bytes Processed	Number of bytes processed. Indicates how many bytes of "Bytes Scanned" were actually processed or acted upon by an S3 Select job.
BYRT	Bytes Returned	Number of bytes that an S3 Select job returned to the client.
REPR	Records Processed	Number of records or rows that an S3 Select job received from Storage Nodes.
RERT	Records Returned	Number of records or rows an S3 Select job returned to the client.
JOFI	Job Finished	Indicates if the S3 Select job finished processing or not. If this is false, then the job failed to finish and the error fields will likely have data in them. The client might have received partial results, or no results at all.
REID	Request ID	Identifier for the S3 Select request.
EXTM	Execution Time	The time, in seconds, it took for the S3 Select Job to complete.
ERMG	Error Message	Error message that the S3 Select job generated.
ERTY	Error Type	Error type that the S3 Select job generated.
ERST	Error Stacktrace	Error Stacktrace that the S3 Select job generated.
S3BK	S3 bucket	The S3 bucket name.

Code	Field	Description
S3AK	S3 Access Key ID (request sender)	The S3 access key ID for the user that sent the request.
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request.
S3KY	S3 Key	The S3 key name, not including the bucket name.

SADD: Security Audit Disable

This message indicates that the originating service (node ID) has turned off audit message logging; audit messages are no longer being collected or delivered.

Code	Field	Description
AETM	Enable Method	The method used to disable the audit.
AEUN	User Name	The user name that executed the command to disable audit logging.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.

The message implies that logging was previously enabled, but has now been disabled. This is typically used only during bulk ingest to improve system performance. Following the bulk activity, auditing is restored (SADE) and the capability to disable auditing is then permanently blocked.

SADE: Security Audit Enable

This message indicates that the originating service (node ID) has restored audit message logging; audit messages are again being collected and delivered.

Code	Field	Description
AETM	Enable Method	The method used to enable the audit.
AEUN	User Name	The user name that executed the command to enable audit logging.
RSLT	Result	This field has the value NONE. RSLT is a mandatory message field, but is not relevant for this message. NONE is used rather than SUCS so that this message is not filtered.

The message implies that logging was previously disabled (SADD), but has now been restored. This is typically only used during bulk ingest to improve system performance. Following the bulk activity, auditing is restored and the capability to disable auditing is then permanently blocked.

SCMT: Object Store Commit

Grid content is not made available or recognized as stored until it has been committed (meaning it has been stored persistently). Persistently stored content has been completely written to disk, and has passed related integrity checks. This message is issued when a content block is committed to storage.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block committed to permanent storage.
RSLT	Result Code	Status at the time the object was stored to disk: SUCS: Object successfully stored.

This message means a given content block has been completely stored and verified, and can now be requested. It can be used to track data flow within the system.

SDEL: S3 DELETE

When an S3 client issues a DELETE transaction, a request is made to remove the specified object or bucket, or to remove a bucket/object subresource. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don't include this field.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the deleted object in bytes. Operations on buckets don't include this field.
DMRK	Delete Marker Version ID	The version ID of the delete marker created when deleting an object from a versioned bucket. Operations on buckets don't include this field.
GFID	Grid Federation Connection ID	The connection ID of the grid federation connection associated with a cross-grid replication delete request. Only included in audit logs on the destination grid.

Code	Field	Description
GFSA	Grid Federation Source Account ID	The account ID of the tenant on the source grid for a cross-grid replication delete request. Only included in audit logs on the destination grid.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field). x-amz-bypass-governance-retention is automatically included if it is present in the request.
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	Result of the DELETE transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.

Code	Field	Description
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SGRP	Site (Group)	If present, the object was deleted at the site specified, which is not the site where the object was ingested.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code> Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUDM	Universally Unique Identifier for a Delete Marker	The identifier of a delete marker. Audit log messages specify either UUDM or UUID, where UUDM indicates a delete marker created as a result of an object delete request, and UUID indicates an object.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was deleted. Operations on buckets and objects in unversioned buckets don't include this field.

SGET: S3 GET

When an S3 client issues a GET transaction, a request is made to retrieve an object or list the objects in a bucket, or to remove a bucket/object subresource. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don't include this field.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.

Code	Field	Description
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field).
LITY	ListObjectsV2	A <i>v2 format</i> response was requested. For details, see AWS ListObjectsV2 . For GET bucket operations only.
NCHD	Number of Children	Includes keys and common prefixes. For GET bucket operations only.
RANG	Range Read	For range read operations only. Indicates the range of bytes that was read by this request. The value after the slash (/) shows the size of the entire object.
RLT	Result Code	Result of the GET transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.

Code	Field	Description
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code> Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
TRNC	Truncated or Not Truncated	Set to false if all results were returned. Set to true if more results are available to return. For GET bucket operations only.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets don't include this field.

SHEA: S3 HEAD

When an S3 client issues a HEAD transaction, a request is made to check for the existence of an object or bucket and retrieve the metadata about an object. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don't include this field.

Code	Field	Description
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the checked object in bytes. Operations on buckets don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field).
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.

Code	Field	Description
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code> Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets don't include this field.

SPOS: S3 POST

When an S3 client issues a POST Object request, this message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field). (Not expected for SPOS).

Code	Field	Description
RSLT	Result Code	Result of the RestoreObject request. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable. Set to "select" for an S3 Select operation.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SRCF	Subresource Configuration	Restore information.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code> Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.

Code	Field	Description
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object that was requested. Operations on buckets and objects in unversioned buckets don't include this field.

SPUT: S3 PUT

When an S3 client issues a PUT transaction, a request is made to create a new object or bucket, or to remove a bucket/object subresource. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don't include this field.
CMPS	Compliance Settings	The compliance settings used when creating the bucket, if present in the request (truncated to the first 1024 characters).
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets don't include this field.
GFID	Grid Federation Connection ID	The connection ID of the grid federation connection associated with a cross-grid replication PUT request. Only included in audit logs on the destination grid.
GFSA	Grid Federation Source Account ID	The account ID of the tenant on the source grid for a cross-grid replication PUT request. Only included in audit logs on the destination grid.

Code	Field	Description
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field). x-amz-bypass-governance-retention is automatically included if it is present in the request.
LKEN	Object Lock Enabled	Value of the request header x-amz-bucket-object-lock-enabled, if present in the request.
LKLH	Object Lock Legal Hold	Value of the request header x-amz-object-lock-legal-hold, if present in the PutObject request.
LKMD	Object Lock Retention Mode	Value of the request header x-amz-object-lock-mode, if present in the PutObject request.
LKRU	Object Lock Retain Until Date	Value of the request header x-amz-object-lock-retain-until-date, if present in the PutObject request. Values are limited to within 100 years of the date the object was ingested.
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	Result of the PUT transaction. Result is always: SUCS: Successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.
S3SR	S3 Subresource	The bucket or object subresource being operated on, if applicable.

Code	Field	Description
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SRCF	Subresource Configuration	The new subresource configuration (truncated to the first 1024 characters).
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code> Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
ULID	Upload ID	Included only in SPUT messages for CompleteMultipartUpload operations. Indicates that all parts have been uploaded and assembled.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of a new object created in a versioned bucket. Operations on buckets and objects in unversioned buckets don't include this field.
VSST	Versioning State	The new versioning state of a bucket. Two states are used: "enabled" or "suspended." Operations on objects don't include this field.

SREM: Object Store Remove

This message is issued when content is removed from persistent storage and is no longer accessible through regular APIs.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block deleted from permanent storage.
RSLT	Result Code	Indicates the result of the content removal operations. The only defined value is: SUCS: Content removed from persistent storage

This audit message means a given content block has been deleted from a node and can no longer be requested directly. The message can be used to track the flow of deleted content within the system.

SUPD: S3 Metadata Updated

This message is generated by the S3 API when an S3 client updates the metadata for an ingested object. The message is issued by the server if the metadata update is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on buckets don't include this field.
CNCH	Consistency Control Header	The value of the Consistency-Control HTTP request header, if present in the request, when updating a bucket's compliance settings.
CNID	Connection Identifier	The unique system identifier for the TCP/IP connection.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on buckets don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field).
RSLT	Result Code	Result of the GET transaction. Result is always: SUCS: successful
S3AI	S3 tenant account ID (request sender)	The tenant account ID of the user who sent the request. An empty value indicates anonymous access.

Code	Field	Description
S3AK	S3 Access Key ID (request sender)	The hashed S3 access key ID for the user that sent the request. An empty value indicates anonymous access.
S3BK	S3 Bucket	The S3 bucket name.
S3KY	S3 Key	The S3 key name, not including the bucket name. Operations on buckets don't include this field.
SACC	S3 tenant account name (request sender)	The name of the tenant account for the user who sent the request. Empty for anonymous requests.
SAIP	IP address (request sender)	The IP address of the client application that made the request.
SBAC	S3 tenant account name (bucket owner)	The tenant account name for the bucket owner. Used to identify cross-account or anonymous access.
SBAI	S3 tenant account ID (bucket owner)	The tenant account ID of the owner of the target bucket. Used to identify cross-account or anonymous access.
SUSR	S3 User URN (request sender)	The tenant account ID and the user name of the user making the request. The user can either be a local user or an LDAP user. For example: <code>urn:sgws:identity::03393893651506583485:root</code> Empty for anonymous requests.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
VSID	Version ID	The version ID of the specific version of an object whose metadata was updated. Operations on buckets and objects in unversioned buckets don't include this field.

SVRF: Object Store Verify Fail

This message is issued whenever a content block fails the verification process. Each time replicated object data is read from or written to disk, several verification and integrity

checks are performed to ensure the data sent to the requesting user is identical to the data originally ingested into the system. If any of these checks fail, the system automatically quarantines the corrupt replicated object data to prevent it from being retrieved again.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block which failed verification.
RSLT	Result Code	Verification failure type: CRCF: Cyclic redundancy check (CRC) failed. HMAC: Hash-based message authentication code (HMAC) check failed. EHSR: Unexpected encrypted content hash. PHSR: Unexpected original content hash. SEQC: Incorrect data sequence on disk. PERR: Invalid structure of disk file. DERR: Disk error. FNAM: Bad file name.



This message should be monitored closely. Content verification failures can indicate impending hardware failures.

To determine what operation triggered the message, see the value of the AMID (Module ID) field. For example, an SVFY value indicates that the message was generated by the Storage Verifier module, that is, background verification, and STOR indicates that the message was triggered by content retrieval.

SVRU: Object Store Verify Unknown

The LDR service's Storage component continuously scans all copies of replicated object data in the object store. This message is issued when an unknown or unexpected copy of replicated object data is detected in the object store and moved to the quarantine directory.

Code	Field	Description
FPTH	File Path	The file path of the unexpected object copy.
RSLT	Result	This field has the value 'NONE'. RSLT is a mandatory message field, but is not relevant for this message. 'NONE' is used rather than 'SUCC' so that this message is not filtered.



The SVRU: Object Store Verify Unknown audit message should be monitored closely. It means unexpected copies of object data were detected in the object store. This situation should be investigated immediately to determine how these copies were created, because it can indicate impending hardware failures.

SYSD: Node Stop

When a service is stopped gracefully, this message is generated to indicate the shutdown was requested. Typically this message is sent only after a subsequent restart, because the audit message queue is not cleared before shutdown. Look for the SYST message, sent at the beginning of the shutdown sequence, if the service has not restarted.

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shutdown.

The message does not indicate if the host server is being stopped, only the reporting service. The RSLT of a SYSD can't indicate a "dirty" shutdown, because the message is generated only by "clean" shutdowns.

SYST: Node Stopping

When a service is gracefully stopped, this message is generated to indicate the shutdown was requested and that the service has initiated its shutdown sequence. SYST can be used to determine if the shutdown was requested, before the service is restarted (unlike SYSD, which is typically sent after the service restarts.)

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shutdown.

The message does not indicate if the host server is being stopped, only the reporting service. The RSLT code of a SYST message can't indicate a "dirty" shutdown, because the message is generated only by "clean" shutdowns.

SYSU: Node Start

When a service is restarted, this message is generated to indicate if the previous shutdown was clean (commanded) or disorderly (unexpected).

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS: System was cleanly shut down. DSDN: System was not cleanly shut down. VRGN: System was started for the first time after server installation (or re-installation).

The message does not indicate if the host server was started, only the reporting service. This message can be used to:

- Detect discontinuity in the audit trail.
- Determine if a service is failing during operation (as the distributed nature of the StorageGRID system can mask these failures). Server Manager restarts a failed service automatically.

WDEL: Swift DELETE

When a Swift client issues a DELETE transaction, a request is made to remove the specified object or container. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on containers don't include this field.
CSIZ	Content Size	The size of the deleted object in bytes. Operations on containers don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field).
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	Result of the DELETE transaction. Result is always: SUCS: Successful
SAIP	IP address of requesting client	The IP address of the client application that made the request.

Code	Field	Description
SGRP	Site (Group)	If present, the object was deleted at the site specified, which is not the site where the object was ingested.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers don't include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

WGET: Swift GET

When a Swift client issues a GET transaction, a request is made to retrieve an object, list the objects in a container, or list the containers in an account. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on accounts and containers don't include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field).

Code	Field	Description
RSLT	Result Code	Result of the GET transaction. Result is always SUCS: successful
SAIP	IP address of requesting client	The IP address of the client application that made the request.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name. Operations on accounts don't include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers don't include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

WHEA: Swift HEAD

When a Swift client issues a HEAD transaction, a request is made to check for the existence of an account, container, or object, and retrieve any relevant metadata. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on accounts and containers don't include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers don't include this field.

Code	Field	Description
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field).
RSLT	Result Code	Result of the HEAD transaction. Result is always: SUCS: successful
SAIP	IP address of requesting client	The IP address of the client application that made the request.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name. Operations on accounts don't include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers don't include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

WPUT: Swift PUT

When a Swift client issues a PUT transaction, a request is made to create a new object or container. This message is issued by the server if the transaction is successful.

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block requested. If the CBID is unknown, this field is set to 0. Operations on containers don't include this field.

Code	Field	Description
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on containers don't include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration. X-Forwarded-For is automatically included if it is present in the request and if the X-Forwarded-For value is different from the request sender IP address (SAIP audit field).
MTME	Last Modified Time	The Unix timestamp, in microseconds, indicating when the object was last modified.
RSLT	Result Code	Result of the PUT transaction. Result is always: SUCS: successful
SAIP	IP address of requesting client	The IP address of the client application that made the request.
TIME	Time	Total processing time for the request in microseconds.
TLIP	Trusted Load Balancer IP Address	If the request was routed by a trusted Layer 7 load balancer, the IP address of the load balancer.
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.
WACC	Swift Account ID	The unique account ID as specified by the StorageGRID system.
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers don't include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

Expand a grid

Expansion types

You can expand the capacity or capabilities of your StorageGRID system without interrupting system operations.

A StorageGRID expansion allows you to add:

- Storage volumes to Storage Nodes
- New grid nodes to an existing site
- An entire new site

The reason you are performing the expansion determines how many new nodes of each type you must add and the location of those new nodes. For example, there are different node requirements if you are performing an expansion to increase storage capacity, add metadata capacity, or add redundancy or new capabilities.

Follow the steps for the type of expansion you're performing:

Add storage volumes

Follow the steps for [adding storage volumes to Storage Nodes](#).

Add grid nodes

1. Follow the steps for [adding grid nodes to an existing site](#).
2. [Update the subnets](#).
3. Deploy grid nodes:
 - [Appliances](#)
 - [VMware](#)
 - [Linux](#)



"Linux" refers to a Red Hat Enterprise Linux, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

4. [Perform the expansion](#).
5. [Configure the expanded system](#).

Add new site

1. Follow the steps for [Adding a new site](#).
2. [Update the subnets](#).
3. Deploy grid nodes:
 - [Appliances](#)
 - [VMware](#)
 - [Linux](#)



"Linux" refers to a Red Hat Enterprise Linux, Ubuntu, or Debian deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

4. [Perform the expansion](#).
5. [Configure the expanded system](#).

Plan StorageGRID expansion

Add storage capacity

Guidelines for adding object capacity

You can expand the object storage capacity of your StorageGRID system by adding storage volumes to existing Storage Nodes or by adding new Storage Nodes to existing sites. You must add storage capacity in a way that meets the requirements of your information lifecycle management (ILM) policy.

Guidelines for adding storage volumes

Before adding storage volumes to existing Storage Nodes, review the following guidelines and limitations:

- You must examine your current ILM rules to determine where and when to [add storage volumes](#) to increase the storage available for [replicated objects](#) or [erasure-coded objects](#).
- You can't increase the metadata capacity of your system by adding storage volumes because object metadata is stored only on volume 0.
- Each software-based Storage Node can support a maximum of 48 storage volumes. If you need to add capacity beyond that, you must add new Storage Nodes.
- You can add one or two expansion shelves to each SG6060 appliance. Each expansion shelf adds 16 storage volumes. With both expansion shelves installed, the SG6060 can support a total of 48 storage volumes.
- You can add one or two expansion shelves to each SG6160 appliance. Each expansion shelf adds 60 storage volumes. With both expansion shelves installed, the SG6160 can support a total of 180 storage volumes.
- You can't add storage volumes to any other storage appliance.
- You can't increase the size of an existing storage volume.
- You can't add storage volumes to a Storage Node at the same time you are performing a system upgrade, recovery operation, or another expansion.

After you have decided to add storage volumes and have determined which Storage Nodes you must expand to satisfy your ILM policy, follow the instructions for your type of Storage Node:

- To add one or two expansion shelves to an SG6060 storage appliance, go to [Add expansion shelf to deployed SG6060](#).
- To add one or two expansion shelves to an SG6160 storage appliance, go to [Add expansion shelf to deployed SG6160](#)
- For a software-based node, follow the instructions for [adding storage volumes to Storage Nodes](#).

Guidelines for adding Storage Nodes

Before adding Storage Nodes to existing sites, review the following guidelines and limitations:

- You must examine your current ILM rules to determine where and when to add Storage Nodes to increase the storage available for [replicated objects](#) or [erasure-coded objects](#).
- You should not add more than 10 Storage Nodes in a single expansion procedure.
- You can add Storage Nodes to more than one site in a single expansion procedure.
- You can add Storage Nodes and other types of nodes in a single expansion procedure.
- Before starting the expansion procedure, you must confirm that all data-repair operations performed as part of a recovery are complete. See [Check data repair jobs](#).
- If you need to remove Storage Nodes before or after performing an expansion, you should not decommission more than 10 Storage Nodes in a single Decommission Node procedure.

Guidelines for ADC service on Storage Nodes

When configuring the expansion, you must choose whether to include the Administrative Domain Controller (ADC) service on each new Storage Node. The ADC service keeps track of the location and availability of grid

services.

- The StorageGRID system requires a [quorum of ADC services](#) to be available at each site and at all times.
- At least three Storage Nodes at each site must include the ADC service.
- Adding the ADC service to every Storage Node is not recommended. Including too many ADC services can cause slowdowns due to the increased amount of communication between nodes.
- A single grid should not have more than 48 Storage Nodes with the ADC service. This is equivalent to 16 sites with three ADC services at each site.
- In general, when you select the **ADC Service** setting for a new node, you should select **Automatic**. Select **Yes** only if the new node will replace another Storage Node that includes the ADC service. Because you can't decommission a Storage Node if too few ADC services would remain, this ensures that a new ADC service is available before the old service is removed.
- You can't add the ADC service to a node after it is deployed.

Add storage capacity for replicated objects

If the information lifecycle management (ILM) policy for your deployment includes a rule that creates replicated copies of objects, you must consider how much storage to add and where to add the new storage volumes or Storage Nodes.

For guidance on where to add additional storage, examine the ILM rules that create replicated copies. If ILM rules create two or more object copies, plan to add storage in each location where object copies are made. As a simple example, if you have a two-site grid and an ILM rule that creates one object copy at each site, you must [add storage](#) to each site to increase the overall object capacity of the grid. For information about object replication, see [What is replication](#).

For performance reasons, you should attempt to keep storage capacity and compute power balanced across sites. So, for this example, you should add the same number of Storage Nodes to each site or additional storage volumes at each site.

If you have a more complex ILM policy that includes rules that place objects in different locations based on criteria such as bucket name, or rules that change object locations over time, your analysis of where storage is required for the expansion will be similar, but more complex.

Charting how quickly overall storage capacity is being consumed can help you understand how much storage to add in the expansion, and when the additional storage space will be required. You can use the Grid Manager to [monitor and chart storage capacity](#).

When planning the timing of an expansion, remember to consider how long it might take to procure and install additional storage.

Add storage capacity for erasure-coded objects

If your ILM policy includes a rule that makes erasure-coded copies, you must plan where to add new storage and when to add new storage. The amount of storage you add and the timing of the addition can affect the grid's usable storage capacity.

The first step in planning a storage expansion is to examine the rules in your ILM policy that create erasure-coded objects. Because StorageGRID creates $k+m$ fragments for every erasure-coded object and stores each fragment on a different Storage Node, you must ensure that at least $k+m$ Storage Nodes have space for new erasure-coded data after the expansion. If the erasure-coding profile provides site-loss protection, you must

add storage to each site. See [What are erasure-coding schemes](#) for information about erasure-coding profiles.

The number of nodes you need to add also depends on how full the existing nodes are when you perform the expansion.

General recommendation for adding storage capacity for erasure-coded objects

If you want to avoid detailed calculations, you can add two Storage Nodes per site when existing Storage Nodes reach 70% capacity.

This general recommendation provides reasonable results across a wide range of erasure-coding schemes for both single-site grids and for grids where erasure coding provides site-loss protection.

To better understand the factors that led to this recommendation or to develop a more precise plan for your site, see [Considerations for rebalancing erasure-coded data](#). For a custom recommendation optimized for your situation, contact your NetApp Professional Services consultant.

Considerations for rebalancing erasure-coded data

If you are performing an expansion to add Storage Nodes and you use ILM rules to erasure code data, you might need to perform the erasure coding (EC) rebalance procedure if you can't add enough Storage Nodes for the erasure-coding scheme you are using.

After reviewing these considerations, perform the expansion, and then go to [Rebalance erasure-coded data after adding Storage Nodes](#) to run the procedure.

What is EC rebalancing?

EC rebalancing is a StorageGRID procedure that might be required after a Storage Node expansion. The procedure is run as a command-line script from the primary Admin Node. When you run the EC rebalance procedure, StorageGRID redistributes erasure-coded fragments among the existing and the newly added Storage Nodes at a site.

The EC rebalance procedure:


- Only moves erasure-coded object data. It does not move replicated object data.
- Redistributes the data within a site. It does not move data between sites.
- Redistributes data among all Storage Nodes at a site. It does not redistribute data within storage volumes.
- Does not consider the replicated data usage on each Storage Node when determining where to move erasure-coded data.
- Redistributes erasure-coded data evenly between Storage Nodes without considering the relative capacities of each node.
- Will not distribute erasure-coded data to Storage Nodes that are more than 80% full.
- Might decrease the performance of ILM operations and S3 client operations when it runs—additional resources are required to redistribute the erasure-coding fragments.

When the EC rebalance procedure is complete:

- Erasure-coded data will have moved from Storage Nodes with less available space to Storage Nodes with more available space.

- The data protection of erasure-coded objects will be unchanged.
- Used (%) values might be different between Storage Nodes for two reasons:
 - Replicated object copies will continue to consume space on the existing nodes—the EC rebalance procedure does not move replicated data.
 - Larger-capacity nodes will be relatively less full than smaller-capacity nodes, even though all nodes will end up with approximately the same amount of erasure-coded data.


For example, suppose three 200-TB nodes are each filled to 80% ($200 \times 0.8 = 160$ TB on each node, or 480 TB for the site). If you add a 400-TB node and run the rebalance procedure, all nodes will now have approximately the same amount of erasure-code data (480/4 = 120 TB). However, the Used (%) for the larger node will be less than the Used (%) for the smaller nodes.



When to rebalance erasure-coded data


Consider the following scenario:

- StorageGRID is running at a single site, which contains three Storage Nodes.
- The ILM policy uses a 2+1 erasure-coding rule for all objects larger than 1.0 MB and a 2-copy replication rule for smaller objects.
- All Storage Nodes have become completely full. The **Low Object Storage** alert has been triggered at the major severity level.



Rebalance is not required if you add enough nodes


To understand when EC rebalance is not required, suppose you added three (or more) new Storage Nodes. In this case, you don't need to perform EC rebalance. The original Storage Nodes will remain full, but new objects will now use the three new nodes for 2+1 erasure coding—the two data fragments and the one parity fragment can each be stored on a different node.




While you can run the EC rebalance procedure in this case, moving the existing erasure-coded data will temporarily decrease the grid's performance, which might impact client operations.

Rebalance is required if you can't add enough nodes

To understand when EC rebalance is required, suppose you can only add two Storage Nodes, instead of three. Because the 2+1 scheme requires at least three Storage Nodes to have space available, the empty nodes can't be used for new erasure-coded data.



To make use of the new Storage Nodes, you should run the EC rebalance procedure. When this procedure runs, StorageGRID redistributes existing erasure-coded data and parity fragments among all Storage Nodes at the site. In this example, when the EC rebalance procedure is complete, all five nodes are now only 60% full, and objects can continue to be ingested into the 2+1 erasure-coding scheme on all Storage Nodes.



Recommendations for EC rebalancing

NetApp requires EC rebalancing if *all* of the following statements are true:

- You use erasure coding for your object data.
- The **Low Object Storage** alert has been triggered for one or more Storage Nodes at a site, indicating that the nodes are 80% or more full.
- You are unable to add enough new Storage Nodes for the erasure-coding scheme in use. See [Add storage capacity for erasure-coded objects](#).
- Your S3 clients can tolerate lower performance for their write and read operations while the EC rebalance procedure is running.

You can optionally run the EC rebalance procedure if you prefer Storage Nodes to be filled to similar levels and your S3 clients can tolerate lower performance for their write and read operations while the EC rebalance procedure is running.

How EC rebalance procedure interacts with other maintenance tasks

You can't perform certain maintenance procedures at the same time you are running the EC rebalance procedure.

Procedure	Allowed during EC rebalance procedure?
Additional EC rebalance procedures	No. You can only run one EC rebalance procedure at a time.
Decommission procedure EC data repair job	No. • You are prevented from starting a decommission procedure or an EC data repair while the EC rebalance procedure is running. • You are prevented from starting the EC rebalance procedure while a Storage Node decommission procedure or an EC data repair is running.
Expansion procedure	No. If you need to add new Storage Nodes in an expansion, run the EC rebalance procedure after adding all new nodes.
Upgrade procedure	No. If you need to upgrade StorageGRID software, perform the upgrade procedure before or after running the EC rebalance procedure. As required, you can terminate the EC rebalance procedure to perform a software upgrade.
Appliance node clone procedure	No. If you need to clone an appliance Storage Node, run the EC rebalance procedure after adding the new node.

Procedure	Allowed during EC rebalance procedure?
Hotfix procedure	Yes. You can apply a StorageGRID hotfix while the EC rebalance procedure is running.
Other maintenance procedures	No. You must terminate the EC rebalance procedure before running other maintenance procedures.

How EC rebalance procedure interacts with ILM

While the EC rebalance procedure is running, avoid making ILM changes that might change the location of existing erasure-coded objects. For example, don't start using an ILM rule that has a different erasure-coding profile. If you need to make such ILM changes, you should terminate the EC rebalance procedure.

Add metadata capacity

To ensure that adequate space is available for object metadata, you might need to perform an expansion procedure to add new Storage Nodes at each site.

StorageGRID reserves space for object metadata on volume 0 of each Storage Node. Three copies of all object metadata are maintained at each site, evenly distributed across all Storage Nodes.

You can use the Grid Manager to monitor the metadata capacity of Storage Nodes and to estimate how quickly metadata capacity is being consumed. In addition, the **Low metadata storage** alert is triggered for a Storage Node when the used metadata space reaches certain thresholds.

Note that a grid's object metadata capacity might be consumed faster than its object storage capacity, depending on how you use the grid. For example, if you typically ingest large numbers of small objects or add large quantities of user metadata or tags to objects, you might need to add Storage Nodes to increase metadata capacity even though sufficient object storage capacity remains.

For more information, see the following:

- [Manage object metadata storage](#)
- [Monitor object metadata capacity for each Storage Node](#)

Guidelines for increasing metadata capacity

Before adding Storage Nodes to increase metadata capacity, review the following guidelines and limitations:

- Assuming sufficient object storage capacity is available, having more space available for object metadata increases the number of objects you can store in your StorageGRID system.
- You can increase a grid's metadata capacity by adding one or more Storage Nodes to each site.
- The actual space reserved for object metadata on any given Storage Node depends on the Metadata Reserved Space storage option (system-wide setting), the amount of RAM allocated to the node, and the size of the node's volume 0.
- You can't increase metadata capacity by adding storage volumes to existing Storage Nodes, because

metadata is stored only on volume 0.

- You can't increase metadata capacity by adding a new site.
- StorageGRID keeps three copies of all object metadata at every site. For this reason, the metadata capacity for your system is limited by the metadata capacity of your smallest site.
- When adding metadata capacity, you should add the same number of Storage Nodes to each site.

Metadata-only Storage Nodes have specific hardware requirements:

- When using StorageGRID appliances, metadata-only nodes can be configured only on SGF6112 appliances with twelve 1.9 TB or twelve 3.8 TB drives.
- When using software-based nodes, metadata-only node resources must match the existing Storage Nodes resources. For example:
 - If the existing StorageGRID site is using SG6000 or SG6100 appliances, the software-based metadata-only nodes must meet the following minimum requirements:
 - 128 GB RAM
 - 8 core CPU
 - 8 TB SSD or equivalent storage for the Cassandra database (rangedb/0)
 - If the existing StorageGRID site is using virtual Storage Nodes with 24 GB RAM, 8 core CPU, and 3 TB or 4TB of metadata storage, the software-based metadata-only nodes should use similar resources (24 GB RAM, 8 core CPU, and 4TB of metadata storage (rangedb/0).
- When adding a new StorageGRID site, the new site total metadata capacity should, at minimum, match existing StorageGRID sites and new site resources should match the Storage Nodes at existing StorageGRID sites.


See the [description of what Metadata Reserved Space is](#).

How metadata is redistributed when you add Storage Nodes

When you add Storage Nodes in an expansion, StorageGRID redistributes the existing object metadata to the new nodes at each site, which increases the overall metadata capacity of the grid. No user action is required.

The following figure shows how StorageGRID redistributes object metadata when you add Storage Nodes in an expansion. The left side of the figure represents volume 0 of three Storage Nodes before an expansion. Metadata is consuming a relatively large portion of each node's available metadata space, and the **Low metadata storage** alert has been triggered.

The right side of the figure shows how the existing metadata is redistributed after two Storage Nodes are added to the site. The amount of metadata on each node has decreased, the **Low metadata storage** alert is no longer triggered, and the space available for metadata has increased.



Add grid nodes to add capabilities to your system

You can add redundancy or additional capabilities to a StorageGRID system by adding new grid nodes to existing sites.

For example, you might choose to add Gateway Nodes to use in a high availability (HA) group, or you might add an Admin Node at a remote site to permit monitoring using a local node.

You can add one or more of the following types of nodes to one or more existing sites in a single expansion operation:

- Non-primary Admin Nodes
- Storage Nodes
- Gateway Nodes

When preparing to add grid nodes, be aware of the following limitations:

- The primary Admin Node is deployed during the initial installation. You can't add a primary Admin Node during an expansion.
- You can add Storage Nodes and other types of nodes in the same expansion.
- When adding Storage Nodes, you must carefully plan the number and location of the new nodes. See [Guidelines for adding object capacity](#).
- If the **Set new node default** option is **Untrusted** on the Untrusted Client Networks tab on the Firewall control page, client applications that connect to expansion nodes using the Client Network must connect using a load balancer endpoint port (**CONFIGURATION > Security > Firewall control**). See the instructions to [change the security setting for the new node](#) and to [configure load balancer endpoints](#).

Add a new site

You can expand your StorageGRID system by adding a new site.

Guidelines for adding a site

Before adding a site, review the following requirements and limitations:

- You can only add one site per expansion operation.
- You can't add grid nodes to an existing site as part of the same expansion.
- All sites must include at least three Storage Nodes.
- Adding a new site does not automatically increase the number of objects you can store. The total object capacity of a grid depends on the amount of available storage, the ILM policy, and the metadata capacity at each site.
- When sizing a new site, you must ensure that it includes enough metadata capacity.

StorageGRID keeps a copy of all object metadata at every site. When you add a new site, you must ensure that it includes enough metadata capacity for the existing object metadata and enough metadata capacity for growth.

For more information, see the following:

- [Manage object metadata storage](#)
- [Monitor object metadata capacity for each Storage Node](#)
- You must consider the available network bandwidth between sites, and the level of network latency. Metadata updates are continually replicated between sites even if all objects are stored only at the site where they are ingested.
- Because your StorageGRID system remains operational during the expansion, you must review ILM rules before starting the expansion procedure. You must ensure that object copies aren't stored to the new site until the expansion procedure is complete.

For example, before you begin the expansion, determine if any rules use the default storage pool (All Storage Nodes). If they do, you must create a new storage pool that contains the existing Storage Nodes and update your ILM rules to use the new storage pool. Otherwise, objects will be copied to the new site as soon as the first node at that site becomes active.

For more information about changing ILM when adding a new site, see the [example for changing an ILM policy](#).

Gather required materials

Before performing an expansion operation, gather the materials and install and configure any new hardware and networks.

Item	Notes
StorageGRID installation archive	If you are adding new grid nodes or a new site, you must download and extract the StorageGRID installation archive. You must use the same version that is currently running on the grid. For details, see the instructions for downloading and extracting the StorageGRID installation files. Note: You don't need to download files if you are adding new storage volumes to existing Storage Nodes or installing a new StorageGRID appliance.
Service laptop	The service laptop has the following: • Network port • SSH client (for example, PuTTY) • Supported web browser
Passwords.txt file	Contains the passwords required to access grid nodes on the command line. Included in the Recovery Package.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the Passwords.txt file.
StorageGRID documentation	• Administer StorageGRID • Release notes • Installation instructions for your platform ◦ Install StorageGRID on Red Hat Enterprise Linux ◦ Install StorageGRID on Ubuntu or Debian ◦ Install StorageGRID on VMware
Current documentation for your platform	For supported versions, see the Interoperability Matrix Tool (IMT) .

Download and extract the StorageGRID installation files

Before you can add new grid nodes or a new site, you must download the appropriate StorageGRID installation archive and extract the files.

About this task

You must perform expansion operations using the version of StorageGRID that is currently running on the grid.

Steps

1. Go to [NetApp Downloads: StorageGRID](#).

2. Select the version of StorageGRID that is currently running on the grid.
3. Sign in with the username and password for your NetApp account.
4. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.
5. In the **Install StorageGRID** column of the download page, select the `.tgz` or `.zip` file for your platform.

The version shown in the installation archive file must match the version of the software that is currently installed.

Use the `.zip` file if you are running Windows on the service laptop.

Platform	Installation archive
Red Hat Enterprise Linux	<code>StorageGRID-Webscale-version-RPM-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-RPM-uniqueID.tgz</code>
Ubuntu or Debian or Appliances	<code>StorageGRID-Webscale-version-DEB-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-DEB-uniqueID.tgz</code>
VMware	<code>StorageGRID-Webscale-version-VMware-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-VMware-uniqueID.tgz</code>
OpenStack/other Hypervisor	To expand an existing deployment on OpenStack, you must deploy a virtual machine running one of the supported Linux distributions listed above and follow the appropriate instructions for Linux.

6. Download and extract the archive file.
7. Follow the appropriate step for your platform to choose the files you need, based on your platform, planned grid topology, and how you will expand your StorageGRID system.

The paths listed in the step for each platform are relative to the top-level directory installed by the archive file.

8. If you are expanding a Red Hat Enterprise Linux system, select the appropriate files.

Path and file name	Description
<code>./rpms/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./rpms/NLF000000.txt</code>	A free license that does not provide any support entitlement for the product.
<code>./rpms/StorageGRID-Webscale-Images-version-SHA.rpm</code>	RPM package for installing the StorageGRID node images on your RHEL hosts.
<code>./rpms/StorageGRID-Webscale-Service-version-SHA.rpm</code>	RPM package for installing the StorageGRID host service on your RHEL hosts.

Path and file name	Description
Deployment scripting tool	Description
<code>./rpms/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.
<code>./rpms/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./rpms/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./rpms/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate integration.
<code>./rpms/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./rpms/extras/ansible</code>	Example Ansible role and playbook for configuring RHEL hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
<code>./rpms/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
<code>./rpms/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.
<code>./rpms/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

9. If you are expanding an Ubuntu or Debian system, select the appropriate files.

Path and file name	Description
<code>./debs/README</code>	A text file that describes all of the files contained in the StorageGRID download file.

Path and file name	Description
<code>./debs/NLF000000.txt</code>	A non-production NetApp License File that you can use for testing and proof of concept deployments.
<code>./debs/storagegrid-webscale-images-version-SHA.deb</code>	DEB package for installing the StorageGRID node images on Ubuntu or Debian hosts.
<code>./debs/storagegrid-webscale-images-version-SHA.deb.md5</code>	MD5 checksum for the file <code>./debs/storagegrid-webscale-images-version-SHA.deb</code> .
<code>./debs/storagegrid-webscale-service-version-SHA.deb</code>	DEB package for installing the StorageGRID host service on Ubuntu or Debian hosts.
Deployment scripting tool	Description
<code>./debs/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.
<code>./debs/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./debs/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate integration.
<code>./debs/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./debs/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./debs/extras/ansible</code>	Example Ansible role and playbook for configuring Ubuntu or Debian hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
<code>./debs/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
<code>./debs/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.

Path and file name	Description
<code>./debs/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

10. If you are expanding a VMware system, select the appropriate files.

Path and file name	Description
<code>./vsphere/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./vsphere/NLF000000.txt</code>	A free license that does not provide any support entitlement for the product.
<code>./vsphere/NetApp-SG-version-SHA.vmdk</code>	The virtual machine disk file that is used as a template for creating grid node virtual machines.
<code>./vsphere/vsphere-primary-admin.ovf</code> <code>./vsphere/vsphere-primary-admin.mf</code>	The Open Virtualization Format template file (.ovf) and manifest file (.mf) for deploying the primary Admin Node.
<code>./vsphere/vsphere-non-primary-admin.ovf</code> <code>./vsphere/vsphere-non-primary-admin.mf</code>	The template file (.ovf) and manifest file (.mf) for deploying non-primary Admin Nodes.
<code>./vsphere/vsphere-gateway.ovf</code> <code>./vsphere/vsphere-gateway.mf</code>	The template file (.ovf) and manifest file (.mf) for deploying Gateway Nodes.
<code>./vsphere/vsphere-storage.ovf</code> <code>./vsphere/vsphere-storage.mf</code>	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
Deployment scripting tool	Description
<code>./vsphere/deploy-vsphere-ovftool.sh</code>	A Bash shell script used to automate the deployment of virtual grid nodes.
<code>./vsphere/deploy-vsphere-ovftool-sample.ini</code>	An example configuration file for use with the <code>deploy-vsphere-ovftool.sh</code> script.
<code>./vsphere/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.

Path and file name	Description
<code>./vsphere/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./vsphere/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled. You can also use this script for Ping Federate integration.
<code>./vsphere/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./vsphere/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./vsphere/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
<code>./vsphere/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.
<code>./vsphere/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

11. If you are expanding a StorageGRID appliance-based system, select the appropriate files.

Path and file name	Description
<code>./debs/storagegrid-webscale-images-version-SHA.deb</code>	DEB package for installing the StorageGRID node images on your appliances.
<code>./debs/storagegrid-webscale-images-version-SHA.deb.md5</code>	MD5 checksum for the file <code>/debs/storagegridwebscale-images-version-SHA.deb</code> .



For appliance installation, these files are only required if you need to avoid network traffic. The appliance can download the required files from the primary Admin Node.

Verify hardware and networking

Before beginning the expansion of your StorageGRID system, ensure the following:

- The hardware needed to support the new grid nodes or new site has been installed and configured.
- All new nodes have bidirectional communication paths to all existing and new nodes (a requirement for the Grid Network). In particular, confirm that the following TCP ports are open between the new nodes you are adding in the expansion and the primary Admin Node:
 - 1055
 - 7443
 - 8011
 - 10342

See [Internal grid node communications](#).

- The primary Admin Node can communicate with all expansion servers that are intended to host the StorageGRID system.
- If any of the new nodes has a Grid Network IP address on a subnet not previously used, you have already [added the new subnet](#) to the Grid Network subnet list. Otherwise, you will have to cancel the expansion, add the new subnet, and start the procedure again.
- You aren't using network address translation (NAT) on the Grid Network between grid nodes or between StorageGRID sites. When you use private IPv4 addresses for the Grid Network, those addresses must be directly routable from every grid node at every site. Using NAT to bridge the Grid Network across a public network segment is supported only if you use a tunneling application that is transparent to all nodes in the grid, meaning the grid nodes require no knowledge of public IP addresses.

This NAT restriction is specific to grid nodes and the Grid Network. As required, you can use NAT between external clients and grid nodes, such as to provide a public IP address for a Gateway Node.

Add storage volumes

Add storage volumes to Storage Nodes

You can expand the storage capacity of Storage Nodes that are below the maximum supported number of volumes. You might need to add storage volumes to more than one Storage Node to satisfy ILM requirements for replicated or erasure-coded copies.

Before you begin

Before adding storage volumes, review the [guidelines for adding object capacity](#) to ensure that you know where to add volumes to meet the requirements of your ILM policy.



These instructions apply to software-based Storage Nodes only. See [Add expansion shelf to deployed SG6060](#) or [Add expansion shelf to deployed SG6160](#) to learn how to add storage volumes to the SG6060 or SG6160 by installing expansion shelves. Other appliance Storage Nodes can't be expanded.

About this task

The underlying storage of a Storage Node is divided into storage volumes. Storage volumes are block-based

storage devices that are formatted by the StorageGRID system and mounted to store objects. Each Storage Node can support up to 48 storage volumes, which are called *object stores* in the Grid Manager.



Object metadata is always stored in object store 0.

Each object store is mounted on a volume that corresponds to its ID. For example, the object store with an ID of 0000 corresponds to the `/var/local/rangedb/0` mount point.

Before adding new storage volumes, use the Grid Manager to view the current object stores for each Storage Node as well as the corresponding mount points. You can use this information when adding storage volumes.

Steps

1. Select **NODES** > **site** > **Storage Node** > **Storage**.
2. Scroll down to view the amounts of available storage for each volume and object store.

For appliance Storage Nodes, the Worldwide Name for each disk matches the volume world-wide identifier (WWID) that appears when you view standard volume properties in SANtricity OS (the management software connected to the appliance's storage controller).

To help you interpret disk read and write statistics related to volume mount points, the first portion of the name shown in the **Name** column of the Disk Devices table (that is, *sd*, *sd*, *sd*, and so on) matches the value shown in the **Device** column of the Volumes table.

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

3. Follow the instructions for your platform to add new storage volumes to the Storage Node.

- [VMware: Add storage volumes to Storage Node](#)
- [Linux: Add direct-attached or SAN volumes to Storage Node](#)

VMware: Add storage volumes to Storage Node

If a Storage Node includes fewer than 16 storage volumes, you can increase its capacity by using VMware vSphere to add volumes.

Before you begin

- You have access to the instructions for installing StorageGRID for VMware deployments.
 - [Install StorageGRID on VMware](#)
- You have the `Passwords.txt` file.
- You have [specific access permissions](#).



Don't attempt to add storage volumes to a Storage Node while a software upgrade, recovery procedure, or another expansion procedure is active.

About this task

The Storage Node is unavailable for a brief time when you add storage volumes. You should perform this procedure on one Storage Node at a time to avoid impacting client-facing grid services.

Steps

1. If necessary, install new storage hardware and create new VMware datastores.
2. Add one or more hard disks to the virtual machine for use as storage (object stores).
 - a. Open VMware vSphere Client.
 - b. Edit the virtual machine settings to add one or more additional hard disks.

The hard disks are typically configured as Virtual Machine Disks (VMDKs). VMDKs are more commonly used and are easier to manage, while RDMS might provide better performance for workloads that use larger object sizes (for example, greater than 100 MB). For more information about adding hard disks to virtual machines, see the VMware vSphere documentation.

3. Restart the virtual machine by using the **Restart Guest OS** option in the VMware vSphere Client, or by entering the following command in an ssh session to the virtual machine: `sudo reboot`



Don't use **Power Off** or **Reset** to restart the virtual machine.

4. Configure the new storage for use by the Storage Node:

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

- b. Configure the new storage volumes:

```
sudo add_rangedbs.rb
```

This script finds any new storage volumes and prompts you to format them.

- c. Enter **y** to accept the formatting.
- d. If any of the volumes have previously been formatted, decide if you want to reformat them.
 - Enter **y** to reformat.
 - Enter **n** to skip reformatting.

The `setup_rangedbs.sh` script runs automatically.

5. Check that the services start correctly:

- a. View a list of the status of all services on the server:

```
sudo storagegrid-status
```

The status is updated automatically.

- b. Wait until all services are Running or Verified.
- c. Exit the status screen:

```
Ctrl+C
```

6. Verify that the Storage Node is online:

- a. Sign in to the Grid Manager using a [supported web browser](#).
- b. Select **SUPPORT > Tools > Grid topology**.
- c. Select **site > Storage Node > LDR > Storage**.
- d. Select the **Configuration** tab and then the **Main** tab.
- e. If the **Storage State - Desired** drop-down list is set to Read-only or Offline, select **Online**.
- f. Select **Apply Changes**.

7. To see the new object stores:

- a. Select **NODES > site > Storage Node > Storage**.
- b. View the details in the **Object Stores** table.

Result

You can use the expanded capacity of the Storage Nodes to save object data.

Linux: Add direct-attached or SAN volumes to Storage Node

If a Storage Node includes fewer than 48 storage volumes, you can increase its capacity by adding new block storage devices, making them visible to the Linux hosts, and adding the new block device mappings to the StorageGRID configuration file used for the Storage Node.

Before you begin

- You have access to the instructions for installing StorageGRID for your Linux platform.
 - [Install StorageGRID on Red Hat Enterprise Linux](#)
 - [Install StorageGRID on Ubuntu or Debian](#)
- You have the `Passwords.txt` file.
- You have [specific access permissions](#).



Don't attempt to add storage volumes to a Storage Node while a software upgrade, recovery procedure, or another expansion procedure is active.

About this task

The Storage Node is unavailable for a brief time when you add storage volumes. You should perform this procedure on one Storage Node at a time to avoid impacting client-facing grid services.

Steps

1. Install the new storage hardware.

For more information, see the documentation provided by your hardware vendor.

2. Create new block storage volumes of the desired sizes.

- Attach the new drives and update the RAID controller configuration as needed, or allocate the new SAN LUNs on the shared storage arrays and allow the Linux host to access them.
- Use the same persistent naming scheme you used for the storage volumes on the existing Storage Node.
- If you use the StorageGRID node migration feature, make the new volumes visible to other Linux hosts that are migration targets for this Storage Node.

For more information, see the instructions for installing StorageGRID for your Linux platform.

3. Log in to the Linux host supporting the Storage Node as root or with an account that has sudo permission.
4. Confirm that the new storage volumes are visible on the Linux host.

You might have to rescan for devices.

5. Run the following command to temporarily disable the Storage Node:

```
sudo storagegrid node stop <node-name>
```

6. Using a text editor such as vim or pico, edit the node configuration file for the Storage Node, which can be found at `/etc/storagegrid/nodes/<node-name>.conf`.
7. Locate the section of the node configuration file that contains the existing object storage block device mappings.

In the example, `BLOCK_DEVICE_RANGEDB_00` to `BLOCK_DEVICE_RANGEDB_03` are the existing object storage block device mappings.

```

NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/sgws-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/sgws-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/sgws-sn1-rangedb-3
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

```

8. Add new object storage block device mappings corresponding to the block storage volumes you added for this Storage Node.

Make sure to start at the next `BLOCK_DEVICE_RANGEDB_nn`. Don't leave a gap.

- Based on the example above, start at `BLOCK_DEVICE_RANGEDB_04`.
- In the example below, four new block storage volumes have been added to the node:
`BLOCK_DEVICE_RANGEDB_04` to `BLOCK_DEVICE_RANGEDB_07`.

```

NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/sgws-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/sgws-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/sgws-sn1-rangedb-3
BLOCK_DEVICE_RANGEDB_04 = /dev/mapper/sgws-sn1-rangedb-4
BLOCK_DEVICE_RANGEDB_05 = /dev/mapper/sgws-sn1-rangedb-5
BLOCK_DEVICE_RANGEDB_06 = /dev/mapper/sgws-sn1-rangedb-6
BLOCK_DEVICE_RANGEDB_07 = /dev/mapper/sgws-sn1-rangedb-7
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003
GRID_NETWORK_IP = 10.1.0.3
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1

```

9. Run the following command to validate your changes to the node configuration file for the Storage Node:

```
sudo storagegrid node validate <node-name>
```

Address any errors or warnings before proceeding to the next step.

If you observe an error similar to the following, it means that the node configuration file is attempting to map the block device used by <node-name> for <PURPOSE> to the given <path-name> in the Linux file system, but there is not a valid block device special file (or softlink to a block device special file) at that location.



```
Checking configuration file for node <node-name>...  
ERROR: BLOCK_DEVICE_<PURPOSE> = <path-name>  
<path-name> is not a valid block device
```

Verify that you entered the correct <path-name>.

10. Run the following command to restart the node with the new block device mappings in place:

```
sudo storagegrid node start <node-name>
```

11. Log in to the Storage Node as admin using the password listed in the `Passwords.txt` file.

12. Check that the services start correctly:

- a. View a list of the status of all services on the server:

```
sudo storagegrid-status
```

The status is updated automatically.

- b. Wait until all services are Running or Verified.

- c. Exit the status screen:

```
Ctrl+C
```

13. Configure the new storage for use by the Storage Node:

- a. Configure the new storage volumes:

```
sudo add_rangedbs.rb
```

This script finds any new storage volumes and prompts you to format them.

- b. Enter **y** to format the storage volumes.

- c. If any of the volumes have previously been formatted, decide if you want to reformat them.

- Enter **y** to reformat.

- Enter **n** to skip reformatting.

The `setup_rangedbs.sh` script runs automatically.

14. Verify that the Storage Node's storage state is online:

- a. Sign in to the Grid Manager using a [supported web browser](#).

- b. Select **SUPPORT > Tools > Grid topology**.

- c. Select **site > Storage Node > LDR > Storage**.

- d. Select the **Configuration** tab and then the **Main** tab.
 - e. If the **Storage State - Desired** drop-down list is set to Read-only or Offline, select **Online**.
 - f. Click **Apply Changes**.
15. To see the new object stores:
- a. Select **NODES > site > Storage Node > Storage**.
 - b. View the details in the **Object Stores** table.

Result

You can now use the expanded capacity of the Storage Nodes to save object data.

Add grid nodes or site

Add grid nodes to existing site or add new site

Follow this procedure to add grid nodes to existing sites or to add a new site. You can only perform one type of expansion at a time.

Before you begin

- You have the [Root access or Maintenance permission](#).
- All existing nodes in the grid are up and running across all sites.
- Any previous expansion, upgrade, decommissioning, or recovery procedures are complete.



You are prevented from starting an expansion while another expansion, upgrade, recovery, or active decommission procedure is in progress. However, if necessary, you can pause a decommission procedure to start an expansion.

Steps

1. [Update subnets for Grid Network](#).
2. [Deploy new grid nodes](#).
3. [Perform expansion](#).

Update subnets for Grid Network

When you add grid nodes or a new site in an expansion, you might need to update or add subnets to the Grid Network.

StorageGRID maintains a list of the network subnets used to communicate between grid nodes on the Grid Network (eth0). These entries include the subnets used for the Grid Network by each site in your StorageGRID system as well as any subnets used for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

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 the provisioning passphrase.

- You have the network addresses, in CIDR notation, of the subnets you want to configure.

About this task

If any of the new nodes has a Grid Network IP address on a subnet not previously used, you must add the new subnet to the Grid Network subnet list before starting the expansion. Otherwise, you will have to cancel the expansion, add the new subnet, and start the procedure again.



Do not use subnets that contain the following IPv4 addresses for the Grid Network, Admin Network, or Client Network of any node:

- 192.168.130.101
- 192.168.131.101
- 192.168.130.102
- 192.168.131.102
- 198.51.100.2
- 198.51.100.4

For example, do not use the following subnet ranges for the Grid Network, Admin Network, or Client Network of any node:

- 192.168.130.0/24 because this subnet range contains the IP addresses 192.168.130.101 and 192.168.130.102
- 192.168.131.0/24 because this subnet range contains the IP addresses 192.168.131.101 and 192.168.131.102
- 198.51.100.0/24 because this subnet range contains the IP addresses 198.51.100.2 and 198.51.100.4

Steps

1. Select **MAINTENANCE > Network > Grid Network**.
2. Select **Add another subnet** to add a new subnet in CIDR notation.

For example, enter 10.96.104.0/22.

3. Enter the provisioning passphrase, and select **Save**.
4. Wait until the changes are applied, then download a new Recovery Package.
 - a. Select **MAINTENANCE > System > Recovery package**.
 - b. Enter the **Provisioning Passphrase**.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system. It is also used to recover the primary Admin Node.

The subnets you have specified are configured automatically for your StorageGRID system.

Deploy new grid nodes

The steps for deploying new grid nodes in an expansion are the same as the steps used

when the grid was first installed. You must deploy all new grid nodes before you can perform the expansion.

When you expand a grid, the nodes you add don't have to match the existing node types. You can add VMware nodes, Linux container-based nodes, or appliance nodes.

VMware: Deploy grid nodes

You must deploy a virtual machine in VMware vSphere for each VMware node you want to add in the expansion.

Steps

1. [Deploy the new node as virtual machine](#) and connect it to one or more StorageGRID networks.

When you deploy the node, you can optionally remap node ports or increase CPU or memory settings.

2. After you have deployed all new VMware nodes, [perform the expansion procedure](#).

Linux: Deploy grid nodes

You can deploy grid nodes on new Linux hosts or on existing Linux hosts. If you need additional Linux hosts to support the CPU, RAM, and storage requirements of the StorageGRID nodes you want to add to your grid, you prepare them in the same way you prepared the hosts when you first installed them. Then, you deploy the expansion nodes in the same way you deployed grid nodes during installation.

Before you begin

- You have the instructions for installing StorageGRID for your version of Linux, and you have reviewed the hardware and storage requirements.
 - [Install StorageGRID on Red Hat Enterprise Linux](#)
 - [Install StorageGRID on Ubuntu or Debian](#)
- If you plan to deploy new grid nodes on existing hosts, you have confirmed the existing hosts have enough CPU, RAM, and storage capacity for the additional nodes.
- You have a plan to minimize failure domains. For example, you should not deploy all Gateway Nodes on a single physical host.



In a production deployment, don't run more than one Storage Node on a single physical or virtual host. Using a dedicated host for each Storage Node provides an isolated failure domain.

- If the StorageGRID node uses storage assigned from a NetApp ONTAP system, confirm that the volume does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.

Steps

1. If you are adding new hosts, access the installation instructions for deploying StorageGRID nodes.
2. To deploy the new hosts, follow the instructions for preparing the hosts.
3. To create node configuration files and to validate the StorageGRID configuration, follow the instructions for deploying grid nodes.
4. If you are adding nodes to a new Linux host, start the StorageGRID host service.

5. If you are adding nodes to an existing Linux host, start the new nodes using the storagegrid host service

```
CLI:sudo storagegrid node start [<node name>]
```

After you finish

After deploying all new grid nodes, you can [perform the expansion](#).

Appliances: Deploying Storage, Gateway, or non-primary Admin Nodes

To install the StorageGRID software on an appliance node, you use the StorageGRID Appliance Installer, which is included on the appliance. In an expansion, each storage appliance functions as a single Storage Node, and each services appliance functions as a single Gateway Node or non-primary Admin Node. Any appliance can connect to the Grid Network, the Admin Network, and the Client Network.

Before you begin

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- You have completed the [Set up hardware](#) steps.

Setting up appliance hardware includes the required steps for configuring StorageGRID connections (network links and IP addresses) as well the optional steps for enabling node encryption, changing the RAID mode, and remapping network ports.

- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.
- The StorageGRID Appliance Installer firmware on the replacement appliance is compatible with the StorageGRID software version currently running on your grid. If the versions aren't compatible, you must upgrade the StorageGRID Appliance Installer firmware.
- You have a service laptop with a [supported web browser](#).
- You know one of the IP addresses assigned to the appliance's compute controller. You can use the IP address for any attached StorageGRID network.

About this task

The process of installing StorageGRID on an appliance node has the following phases:

- You specify or confirm the IP address of the primary Admin Node and the name of the appliance node.
- You start the installation and wait as volumes are configured and the software is installed.

Partway through appliance installation tasks, the installation pauses. To resume the installation, you sign into the Grid Manager, approve all grid nodes, and complete the StorageGRID installation process.



If you need to deploy multiple appliance nodes at one time, you can automate the installation process by using the `configure-sga.py` Appliance Installation Script.

Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

```
https://Controller_IP:8443
```

The StorageGRID Appliance Installer Home page appears.

2. In the **Primary Admin Node** connection section, determine whether you need to specify the IP address for

the primary Admin Node.

If you have previously installed other nodes in this data center, the StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

3. If this IP address is not shown or you need to change it, specify the address:

Option	Description
Manual IP entry	a. Clear the Enable Admin Node discovery checkbox. b. Enter the IP address manually. c. Click Save. d. Wait for the connection state for the new IP address to become ready.
Automatic discovery of all connected primary Admin Nodes	a. Select the Enable Admin Node discovery checkbox. b. Wait for the list of discovered IP addresses to be displayed. c. Select the primary Admin Node for the grid where this appliance Storage Node will be deployed. d. Click Save. e. Wait for the connection state for the new IP address to become ready.

4. In the **Node name** field, enter the name you want to use for this appliance node, and select **Save**.

The node name is assigned to this appliance node in the StorageGRID system. It is shown on the Nodes page (Overview tab) in the Grid Manager. If required, you can change the name when you approve the node.

5. In the **Installation** section, confirm that the current state is "Ready to start installation of *node name* into grid with primary Admin Node *admin_ip*" and that the **Start Installation** button is enabled.

If the **Start Installation** button is not enabled, you might need to change the network configuration or port settings. For instructions, see the maintenance instructions for your appliance.

6. From the StorageGRID Appliance Installer home page, select **Start Installation**.

NetApp® StorageGRID® Appliance Installer

Home

Configure Networking ▾

Configure Hardware ▾

Monitor Installation

Advanced ▾

Home

The installation is ready to be started. Review the settings below, and then click Start Installation.

Primary Admin Node connection

Enable Admin Node discovery

☐

Primary Admin Node IP

172.16.4.210

Connection state

Connection to 172.16.4.210 ready

Cancel

Save

Node name

Node name

NetApp-SGA

Cancel

Save

Installation

Current state

Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

Start Installation

The Current state changes to "Installation is in progress," and the Monitor Installation page is displayed.

7. If your expansion includes multiple appliance nodes, repeat the previous steps for each appliance.



If you need to deploy multiple appliance Storage Nodes at one time, you can automate the installation process by using the configure-sga.py appliance installation script.

8. If you need to manually access the Monitor Installation page, select **Monitor Installation** from the menu bar.

The Monitor Installation page shows the installation progress.

1. Configure storage			Running
Step	Progress	Status	
Connect to storage controller		Complete	
Clear existing configuration		Complete	
Configure volumes		Creating volume StorageGRID-obj-00	
Configure host settings		Pending	
2. Install OS			Pending
3. Install StorageGRID			Pending
4. Finalize installation			Pending

The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.



The installer ensures that tasks completed in a previous install aren't re-run. If you are re-running an installation, any tasks that don't need to be re-run are shown with a green status bar and a status of "Skipped."

9. Review the progress of first two installation stages.

1. Configure appliance

During this stage, one of the following processes occurs:

- For a storage appliance, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity OS to configure volumes, and configures host settings.
- For a services appliance, the installer clears any existing configuration from the drives in the compute controller, and configures host settings.

2. Install OS

During this stage, the installer copies the base operating system image for StorageGRID to the appliance.

10. Continue monitoring the installation progress until a message appears in the console window, prompting you to use the Grid Manager to approve the node.



Wait until all nodes you added in this expansion are ready for approval before going to the Grid Manager to approve the nodes.

Monitor Installation

1. Configure storage	Complete
2. Install OS	Complete
3. Install StorageGRID	Running
4. Finalize installation	Pending

Connected (unencrypted) to: QEMU

```

/platform.type: Device or resource busy
[2017-07-31T22:09:12.362566] INFO -- [INSG] NOTICE: seeding /var/local with c
ontainer data
[2017-07-31T22:09:12.366205] INFO -- [INSG] Fixing permissions
[2017-07-31T22:09:12.369633] INFO -- [INSG] Enabling syslog
[2017-07-31T22:09:12.511533] INFO -- [INSG] Stopping system logging: syslog-n
g.
[2017-07-31T22:09:12.570096] INFO -- [INSG] Starting system logging: syslog-n
g.
[2017-07-31T22:09:12.576360] INFO -- [INSG] Beginning negotiation for downloa
d of node configuration
[2017-07-31T22:09:12.581363] INFO -- [INSG]
[2017-07-31T22:09:12.585066] INFO -- [INSG]
[2017-07-31T22:09:12.588314] INFO -- [INSG]
[2017-07-31T22:09:12.591851] INFO -- [INSG]
[2017-07-31T22:09:12.594886] INFO -- [INSG]
[2017-07-31T22:09:12.598360] INFO -- [INSG]
[2017-07-31T22:09:12.601324] INFO -- [INSG]
[2017-07-31T22:09:12.604759] INFO -- [INSG]
[2017-07-31T22:09:12.607800] INFO -- [INSG]
[2017-07-31T22:09:12.610985] INFO -- [INSG]
[2017-07-31T22:09:12.614597] INFO -- [INSG]
[2017-07-31T22:09:12.618282] INFO -- [INSG] Please approve this node on the A
dmin Node GMI to proceed...

```

Perform expansion

When you perform the expansion, the new grid nodes are added to your existing StorageGRID deployment.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the provisioning passphrase.
- You have deployed all of the grid nodes that are being added in this expansion.
- You have the [Maintenance or Root access permission](#).

- If you are adding Storage Nodes, you have confirmed that all data-repair operations performed as part of a recovery are complete. See [Check data repair jobs](#).
- If you are adding Storage Nodes and you want to assign a custom storage grade to those nodes, you have already [created the custom storage grade](#). You also have either the Root access permission or both the Maintenance and ILM permissions.
- If you are adding a new site, you have reviewed and updated ILM rules. You must ensure that object copies aren't stored to the new site until after the expansion is complete. For example, if a rule uses the default storage pool (**All Storage Nodes**), you must [create a new storage pool](#) that contains only the existing Storage Nodes and [update ILM rules](#) and the ILM policy to use that new storage pool. Otherwise, objects will be copied to the new site as soon as the first node at that site becomes active.

About this task

Performing the expansion includes these main user tasks:

1. Configure the expansion.
2. Start the expansion.
3. Download a new Recovery Package file.
4. Monitor the expansion steps and stages until all new nodes are installed and configured and all services have started.



Some expansion steps and stages might take a significant amount of time to run on a large grid. For example, streaming Cassandra to a new Storage Node might take only a few minutes if the Cassandra database is empty. However, if the Cassandra database includes a large amount of object metadata, this stage might take several hours or longer. Don't reboot any Storage Nodes during the either the "Expanding the Cassandra cluster" or "Starting Cassandra and streaming data" stages.

Steps

1. Select **MAINTENANCE > Tasks > Expansion**.

The Grid Expansion page appears. The Pending Nodes section lists the nodes that are ready to be added.

Grid Expansion

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Configure Expansion

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

+ Approve

✖ Remove

Search

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
00:50:56:a7:7a:c0	rlco-010-096-106-151	Storage Node	VMware VM	10.96.106.151/22
00:50:56:a7:0f:2e	rlco-010-096-106-156	API Gateway Node	VMware VM	10.96.106.156/22

2. Select **Configure Expansion**.

The Site Selection dialog box appears.

3. Select the type of expansion you are starting:

- If you are adding a new site, select **New**, and enter the name of the new site.
- If you are adding one or more nodes to an existing site, select **Existing**.

4. Select **Save**.

5. Review the **Pending Nodes** list, and confirm that it shows all of the grid nodes you deployed.

As required, you can position your cursor over a node's **Grid Network MAC Address** to see details about that node.

Pending Nodes

Grid nodes are listed as

+

Approve

×

Remove

Grid Network MAC Address

☐ 00:50:56:a7:7a:c0
 ☐ 00:50:56:a7:0f:2e

Approved Nodes

leo-010-096-106-151

Storage Node

Network

Grid Network	10.96.106.151/22	10.96.104.1
Admin Network	Name	Type
Client Network		

Hardware

VMware VM
 4 CPUs
 8 GB RAM

Disks

55 GB
 55 GB
 55 GB



If a node is missing, confirm that it was deployed successfully.

6. From the list of pending nodes, approve the nodes you want to add in this expansion.
 - a. Select the radio button next to the first pending grid node you want to approve.
 - b. Select **Approve**.

The grid node configuration form appears.

- c. As required, modify the general settings:

Field	Description
Site	The name of the site the grid node will be associated with. If you are adding multiple nodes, be sure to select the correct site for each node. If you are adding a new site, all nodes are added to the new site.
Name	The system name for the node. System names are required for internal StorageGRID operations and can't be changed.
Storage Type (Storage Nodes only)	• Data and metadata ("combined"): Object-data and metadata Storage Node • Data-only: Storage Node containing only object data (no metadata) • Metadata-only: Storage Node containing only metadata (no object data)

Field	Description
NTP Role	The Network Time Protocol (NTP) role of the grid node: • Select Automatic (default) to automatically assign the NTP role to the node. The Primary role will be assigned to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. The Client role will be assigned to all other grid nodes. • Select Primary to manually assign the Primary NTP role to the node. At least two nodes at each site should have the Primary role to provide redundant system access to external timing sources. • Select Client to manually assign the Client NTP role to the node.
ADC Service (combined or metadata-only Storage Nodes)	Whether this Storage Node will run the Administrative Domain Controller (ADC) service. The ADC service keeps track of the location and availability of grid services. At least three Storage Nodes at each site must include the ADC service. You can't add the ADC service to a node after it is deployed. • Select Yes if the Storage Node you are replacing includes the ADC service. Because you can't decommission a Storage Node if too few ADC services would remain, this ensures that a new ADC service is available before the old service is removed. • Select Automatic to let the system determine whether this node requires the ADC service. Learn about the ADC quorum.
Storage Grade (combined or data-only Storage Nodes)	Use the Default storage grade, or select the custom storage grade you want to assign to this new node. Storage grades are used by ILM storage pools, so your selection can affect which objects will be placed on the Storage Node.

d. As required, modify the settings for the Grid Network, Admin Network, and Client Network.

- **IPv4 Address (CIDR):** The CIDR network address for the network interface. For example: 172.16.10.100/24



If you discover that nodes have duplicate IP addresses on the Grid Network while you are approving nodes, you must cancel the expansion, redeploy the virtual machines or appliances with a non-duplicate IP, and restart the expansion.

- **Gateway:** The default gateway of the grid node. For example: 172.16.10.1
- **Subnets (CIDR):** One or more subnetworks for the Admin Network.

e. Select **Save**.

The approved grid node moves to the Approved Nodes list.

- To modify the properties of an approved grid node, select its radio button, and select **Edit**.
- To move an approved grid node back to the Pending Nodes list, select its radio button, and select **Reset**.
- To permanently remove an approved grid node, power the node off. Then, select its radio button, and select **Remove**.

f. Repeat these steps for each pending grid node you want to approve.



If possible, you should approve all pending grid notes and perform a single expansion. More time will be required if you perform multiple small expansions.

7. When you have approved all grid nodes, enter the **Provisioning Passphrase**, and select **Expand**.

After a few minutes, this page updates to display the status of the expansion procedure. When tasks that affect individual grid nodes are in progress, the Grid Node Status section lists the current status for each grid node.



During the "Installing grid nodes" step for a new appliance, the StorageGRID Appliance Installer shows installation moving from Stage 3 to Stage 4, Finalize Installation. When Stage 4 completes, the controller is rebooted.

Expansion Progress

Lists the status of grid configuration tasks required to change the grid topology. These grid configuration tasks are run automatically by the StorageGRID system.

1. Installing grid nodes

In Progress

Grid Node Status

Lists the installation and configuration status of each grid node included in the expansion.

Search

Name	Site	Grid Network IPv4 Address	Progress	Stage
rleo-010-096-106-151	Data Center 1	10.96.106.151/22		Waiting for Dynamic IP Service peers
rleo-010-096-106-156	Data Center 1	10.96.106.156/22		Waiting for NTP to synchronize

2. Initial configuration

Pending

3. Distributing the new grid node's certificates to the StorageGRID system.

Pending

4. Assigning Storage Nodes to storage grade

Pending

5. Starting services on the new grid nodes

Pending

6. Starting background process to clean up unused Cassandra keys

Pending



A site expansion includes an additional task to configure Cassandra for the new site.

8. As soon as the **Download Recovery Package** link appears, download the Recovery Package file.

You must download an updated copy of the Recovery Package file as soon as possible after making grid

topology changes to the StorageGRID system. The Recovery Package file allows you to restore the system if a failure occurs.

- a. Select the download link.
- b. Enter the provisioning passphrase, and select **Start Download**.
- c. When the download completes, open the `.zip` file and confirm that you can access the contents, including the `Passwords.txt` file.
- d. Copy the downloaded Recovery Package file (`.zip`) to two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

9. If you are adding Storage Nodes to an existing site or adding a site, monitor the Cassandra stages, which occur when services are started on the new grid nodes.



Don't reboot any Storage Nodes during either the "Expanding the Cassandra cluster" or "Starting Cassandra and streaming data" stages. These stages might take many hours to complete for each new Storage Node, especially if existing Storage Nodes contain a large amount of object metadata.

Adding Storage Nodes

If you are adding Storage Nodes to an existing site, review the percentage shown in the "Starting Cassandra and streaming data" status message.

5. Starting services on the new grid nodes

In Progress

Grid Node Status

Lists the installation and configuration status of each grid node included in the expansion.

⚠ Do not reboot any Storage Nodes during Step 4. The "Starting Cassandra and streaming data" stage might take hours, especially if existing Storage Nodes contain a large amount of object metadata.

Name	Site	Grid Network IPv4 Address	Progress	Stage
rleo-010-096-106-151	Data Center 1	10.96.106.151/22		Starting Cassandra and streaming data (20.4% streamed)
rleo-010-096-106-156	Data Center 1	10.96.106.156/22		Starting services

This percentage estimates how complete the Cassandra streaming operation is, based on the total amount of Cassandra data available and the amount that has already been written to the new node.

Adding site

If you are adding a new site, use `nodetool status` to monitor the progress of Cassandra streaming and to see how much metadata has been copied to the new site during the "Expanding the Cassandra cluster" stage. The total Data Load on the new site should be within about 20% of the total of a current site.

- Continue monitoring the expansion until all tasks are complete and the **Configure Expansion** button reappears.

After you finish

Depending on which types of grid nodes you added, perform additional integration and configuration steps. See [Configuration steps after expansion](#).

Configure expanded system

Configuration steps after expansion

After completing an expansion, you must perform additional integration and configuration steps.

About this task

You must complete the configuration tasks listed below for the grid nodes or sites you are adding in your expansion. Some tasks might be optional, depending on the options selected when installing and administering your system, and how you want to configure the nodes and sites added during the expansion.

Steps

1. If you added a site:

- [Create a storage pool](#) for the site and each storage grade you selected for the new Storage Nodes.
- Confirm that the ILM policy meets the new requirements. If rule changes are required, [create new rules](#) and [update the ILM policy](#). If the rules are already correct, [activate a new policy](#) with no rule changes to ensure StorageGRID uses the new nodes.
- Confirm that Network Time Protocol (NTP) servers are accessible from that site. See [Manage NTP servers](#).



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

2. If you added one or more Storage Nodes to an existing site:

- [View storage pool details](#) to confirm that each node you added is included in the expected storage pools and used in the expected ILM rules.
- Confirm that the ILM policy meets the new requirements. If rule changes are required, [create new rules](#) and [update the ILM policy](#). If the rules are already correct, [activate a new policy](#) with no rule changes to ensure StorageGRID uses the new nodes.
- [Verify that the Storage Node is active](#) and able to ingest objects.
- If you were unable to add the recommended number of Storage Nodes, rebalance erasure-coded data. See [Rebalance erasure-coded data after adding Storage Nodes](#).

3. If you added a Gateway Node:

- If high availability (HA) groups are used for client connections, optionally add the Gateway Node to an HA group. Select **CONFIGURATION > Network > High availability groups** to review the list of existing HA groups and to add the new node. See [Configure high availability groups](#).

4. If you added an Admin Node:

- a. If single sign-on (SSO) is enabled for your StorageGRID system, create a relying party trust for the new Admin Node. You can't sign in to the node until you create this relying party trust. See [Configure single sign-on](#).
- b. If you plan to use the Load Balancer service on Admin Nodes, optionally add the new Admin Node to an HA group. Select **CONFIGURATION > Network > High availability groups** to review the list of existing HA groups and to add the new node. See [Configure high availability groups](#).
- c. Optionally, copy the Admin Node database from the primary Admin Node to the expansion Admin Node if you want to keep the attribute and audit information consistent on each Admin Node. See [Copy the Admin Node database](#).
- d. Optionally, copy the Prometheus database from the primary Admin Node to the expansion Admin Node if you want to keep the historical metrics consistent on each Admin Node. See [Copy Prometheus metrics](#).
- e. Optionally, copy the existing audit logs from the primary Admin Node to the expansion Admin Node if you want to keep the historical log information consistent on each Admin Node. See [Copy audit logs](#).

5. To check if expansion nodes were added with an untrusted Client Network or to change whether a node's Client Network is untrusted or trusted, go to **CONFIGURATION > Security > Firewall control**.

If the Client Network on the expansion node is untrusted, then connections to the node on the Client Network must be made using a load balancer endpoint. See [Configure load balancer endpoints](#) and [Manage firewall controls](#).

6. Configure the DNS.

If you have been specifying DNS settings separately for each grid node, you must add custom per-node DNS settings for the new nodes. See [Modify DNS configuration for single grid node](#).

To ensure proper operation, specify two or three DNS servers. If you specify more than three, it is possible that only three will be used because of known OS limitations on some platforms. If you have routing restrictions in your environment, you can [customize the DNS server list](#) for individual nodes (typically all nodes at a site) to use a different set of up to three DNS servers.

If possible, use DNS servers that each site can access locally to ensure that an islanded site can resolve the FQDNs for external destinations.

Verify that Storage Node is active

After an expansion operation that adds new Storage Nodes completes, the StorageGRID system should automatically start using the new Storage Nodes. You must use the StorageGRID system to verify that the new Storage Node is active.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. Select **NODES > Expansion Storage Node > Storage**.
3. Position your cursor over the **Storage Used - Object Data** graph to view the value for **Used**, which is the amount of the Total usable space that has been used for object data.
4. Verify that the value of **Used** is increasing as you move your cursor to the right on the graph.

Copy Admin Node database

When adding Admin Nodes through an expansion procedure, you can optionally copy the database from the primary Admin Node to the new Admin Node. Copying the database allows you to retain historical information about attributes, alerts, and alerts.

Before you begin

- You have completed the required expansion steps to add an Admin Node.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

The StorageGRID software activation process creates an empty database for the NMS service on the expansion Admin Node. When the NMS service starts on the expansion Admin Node, it records information for servers and services that are currently part of the system or added later. This Admin Node database includes the following information:

- Alert history
- Historical attribute data, which is used in legacy-style charts on the Nodes page

To ensure that the Admin Node database is consistent between nodes, you can copy the database from the primary Admin Node to the expansion Admin Node.



Copying the database from the primary Admin Node (*thesource Admin Node*) to an expansion Admin Node can take up to several hours to complete. During this period, the Grid Manager is inaccessible.

Use these steps to stop the MI service and the Management API service on both the primary Admin Node and the expansion Admin Node before copying the database.

Steps

1. Complete the following steps on the primary Admin Node:
 - a. Log in to the Admin 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.
 - b. Run the following command: `recover-access-points`
 - c. Enter the provisioning passphrase.
 - d. Stop the MI service: `service mi stop`
 - e. Stop the Management Application Program Interface (mgmt-api) service: `service mgmt-api stop`
2. Complete the following steps on the expansion Admin Node:
 - a. Log in to the expansion Admin 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.
 - b. Stop the MI service: `service mi stop`
 - c. Stop the mgmt-api service: `service mgmt-api stop`
 - d. Add the SSH private key to the SSH agent. Enter: `ssh-add`
 - e. Enter the SSH Access Password listed in the `Passwords.txt` file.
 - f. Copy the database from the source Admin Node to the expansion Admin Node:
`/usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP`
 - g. When prompted, confirm that you want to overwrite the MI database on the expansion Admin Node.

The database and its historical data are copied to the expansion Admin Node. When the copy operation is done, the script starts the expansion Admin Node.
 - h. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter: `ssh-add -D`
3. Restart the services on the primary Admin Node: `service servermanager start`

Copy Prometheus metrics

After adding a new Admin Node, you can optionally copy the historical metrics maintained by Prometheus from the primary Admin Node to the new Admin Node. Copying the metrics ensures that historical metrics are consistent between Admin Nodes.

Before you begin

- The new Admin Node is installed and running.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

When you add an Admin Node, the software installation process creates a new Prometheus database. You can keep the historical metrics consistent between nodes by copying the Prometheus database from the primary Admin Node (the *source Admin Node*) to the new Admin Node.



Copying the Prometheus database might take an hour or more. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

Steps

1. Log in to the source Admin Node:
 - a. Enter the following command: `ssh admin@grid_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.
2. From the source Admin Node, stop the Prometheus service: `service prometheus stop`
3. Complete the following steps on the new Admin Node:
 - a. Log in to the new Admin 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.
 - b. Stop the Prometheus service: `service prometheus stop`
 - c. Add the SSH private key to the SSH agent. Enter: `ssh-add`
 - d. Enter the SSH Access Password listed in the `Passwords.txt` file.
 - e. Copy the Prometheus database from the source Admin Node to the new Admin Node:
`/usr/local/prometheus/bin/prometheus-clone-db.sh Source_Admin_Node_IP`
 - f. When prompted, press **Enter** to confirm that you want to destroy the new Prometheus database on the new Admin Node.

The original Prometheus database and its historical data are copied to the new Admin Node. When the copy operation is done, the script starts the new Admin Node. The following status appears:

```
Database cloned, starting services
```

- g. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter:

```
ssh-add -D
```

4. Restart the Prometheus service on the source Admin Node.

```
service prometheus start
```

Copy audit logs

When you add a new Admin Node through an expansion procedure, its AMS service only logs events and actions that occur after it joins the system. As required, you can copy audit logs from a previously installed Admin Node to the new expansion Admin Node so that it is in sync with the rest of the StorageGRID system.

Before you begin

- You have completed the required expansion steps to add an Admin Node.
- You have the `Passwords.txt` file.

About this task

To make historical audit messages available on a new Admin Node, you must copy the audit log files manually from an existing Admin Node to the expansion Admin Node.

By default, audit information is sent to the audit log on Admin Nodes. You can skip these steps if either of the following applies:



- You configured an external syslog server and audit logs are now being sent to the syslog server instead of to Admin Nodes.
- You explicitly specified that audit messages should be saved only on the local nodes that generated them.

See [Configure audit messages and log destinations](#) for details.

Steps

1. Log in to the primary Admin Node:

- a. Enter the following command: `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 `#`.

2. Stop the AMS service to prevent it from creating a new file: `service ams stop`
3. Navigate to the audit export directory:

```
cd /var/local/log
```

4. Rename the source `audit.log` file to ensure that it does not overwrite the file on the expansion Admin Node you are copying it to:

```
ls -l  
mv audit.log _new_name_.txt
```

5. Copy all audit log files to the destination location on the expansion Admin Node:

```
scp -p * IP_address:/var/local/log
```

6. If prompted for the passphrase for `/root/.ssh/id_rsa`, enter the SSH Access Password for the Primary Admin Node listed in the `Passwords.txt` file.

7. Restore the original `audit.log` file:

```
mv new_name.txt audit.log
```

8. Start the AMS service:

```
service ams start
```

9. Log out from the server:

```
exit
```

10. Log in to the expansion Admin Node:

- a. Enter the following command: `ssh admin@expansion_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 `#`.

11. Update the user and group settings for the audit log files:

```
cd /var/local/log  
  
chown ams-user:bycast *
```

12. Log out from the server:

```
exit
```

Rebalance erasure-coded data after adding Storage Nodes

After you add Storage Nodes, you can use the erasure coding (EC) rebalance procedure to redistribute erasure-coded fragments among the existing and new Storage Nodes.

Before you begin

- You have completed the expansion steps to add the new Storage Nodes.
- You have reviewed the [considerations for rebalancing erasure-coded data](#).
- You understand that replicated object data will not be moved by this procedure and that the EC rebalance procedure does not consider the replicated data usage on each Storage Node when determining where to move erasure-coded data.
- You have the `Passwords.txt` file.

What happens when this procedure runs

Before starting the procedure, note the following:

- The EC rebalance procedure will not start if one or more volumes are offline (unmounted) or if they are online (mounted) but in an error state.
- The EC rebalance procedure temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the rebalance is complete. If there is not enough storage for the reservation, the EC rebalance procedure will fail. Storage reservations are released when the EC rebalance procedure completes, whether the procedure failed or succeeded.
- If a volume goes offline while the EC rebalance procedure is in process, the rebalance procedure will terminate. Any data fragments that were already moved will remain in their new locations, and no data will be lost.

You can rerun the procedure after all volumes are back online.

- When the EC rebalance procedure is running, the performance of ILM operations and S3 client operations might be impacted.



S3 API operations to upload objects (or object parts) might fail during the EC rebalance procedure if they require more than 24 hours to complete. Long-duration PUT operations will fail if the applicable ILM rule uses Balanced or Strict placement on ingest. The following error will be reported: 500 Internal Server Error.

- During this procedure all nodes have a storage capacity limit of 80%. Nodes that exceed this limit, but still store below the target data partition, are excluded from:
 - The site imbalance value
 - Any job completion conditions



The target data partition is calculated by dividing the total data for a site by the number of nodes.

- **Job completion conditions.** The EC rebalance procedure is considered complete when any of the following is true:
 - It can't move any more erasure-coded data.
 - The data in all nodes is within a 5% deviation of the target data partition.
 - The procedure has been running for 30 days.

Steps

1. Review the current object storage details for the site you plan to rebalance.

- a. Select **NODES**.
 - b. Select the first Storage Node at the site.
 - c. Select the **Storage** tab.
 - d. Position your cursor over the Storage Used - Object Data chart to see the current amount of replicated data and erasure-coded data on the Storage Node.
 - e. Repeat these steps to view the other Storage Nodes at the site.
2. Log in to the primary Admin Node:
 - a. Enter the following command: `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 `#`.

3. Start the procedure:

```
`rebalance-data start --site "site-name"
```

For `"site-name"`, specify the first site where you added new Storage Node or nodes. Enclose `site-name` in quotes.

The EC rebalance procedure starts, and a job ID is returned.

4. Copy the job ID.
5. Monitor the status of the EC rebalance procedure.

- To view the status of a single EC rebalance procedure:

```
rebalance-data status --job-id job-id
```

For `job-id`, specify the ID that was returned when you started the procedure.

- To view the status of the current EC rebalance procedure and any previously completed procedures:

```
rebalance-data status
```



To get help on the `rebalance-data` command:

```
rebalance-data --help
```

6. Perform additional steps, based on the status returned:

- If `State` is `In progress`, the EC rebalance operation is still running. You should periodically monitor the procedure until it completes.

Use the `Site Imbalance` value to assess how unbalanced erasure-code data usage is across the Storage Nodes at the site. This value can range from 1.0 to 0, with 0 indicating that erasure-coding data usage is completely balanced across all Storage Nodes at the site.

The EC rebalance job is considered complete and will stop when the data in all nodes is within a 5%

deviation of the target data partition.

- If State is Success, optionally [review object storage](#) to see the updated details for the site.

Erasure-coded data should now be more balanced among the Storage Nodes at the site.

- If State is Failure:
 - a. Confirm that all Storage Nodes at the site are connected to the grid.
 - b. Check for and resolve any alerts that might be affecting these Storage Nodes.
 - c. Restart the EC rebalance procedure:

```
rebalance-data start --job-id job-id
```

- d. [View the status](#) of the new procedure. If State is still Failure, contact technical support.

7. If the EC rebalance procedure is generating too much load (for example, ingest operations are affected), pause the procedure.

```
rebalance-data pause --job-id job-id
```

8. If you need to terminate the EC rebalance procedure (for example, so you can perform a StorageGRID software upgrade), enter the following:

```
rebalance-data terminate --job-id job-id
```



When you terminate an EC rebalance procedure, any data fragments that have already been moved remain in their new locations. Data is not moved back to the original location.

9. If you are using erasure coding at more than one site, run this procedure for all other affected sites.

Troubleshoot expansion

If you encounter errors during the grid expansion process that you are unable to resolve, or if a grid task fails, collect the log files and contact technical support.

Before you contact technical support, collect the required log files to assist in troubleshooting.

Steps

1. Connect to the expansion node that has experienced failures:

- a. Enter the following command: `ssh -p 8022 admin@grid_node_IP`



Port 8022 is the SSH port of the base OS, while port 22 is the SSH port of the container engine running StorageGRID.

- 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.

After you log in as root, the prompt changes from `$` to `#`.

2. Depending on the stage the installation reached, retrieve any of the following logs that are available on the grid node:

Platform	Logs
VMware	• /var/log/daemon.log • /var/log/storagegrid/daemon.log • /var/log/storagegrid/nodes/<node-name>.log
Linux	• /var/log/storagegrid/daemon.log • /etc/storagegrid/nodes/<node-name>.conf (for each failed node) • /var/log/storagegrid/nodes/<node-name>.log (for each failed node; might not exist)

Maintain a StorageGRID system

Grid maintenance

Grid maintenance tasks include decommissioning a node or site, renaming a grid, node or site, and maintaining networks. You can also perform host and middleware procedures and grid node procedures.



In these instructions, "Linux" refers to a Red Hat® Enterprise Linux®, Ubuntu®, or Debian® deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool](#).

Before you begin

- You have a broad understanding of the StorageGRID system.
- You have reviewed your StorageGRID system's topology and you understand the grid configuration.
- You understand that you must follow all instructions exactly and heed all warnings.
- You understand that maintenance procedures not described aren't supported or require a services engagement.

Maintenance procedures for appliances

For hardware procedures, see the [maintenance instructions for your StorageGRID appliance](#).

Download Recovery Package

The Recovery Package file allows you to restore the StorageGRID system if a failure occurs.

Before you begin

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

Download the current Recovery Package file before making grid topology changes to the StorageGRID system or before upgrading software. Then, download a new copy of the Recovery Package after making grid topology changes or after upgrading software.

Steps

1. Select **MAINTENANCE > System > Recovery package**.
2. Enter the provisioning passphrase, and select **Start download**.

The download starts immediately.

3. When the download completes, open the `.zip` file and confirm that you can access the contents, including the `Passwords.txt` file.
4. Copy the downloaded Recovery Package file (`.zip`) to two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

Decommission nodes or site

Decommission node or site

You can perform a decommission procedure to permanently remove grid nodes or an entire site from the StorageGRID system.

To remove a grid node or a site, you perform one of the following decommission procedures:

- Perform a [grid node decommission](#) to remove one or more nodes, which can be at one or more sites. The nodes you remove can be online and connected to the StorageGRID system, or they can be offline and disconnected.
- Perform a [site decommission](#) to remove a site. You perform a **connected site decommission** if all nodes are connected to StorageGRID. You perform a **disconnected site decommission** if all nodes are disconnected from StorageGRID. If the site contains a mixture of connected and disconnected nodes, you must bring all offline nodes back online.



Before performing a disconnected site decommission, contact your NetApp account representative. NetApp will review your requirements before enabling all steps in the Decommission Site wizard. You should not attempt a disconnected site decommission if you believe it might be possible to recover the site or to recover object data from the site.

Decommission nodes

Grid node decommission

You can use the node decommission procedure to remove one or more grid nodes at one or more sites. You can't decommission the primary Admin Node.

When to decommission a node

Use the node decommission procedure when any of the following are true:

- You added a larger Storage Node in an expansion and you want to remove one or more smaller Storage Nodes, while at the same time preserving objects.



If you want to replace an older appliance with a newer appliance, consider [cloning the appliance node](#) instead of adding a new appliance in an expansion and then decommissioning the old appliance.

- You require less total storage.
- You no longer require a Gateway Node.
- You no longer require a non-primary Admin Node.
- Your grid includes a disconnected node that you can't recover or bring back online.

- Your grid includes an Archive Node.

How to decommission a node

You can decommission connected grid nodes or disconnected grid nodes.

Decommission connected nodes

In general, you should decommission grid nodes only when they are connected to the StorageGRID system and only when all nodes are in normal health (have green icons on the **NODES** pages and on the **Decommission Nodes** page).

For instructions, see [Decommission connected grid nodes](#).

Decommission disconnected nodes

In some cases, you might need to decommission a grid node that is not currently connected to the grid (one whose Health is Unknown or Administratively Down).

For instructions, see [Decommission disconnected grid nodes](#).

What to consider before decommissioning a node

Before performing either procedure, review the considerations for each type of node:

- [Considerations for Admin or Gateway Node decommission](#)
- [Considerations for Storage Node decommission](#)

Considerations for decommissioning Admin or Gateway Nodes

Review the considerations for decommissioning an Admin Node or Gateway Node.

Considerations for Admin Node

- You can't decommission the primary Admin Node.
- You can't decommission an Admin Node if one of its network interfaces is part of a high availability (HA) group. You must first remove the network interfaces from the HA group. See the instructions for [managing HA groups](#).
- As required, you can safely change ILM policies while decommissioning an Admin Node.
- If you decommission an Admin Node and single sign-on (SSO) is enabled for your StorageGRID system, you must remember to remove the node's relying party trust from Active Directory Federation Services (AD FS).
- If you use [grid federation](#), ensure that the IP address of the node you are decommissioning was not specified for a grid federation connection.
- When you decommission a disconnected Admin Node, you will lose the audit logs from that node; however, these logs should also exist on the primary Admin Node.

Considerations for Gateway Node

- You can't decommission a Gateway Node if one of its network interfaces is part of a high availability (HA) group. You must first remove the network interfaces from the HA group. See the instructions for [managing HA groups](#).
- As required, you can safely change ILM policies while decommissioning a Gateway Node.

- If you use [grid federation](#), ensure that the IP address of the node you are decommissioning was not specified for a grid federation connection.
- You can safely decommission a Gateway Node while it is disconnected.

Considerations for Storage Nodes

Considerations for decommissioning Storage Nodes

Before decommissioning a Storage Node, consider whether you can clone the node instead. Then, if you do decide to decommission the node, review how StorageGRID manages objects and metadata during the decommission procedure.

When to clone a node instead of decommissioning it

If you want to replace an older appliance Storage Node with a newer or larger appliance, consider cloning the appliance node instead of adding a new appliance in an expansion and then decommissioning the old appliance.

Appliance node cloning lets you easily replace an existing appliance node with a compatible appliance at the same StorageGRID site. The cloning process transfers all data to the new appliance, places the new appliance in service, and leaves the old appliance in a pre-install state.

You can clone an appliance node if you need to:

- Replace an appliance that is reaching end-of-life.
- Upgrade an existing node to take advantage of improved appliance technology.
- Increase grid storage capacity without changing the number of Storage Nodes in your StorageGRID system.
- Improve storage efficiency, such as by changing the RAID mode.

See [Appliance node cloning](#) for details.

Considerations for connected Storage Nodes

Review the considerations for decommissioning a connected Storage Node.

- You should not decommission more than 10 Storage Nodes in a single Decommission Node procedure.
- The system must, at all times, include enough Storage Nodes to satisfy operational requirements, including the [ADC quorum](#) and the active [ILM policy](#). To satisfy this restriction, you might need to add a new Storage Node in an expansion operation before you can decommission an existing Storage Node.

Use caution when you decommission Storage Nodes in a grid containing software-based metadata-only nodes. If you decommission all nodes configured to store *both* objects and metadata, the ability to store objects is removed from the grid. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

- When you remove a Storage Node, large volumes of object data are transferred over the network. Although these transfers should not affect normal system operations, they can affect the total amount of network bandwidth consumed by the StorageGRID system.
- Tasks associated with Storage Node decommissioning are given a lower priority than tasks associated with normal system operations. This means that decommissioning does not interfere with normal StorageGRID

system operations, and does not need to be scheduled for a period of system inactivity. Because decommissioning is performed in the background, it is difficult to estimate how long the process will take to complete. In general, decommissioning finishes more quickly when the system is quiet, or if only one Storage Node is being removed at a time.

- It might take days or weeks to decommission a Storage Node. Plan this procedure accordingly. While the decommission process is designed to not impact system operations, it can limit other procedures. In general, you should perform any planned system upgrades or expansions before you remove grid nodes.
- If you need to perform another maintenance procedure while Storage Nodes are being removed, you can [pause the decommission procedure](#) and resume it after the other procedure is complete.



The **Pause** button is enabled only when the ILM evaluation or erasure-coded data decommissioning stages are reached; however, ILM evaluation (data migration) will continue to run in the background.

- You can't run data repair operations on any grid nodes when a decommission task is running.
- You should not make any changes to an ILM policy while a Storage Node is being decommissioned.
- To permanently and securely remove data, you must wipe the Storage Node's drives after the decommission procedure is complete.

Considerations for disconnected Storage Nodes

Review the considerations for decommissioning a disconnected Storage Node.

- Never decommission a disconnected node unless you are sure it can't be brought online or recovered.



Don't perform this procedure if you believe it might be possible to recover object data from the node. Instead, contact technical support to determine if node recovery is possible.

- When you decommission a disconnected Storage Node, StorageGRID uses data from other Storage Nodes to reconstruct the object data and metadata that was on the disconnected node.
- Data loss might occur if you decommission more than one disconnected Storage Node. The system might not be able to reconstruct data if not enough object copies, erasure-coded fragments, or object metadata remain available. When decommissioning Storage Nodes in a grid with software-based metadata-only nodes, decommissioning all nodes configured to store both objects and metadata removes all object storage from the grid. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.



If you have more than one disconnected Storage Node that you can't recover, contact technical support to determine the best course of action.

- When you decommission a disconnected Storage Node, StorageGRID starts data repair jobs at the end of the decommissioning process. These jobs attempt to reconstruct the object data and metadata that was stored on the disconnected node.
- When you decommission a disconnected Storage Node, the decommission procedure completes relatively quickly. However, the data repair jobs can take days or weeks to run and aren't monitored by the decommission procedure. You must manually monitor these jobs and restart them as needed. See [Check data repair jobs](#).
- If you decommission a disconnected Storage Node that contains the only copy of an object, the object will be lost. The data repair jobs can only reconstruct and recover objects if at least one replicated copy or enough erasure-coded fragments exist on Storage Nodes that are currently connected.

What is the ADC quorum?

You might not be able to decommission certain Storage Nodes at a site if too few Administrative Domain Controller (ADC) services would remain after the decommissioning.

The ADC service, which is found on some Storage Nodes, maintains grid topology information and provides configuration services to the grid. The StorageGRID system requires a quorum of ADC services to be available at each site and at all times.

You can't decommission a Storage Node if removing the node would cause the ADC quorum to no longer be met. To satisfy the ADC quorum during a decommissioning, a minimum of three Storage Nodes at each site must have the ADC service. If a site has more than three Storage Nodes with the ADC service, a simple majority of those nodes must remain available after the decommissioning: $((0.5 * \text{Storage Nodes with ADC}) + 1)$



Use caution when you decommission Storage Nodes in a grid containing software-based metadata-only nodes. If you decommission all nodes configured to store *both* objects and metadata, the ability to store objects is removed from the grid. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

For example, suppose a site currently includes six Storage Nodes with ADC services and you want to decommission three Storage Nodes. Because of the ADC quorum requirement, you must complete two decommission procedures, as follows:

- In the first decommission procedure, you must ensure that four Storage Nodes with ADC services remain available: $((0.5 * 6) + 1)$. This means that you can only decommission two Storage Nodes initially.
- In the second decommission procedure, you can remove the third Storage Node because the ADC quorum now only requires three ADC services to remain available: $((0.5 * 4) + 1)$.

If you need to decommission a Storage Node but are unable to because of the ADC quorum requirement, add a new Storage Node in an [expansion](#) and specify that it should have an ADC service. Then, decommission the existing Storage Node.

Review ILM policy and storage configuration

If you plan to decommission a Storage Node, you should review your StorageGRID system's ILM policy before starting the decommissioning process.

During decommissioning, all object data is migrated from the decommissioned Storage Node to other Storage Nodes.



The ILM policy you have *during* the decommission will be the one used *after* the decommission. You must ensure this policy meets your data requirements both before you start the decommission and after the decommission is complete.

You should review the rules in each [active ILM policy](#) to ensure that the StorageGRID system will continue to have enough capacity of the correct type and in the correct locations to accommodate the decommissioning of a Storage Node.

Consider the following:

- Will it be possible for ILM evaluation services to copy object data such that ILM rules are satisfied?
- What happens if a site becomes temporarily unavailable while decommissioning is in progress? Can additional copies be made in an alternate location?
- How will the decommissioning process affect the final distribution of content? As described in [Consolidate Storage Nodes](#), you should [add new Storage Nodes](#) before decommissioning old ones. If you add a larger replacement Storage Node after decommissioning a smaller Storage Node, the old Storage Nodes could be close to capacity and the new Storage Node could have almost no content. Most write operations for new object data would then be directed at the new Storage Node, reducing the overall efficiency of system operations.
- Will the system, at all times, include enough Storage Nodes to satisfy the active ILM policies?



An ILM policy that can't be satisfied will lead to backlogs and alerts, and might halt operation of the StorageGRID system.

Verify that the proposed topology that will result from the decommissioning process satisfies the ILM policy by assessing the areas listed in the table.

Area to assess	What to consider
Available capacity	Will there be enough storage capacity to accommodate all of the object data stored in the StorageGRID system, including the permanent copies of object data currently stored on the Storage Node to be decommissioned? Will there be enough capacity to handle the anticipated growth in stored object data for a reasonable interval of time after decommissioning is complete?
Location of storage	If enough capacity remains in the StorageGRID system as a whole, is the capacity in the right locations to satisfy the StorageGRID system's business rules?
Storage type	Will there be enough storage of the appropriate type after decommissioning is complete? For example, ILM rules might move content from one type of storage to another as content ages. In this case, you must ensure that enough storage of the appropriate type is available in the final configuration of the StorageGRID system.

Consolidate Storage Nodes

You can consolidate Storage Nodes to reduce the Storage Node count for a site or deployment while increasing storage capacity.

When you consolidate Storage Nodes, you [expand the StorageGRID system](#) by adding new, larger capacity Storage Nodes and then decommission the old, smaller capacity Storage Nodes. During the decommission procedure, objects are migrated from the old Storage Nodes to the new Storage Nodes.



If you are consolidating older and smaller appliances with new models or larger capacity appliances, consider [cloning the appliance node](#) (or use appliance node cloning and the decommission procedure if you aren't doing a one-to-one replacement).

For example, you might add two new, larger capacity Storage Nodes to replace three older Storage Nodes. You would first use the expansion procedure to add the two new, larger Storage Nodes, and then use the decommission procedure to remove the three old, smaller capacity Storage Nodes.

By adding new capacity before removing existing Storage Nodes, you ensure a more balanced distribution of data across the StorageGRID system. You also reduce the possibility that an existing Storage Node might be pushed beyond the storage watermark level.

Decommission multiple Storage Nodes

If you need to remove more than one Storage Node, you can decommission them either sequentially or in parallel.



Use caution when you decommission Storage Nodes in a grid containing software-based metadata-only nodes. If you decommission all nodes configured to store *both* objects and metadata, the ability to store objects is removed from the grid. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

- If you decommission Storage Nodes sequentially, you must wait for the first Storage Node to complete decommissioning before starting to decommission the next Storage Node.
- If you decommission Storage Nodes in parallel, the Storage Nodes simultaneously process decommission tasks for all Storage Nodes being decommissioned. This can result in a situation where all permanent copies of a file are marked as "read-only," temporarily disabling deletion in grids where this functionality is enabled.

Check data repair jobs

Before decommissioning a grid node, you must confirm that no data repair jobs are active. If any repairs have failed, you must restart them and allow them to complete before performing the decommission procedure.

About this task

If you need to decommission a disconnected Storage Node, you will also complete these steps after the decommission procedure completes to ensure the data repair job has completed successfully. You must ensure that any erasure-coded fragments that were on the removed node have been restored successfully.

These steps only apply to systems that have erasure-coded objects.

Steps

1. Log in to the primary Admin Node:
 - a. Enter the following command: `ssh admin@grid_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 #.

2. Check for running repairs: `repair-data show-ec-repair-status`

- If you have never run a data repair job, the output is `No job found`. You don't need to restart any repair jobs.
- If the data repair job was run previously or is running currently, the output lists information for the repair. Each repair has a unique repair ID.

```
root@ADM1-0:~# repair-data show-ec-repair-status
```

Repair ID	Affected Nodes / Volumes	Start Time	End Time	State	Estimated Bytes Affected	Bytes Repaired	Percentage
4216507958013005550	DC1-S1-0-182 (Volumes: 2)	2022-08-17T21:37:30.051543	2022-08-17T21:37:37.328998	Completed	1015788876	0	0
18214680851049518682	DC1-S1-0-182 (Volumes: 1)	2022-08-17T20:37:58.869362	2022-08-17T20:38:45.299688	Completed	0	0	100
7962734388032289010	DC1-S1-0-182 (Volumes: 0)	2022-08-17T20:42:29.578740		Stopped			Unknown



Optionally, you can use the Grid Manager to monitor restoration processes in progress and display a restoration history. See [Restore object data using Grid Manager](#).

3. If the State for all repairs is `Completed`, you don't need to restart any repair jobs.

4. If the State for any repair is `Stopped`, you must restart that repair.

- Obtain the repair ID for the failed repair from the output.
- Run the `repair-data start-ec-node-repair` command.

Use the `--repair-id` option to specify the Repair ID. For example, if you want to retry a repair with repair ID 949292, run this command: `repair-data start-ec-node-repair --repair-id 949292`

- Continue to track the status of EC data repairs until the State for all repairs is `Completed`.

Gather required materials

Before performing a grid node decommission, you must obtain the following information.

Item	Notes
Recovery Package .zip file	You must download the most recent Recovery Package .zip file (<code>sgws-recovery-package-id-revision.zip</code>). You can use the Recovery Package file to restore the system if a failure occurs.
Passwords.txt file	This file contains the passwords required to access grid nodes on the command line and is included in the Recovery Package.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the Passwords.txt file.
Description of StorageGRID system's topology before decommissioning	If available, obtain any documentation that describes the system's current topology.

Related information

[Web browser requirements](#)

Access Decommission Nodes page

When you access the Decommission Nodes page in the Grid Manager, you can see at a glance which nodes can be decommissioned.

Before you begin

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



Use caution when you decommission Storage Nodes in a grid containing software-based metadata-only nodes. If you decommission all nodes configured to store *both* objects and metadata, the ability to store objects is removed from the grid. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

Steps

1. Select **MAINTENANCE > Tasks > Decommission**.
2. Select **Decommission Nodes**.

The Decommission Nodes page appears. From this page, you can:

- Determine which grid nodes can be decommissioned currently.
- See the health of all grid nodes
- Sort the list in ascending or descending order by **Name**, **Site**, **Type**, or **Has ADC**.
- Enter search terms to quickly find particular nodes.

In this example, the Decommission Possible column indicates that you can decommission the Gateway Node and one of the four Storage Nodes.

Name	Site	Type	Has ADC	Health	Decommission Possible
DC1-ADM1	Data Center 1	Admin Node	-		No, member of HA group(s): HAGroup. Before you can decommission this node, you must remove it from all HA groups.
DC1-ARC1	Data Center 1	Archive Node	-		No, you can't decommission an Archive Node unless the node is disconnected.
☐ DC1-G1	Data Center 1	API Gateway Node	-		
DC1-S1	Data Center 1	Storage Node	Yes		No, site Data Center 1 requires a minimum of 3 Storage Nodes with ADC services.
DC1-S2	Data Center 1	Storage Node	Yes		No, site Data Center 1 requires a minimum of 3 Storage Nodes with ADC services.
DC1-S3	Data Center 1	Storage Node	Yes		No, site Data Center 1 requires a minimum of 3 Storage Nodes with ADC services.
☐ DC1-S4	Data Center 1	Storage Node	No		

3. Review the **Decommission Possible** column for each node you want to decommission.

If a grid node can be decommissioned, this column includes a green check mark, and the left column includes a checkbox. If a node can't be decommissioned, this column describes the issue. If there is more than one reason a node can't be decommissioned, the most critical reason is shown.

Decommission Possible reason	Description	Steps to resolve
No, <i>node type</i> decommissioning is not supported.	You can't decommission the primary Admin Node.	None.
No, at least one grid node is disconnected. Note: This message is shown for connected grid nodes only.	You can't decommission a connected grid node if any grid node is disconnected. The Health column includes one of these icons for grid nodes that are disconnected:  (gray): Administratively Down  (blue): Unknown	You must bring all disconnected nodes back online or decommission all disconnected nodes before you can remove a connected node. Note: If your grid contains multiple disconnected nodes, the software requires you to decommission them all at the same time, which increases the potential for unexpected results.
No, one or more required nodes is currently disconnected and must be recovered. Note: This message is shown for disconnected grid nodes only.	You can't decommission a disconnected grid node if one or more required nodes is also disconnected (for example, a Storage Node that is required for the ADC quorum).	- Review the Decommission Possible messages for all disconnected nodes. - Determine which nodes can't be decommissioned because they are required. - If the Health of a required node is Administratively Down, bring the node back online. - If the health of a required node is Unknown, perform a node recovery procedure to recover the required node.
No, member of HA group(s): <i>group name</i> . Before you can decommission this node, you must remove it from all HA groups.	You can't decommission an Admin Node or a Gateway Node if a node interface belongs to a high availability (HA) group.	Edit the HA group to remove the node's interface or remove the entire HA group. See Configure high availability groups .
No, site x requires a minimum of <i>n</i> Storage Nodes with ADC services.	Storage Nodes only. You can't decommission a Storage Node if insufficient nodes would remain at the site to support ADC quorum requirements.	Perform an expansion. Add a new Storage Node to the site, and specify that it should have an ADC service. See information about the ADC quorum .

Decommission Possible reason	Description	Steps to resolve
No, one or more erasure-coding profiles need at least n Storage Nodes. If the profile is not used in an ILM rule, you can deactivate it.	Storage Nodes only. You can't decommission a Storage Node unless enough nodes would remain for the existing erasure-coding profiles. For example, if an erasure-coding profile exists for 4+2 erasure coding, at least 6 Storage Nodes must remain.	For each affected erasure-coding profile, perform one of the following steps, based on how the profile is being used: • Used in active ILM policies: Perform an expansion. Add enough new Storage Nodes to allow erasure coding to continue. See the instructions for expanding your grid. • Used in an ILM rule but not in active ILM policies: Edit or delete the rule and then deactivate the erasure-coding profile. • Not used in any ILM rule: Deactivate the erasure-coding profile. Note: An error message appears if you attempt to deactivate an erasure-coding profile and object data is still associated with the profile. You might need to wait several weeks before trying the deactivation process again. Learn about deactivating an erasure-coding profile.
No, you can't decommission an Archive Node unless the node is disconnected.	If an Archive Node is still connected, you can't remove it.	Note: Support for Archive Nodes has been removed. If you need to decommission an Archive Node, see Grid node decommissioning (StorageGRID 11.8 doc site)

Decommission disconnected grid nodes

You might need to decommission a node that is not currently connected to the grid (one whose Health is Unknown or Administratively Down).

Before you begin

- You understand the considerations for decommissioning [Admin and Gateway Nodes](#) and the considerations for decommissioning [Storage Nodes](#).
- You have obtained all prerequisite items.
- You have ensured that no data repair jobs are active. See [Check data repair jobs](#).

- You have confirmed that Storage Node recovery is not in progress anywhere in the grid. If it is, you must wait until any Cassandra rebuild performed as part of the recovery is complete. You can then proceed with decommissioning.
- You have ensured that other maintenance procedures will not be run while the node decommission procedure is running, unless the node decommission procedure is paused.
- The **Decommission Possible** column for the disconnected node or nodes you want to decommission includes a green check mark.
- You have the provisioning passphrase.

About this task

You can identify disconnected nodes by looking for the blue Unknown icon  or the gray Administratively down icon  in the **Health** column.

Before decommissioning any disconnected node, note the following:

- This procedure is primarily intended for removing a single disconnected node. If your grid contains multiple disconnected nodes, the software requires you to decommission them all at the same time, which increases the potential for unexpected results.



Data loss might occur if you decommission more than one disconnected Storage Node at a time. See [Considerations for disconnected Storage Nodes](#).



Use caution when you decommission Storage Nodes in a grid containing software-based metadata-only nodes. If you decommission all nodes configured to store *both* objects and metadata, the ability to store objects is removed from the grid. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

- If a disconnected node can't be removed (for example, a Storage Node that is required for the ADC quorum), no other disconnected node can be removed.

Steps

1. Unless you are decommissioning an Archive Node (which must be disconnected), attempt to bring any disconnected grid nodes back online or recover them.

See [Grid node recovery procedures](#) for instructions.

2. If you are unable to recover a disconnected grid node and you want to decommission it while it is disconnected, select the checkbox for that node.



If your grid contains multiple disconnected nodes, the software requires you to decommission them all at the same time, which increases the potential for unexpected results.



Be careful when choosing to decommission more than one disconnected grid node at a time, especially if you are selecting multiple disconnected Storage Nodes. If you have more than one disconnected Storage Node that you can't recover, contact technical support to determine the best course of action.

3. Enter the provisioning passphrase.

The **Start Decommission** button is enabled.

4. Click **Start Decommission**.

A warning appears, indicating that you have selected a disconnected node and that object data will be lost if the node has the only copy of an object.

5. Review the list of nodes, and click **OK**.

The decommission procedure starts, and the progress is displayed for each node. During the procedure, a new Recovery Package is generated containing the grid configuration change.

6. As soon as the new Recovery Package is available, click the link or select **MAINTENANCE > System > Recovery package** to access the Recovery Package page. Then, download the .zip file.

See the instructions for [downloading the Recovery Package](#).



Download the Recovery Package as soon as possible to ensure you can recover your grid if something goes wrong during the decommission procedure.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

7. Periodically monitor the Decommission page to ensure that all selected nodes are decommissioned successfully.

Storage Nodes can take days or weeks to decommission. When all tasks are complete, the node selection list is redisplayed with a success message. If you decommissioned a disconnected Storage Node, an information message indicates that the repair jobs have been started.

8. After the nodes have shut down automatically as part of the decommission procedure, remove any remaining virtual machines or other resources that are associated with the decommissioned node.



Don't perform this step until the nodes have shut down automatically.

9. If you are decommissioning a Storage Node, monitor the status of the **replicated data** and **erasure-coded (EC) data** repair jobs that are automatically started during the decommissioning process.

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the `repair-data` command.

```
repair-data show-replicated-repair-status
```

- To determine if repairs are complete:
 1. Select **NODES > Storage Node being repaired > ILM**.
 2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
 1. Select **SUPPORT > Tools > Grid topology**.
 2. Select **grid > Storage Node being repaired > LDR > Data Store**.
 3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.



Cassandra inconsistencies might be present, and failed repairs aren't tracked.

- **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.



High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

- **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure-coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
 - Select **SUPPORT > Tools > Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.

- Use this command to see the status of a specific `repair-data` operation:

```
repair-data show-ec-repair-status --repair-id repair ID
```

- Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

This command retries a failed node repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-node-repair --repair-id 6949309319275667690
```

This command retries a failed volume repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-volume-repair --repair-id 6949309319275667690
```

After you finish

As soon as the disconnected nodes have been decommissioned and all data repair jobs have been completed, you can decommission any connected grid nodes as required.

Then, complete these steps after you complete the decommission procedure:

- Ensure that the drives of the decommissioned grid node are wiped clean. Use a commercially available data wiping tool or service to permanently and securely remove data from the drives.
- If you decommissioned an appliance node and the data on the appliance was protected using node encryption, use the StorageGRID Appliance Installer to clear the key management server configuration (Clear KMS). You must clear the KMS configuration if you want to add the appliance to another grid. For instructions, see [Monitor node encryption in maintenance mode](#).

Decommission connected grid nodes

You can decommission and permanently remove nodes that are connected to the grid.

Before you begin

- You understand the considerations for decommissioning [Admin and Gateway Nodes](#) and the considerations for decommissioning [Storage Nodes](#).
- You have gathered all required materials.
- You have ensured that no data repair jobs are active.
- You have confirmed that Storage Node recovery is not in progress anywhere in the grid. If it is, wait until any Cassandra rebuild performed as part of the recovery is complete. You can then proceed with decommissioning.
- You have ensured that other maintenance procedures will not be run while the node decommission procedure is running, unless the node decommission procedure is paused.
- You have the provisioning passphrase.
- Grid nodes are connected.
- The **Decommission Possible** column for the node or nodes you want to decommission includes a green check mark.



The decommission will not start if one or more volumes are offline (unmounted) or if they are online (mounted) but in an error state.



If one or more volumes go offline while a decommission is in progress, the decommission process completes after these volumes have come back online.

- All grid nodes have Normal (green) health . If you see one of these icons in the **Health** column, you must try to resolve the issue:

Icon	Color	Severity
	Yellow	Notice
	Light orange	Minor
	Dark orange	Major
	Red	Critical

- If you previously decommissioned a disconnected Storage Node, the data repair jobs have all completed successfully. See [Check data repair jobs](#).



Don't remove a grid node's virtual machine or other resources until instructed to do so in this procedure.



Use caution when you decommission Storage Nodes in a grid containing software-based metadata-only nodes. If you decommission all nodes configured to store *both* objects and metadata, the ability to store objects is removed from the grid. See [Types of Storage Nodes](#) for more information about metadata-only Storage Nodes.

About this task

When a node is decommissioned, its services are disabled and the node automatically shut down.

Steps

1. From the Decommission Nodes page, select the checkbox for each grid node you want to decommission.
2. Enter the provisioning passphrase.

The **Start Decommission** button is enabled.

3. Select **Start Decommission**.
4. Review the list of nodes in the confirmation dialog, and select **OK**.

The node decommission procedure starts, and the progress is displayed for each node.



Don't take a Storage Node offline after the decommission procedure has started. Changing the state might result in some content not being copied to other locations.

5. As soon as the new Recovery Package is available, select the Recovery Package link in the banner or select **MAINTENANCE > System > Recovery package** to access the Recovery Package page. Then,

download the .zip file.

See [downloading the Recovery Package](#).



Download the Recovery Package as soon as possible to ensure you can recover your grid if something goes wrong during the decommission procedure.

6. Periodically monitor the Decommission Nodes page to ensure that all selected nodes are decommissioned successfully.



Storage Nodes can take days or weeks to decommission.

When all tasks are complete, the node selection list is redisplayed with a success message.

After you finish

Complete these steps after you complete the node decommission procedure:

1. Follow the appropriate step for your platform. For example:
 - **Linux:** You might want to detach the volumes and delete the node configuration files you created during installation. See [Install StorageGRID on Red Hat Enterprise Linux](#) and [Install StorageGRID on Ubuntu or Debian](#).
 - **VMware:** You might want to use the vCenter "Delete from Disk" option to delete the virtual machine. You might also need to delete any data disks that are independent of the virtual machine.
 - **StorageGRID appliance:** The appliance node automatically reverts to an undeployed state where you can access the StorageGRID Appliance Installer. You can power off the appliance or add it to another StorageGRID system.
2. Ensure that the drives of the decommissioned grid node are wiped clean. Use a commercially available data wiping tool or service to permanently and securely remove data from the drives.
3. If you decommissioned an appliance node and the data on the appliance was protected using node encryption, use the StorageGRID Appliance Installer to clear the key management server configuration (Clear KMS). You must clear the KMS configuration if you want to add the appliance to another grid. For instructions, see [Monitor node encryption in maintenance mode](#).

Pause and resume decommission process for Storage Nodes

If you need to perform a second maintenance procedure, you can pause the decommission procedure for a Storage Node during certain stages. After the other procedure is finished, you can resume decommissioning.



The **Pause** button is enabled only when the ILM evaluation or erasure-coded data decommissioning stages are reached; however, ILM evaluation (data migration) will continue to run in the background.

Before you begin

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

Steps

1. Select **MAINTENANCE > Tasks > Decommission**.

The Decommission page appears.

2. Select **Decommission Nodes**.

The Decommission Nodes page appears. When the decommission procedure reaches either of the following stages, the **Pause** button is enabled.

- Evaluating ILM
- Decommissioning Erasure-Coded Data

3. Select **Pause** to suspend the procedure.

The current stage is paused, and the **Resume** button is enabled.

Decommission Nodes

 A new Recovery Package has been generated as a result of the configuration change. Go to the [Recovery Package page](#) to download it.

 Decommissioning procedure has been paused. Click 'Resume' to resume the procedure.

The progress for each node is displayed while the decommission procedure is running. When all tasks are complete, the node selection list is redisplayed.

Search 				
Name 	Type 	Progress 	Stage 	
DC1-S5	Storage Node		Evaluating ILM	

Pause

Resume

4. After the other maintenance procedure is finished, select **Resume** to proceed with the decommission.

Decommission site

Considerations for removing a site

Before using the site decommission procedure to remove a site, you must review the considerations.

What happens when you decommission a site

When you decommission a site, StorageGRID permanently removes all nodes at the site and the site itself from the StorageGRID system.

When the site decommission procedure is complete:

- You can no longer use StorageGRID to view or access the site or any of the nodes at the site.
- You can no longer use any storage pools or erasure-coding profiles that referred to the site. When StorageGRID decommissions a site, it automatically removes these storage pools and deactivates these erasure-coding profiles.

Differences between connected site and disconnected site decommission procedures

You can use the site decommission procedure to remove a site in which all nodes are connected to StorageGRID (referred to as a connected site decommission) or to remove a site in which all nodes are disconnected from StorageGRID (referred to as a disconnected site decommission). Before you begin, you must understand the differences between these procedures.



If a site contains a mixture of connected (✓) and disconnected nodes (☾ or ⚙), you must bring all offline nodes back online.

- A connected site decommission allows you to remove an operational site from the StorageGRID system. For example, you can perform a connected site decommission to remove a site that is functional but no longer needed.
- When StorageGRID removes a connected site, it uses ILM to manage the object data at the site. Before you can start a connected site decommission, you must remove the site from all ILM rules and activate a new ILM policy. The ILM processes to migrate object data and the internal processes to remove a site can occur at the same time, but the best practice is to allow the ILM steps to complete before you start the actual decommission procedure.
- A disconnected site decommission allows you to remove a failed site from the StorageGRID system. For example, you can perform a disconnected site decommission to remove a site that has been destroyed by a fire or flood.

When StorageGRID removes a disconnected site, it considers all nodes to be unrecoverable and makes no attempt to preserve data. However, before you can start a disconnected site decommission, you must remove the site from all ILM rules and activate a new ILM policy.



Before performing a disconnected site decommission procedure, you must contact your NetApp account representative. NetApp will review your requirements before enabling all steps in the Decommission Site wizard. You should not attempt a disconnected site decommission if you believe it might be possible to recover the site or to recover object data from the site.

General requirements for removing a connected or a disconnected site

Before removing a connected or disconnected site, you must be aware of the following requirements:

- You can't decommission a site that includes the primary Admin Node.
- You can't decommission a site if any of the nodes have an interface that belongs to a high availability (HA) group. You must either edit the HA group to remove the node's interface or remove the entire HA group.
- You can't decommission a site if it contains a mixture of connected (✓) and disconnected (⚙ or ☾) nodes.
- You can't decommission a site if any node at any other site is disconnected (⚙ or ☾).
- You can't start the site decommission procedure if an ec-node-repair operation is in progress. See [Check data repair jobs](#) to track repairs of erasure-coded data.
- While the site decommission procedure is running:
 - You can't create ILM rules that refer to the site being decommissioned. You also can't edit an existing ILM rule to refer to the site.

- You can't perform other maintenance procedures, such as expansion or upgrade.



If you need to perform another maintenance procedure during a connected site decommission, you can [pause the procedure while the Storage Nodes are being removed](#). The **Pause** button is enabled only when the ILM evaluation or erasure-coded data decommissioning stages are reached; however, ILM evaluation (data migration) will continue to run in the background. After the second maintenance procedure is complete, you can resume decommissioning.

- If you need to recover any node after starting the site decommission procedure, you must contact support.
- You can't decommission more than one site at a time.
- If the site includes one or more Admin Nodes and single sign-on (SSO) is enabled for your StorageGRID system, you must remove all relying party trusts for the site from Active Directory Federation Services (AD FS).

Requirements for information lifecycle management (ILM)

As part of removing a site, you must update your ILM configuration. The Decommission Site wizard guides you through a number of prerequisite steps to ensure the following:

- The site is not referred to by any ILM policy. If it is, you must edit the policies or create and activate policies with new ILM rules.
- No ILM rules refer to the site, even if those rules aren't used in any policy. You must delete or edit all rules that refer to the site.

When StorageGRID decommissions the site, it will automatically deactivate any unused erasure-coding profiles that refer to the site, and it will automatically delete any unused storage pools that refer to the site. If the All Storage Nodes storage pool exists (StorageGRID 11.6 and earlier), it is removed because it uses all sites.



Before you can remove a site, you might be required to create new ILM rules and activate a new ILM policy. These instructions assume that you have a good understanding of how ILM works and that you are familiar with creating storage pools, erasure-coding profiles, ILM rules, and simulating and activating an ILM policy. See [Manage objects with ILM](#).

Considerations for the object data at a connected site

If you are performing a connected site decommission, you must decide what to do with existing object data at the site when you create new ILM rules and a new ILM policy. You can do either or both of the following:

- Move object data from the selected site to one or more other sites in your grid.

Example for moving data: Suppose you want to decommission a site in Raleigh because you added a new site in Sunnyvale. In this example, you want to move all object data from the old site to the new site. Before updating your ILM rules and ILM policies, you must review the capacity at both sites. You must ensure that the Sunnyvale site has enough capacity to accommodate the object data from the Raleigh site and that adequate capacity will remain in Sunnyvale for future growth.



To ensure that adequate capacity is available, you might need to [expand a grid](#) by adding storage volumes or Storage Nodes to an existing site or adding a new site before you perform this procedure.

- Delete object copies from the selected site.

Example for deleting data: Suppose you currently use a 3-copy ILM rule to replicate object data across three sites. Before decommissioning a site, you can create an equivalent 2-copy ILM rule to store data at only two sites. When you activate a new ILM policy that uses the 2-copy rule, StorageGRID deletes the copies from the third site because they no longer satisfy ILM requirements. However, the object data will still be protected and the capacity of the two remaining sites will stay the same.



Never create a single-copy ILM rule to accommodate the removal of a site. An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.

Additional requirements for a connected site decommission

Before StorageGRID can remove a connected site, you must ensure the following:

- All nodes in your StorageGRID system must have a Connection State of **Connected** (✓); however, the nodes can have active alerts.



You can complete Steps 1-4 of the Decommission Site wizard if one or more nodes are disconnected. However, you can't complete Step 5 of the wizard, which starts the decommission process, unless all nodes are connected.

- If the site you plan to remove contains a Gateway Node or an Admin Node that is used for load balancing, you might need to [expand a grid](#) to add an equivalent new node at another site. Be sure clients can connect to the replacement node before starting the site decommission procedure.
- If the site you plan to remove contains any Gateway Node or Admin Nodes that are in an high availability (HA) group, you can complete Steps 1-4 of the Decommission Site wizard. However, you can't complete Step 5 of the wizard, which starts the decommission process, until you remove these nodes from all HA groups. If existing clients connect to an HA group that includes nodes from the site, you must ensure they can continue to connect to StorageGRID after the site is removed.
- If clients connect directly to Storage Nodes at the site you are planning to remove, you must ensure that they can connect to Storage Nodes at other sites before starting the site decommission procedure.
- You must provide sufficient space on the remaining sites to accommodate any object data that will be moved because of changes to any active ILM policy. In some cases, you might need to [expand a grid](#) by adding Storage Nodes, storage volumes, or new sites before you can complete a connected site decommission.
- You must allow adequate time for the decommission procedure to complete. StorageGRID ILM processes might take days, weeks, or even months to move or delete object data from the site before the site can be decommissioned.



Moving or deleting object data from a site might take days, weeks, or even months, depending on the amount of data at the site, the load on your system, network latencies, and the nature of the required ILM changes.

- Whenever possible, you should complete Steps 1-4 of the Decommission Site wizard as early as you can. The decommission procedure will complete more quickly and with fewer disruptions and performance

impacts if you allow data to be moved from the site before starting the actual decommission procedure (by selecting **Start Decommission** in Step 5 of the wizard).

Additional requirements for a disconnected site decommission

Before StorageGRID can remove a disconnected site, you must ensure the following:

- You have contacted your NetApp account representative. NetApp will review your requirements before enabling all steps in the Decommission Site wizard.



You should not attempt a disconnected site decommission if you believe it might be possible to recover the site or to recover any object data from the site. See [How technical support recovers a site](#).

- All nodes at the site must have a Connection State of one of the following:
 - **Unknown** (🌀): For an unknown reason, a node is disconnected or services on the node are unexpectedly down. For example, a service on the node might be stopped, or the node might have lost its network connection because of a power failure or unexpected outage.
 - **Administratively Down** (🌑): The node is not connected to the grid for an expected reason. For example, the node or services on the node have been gracefully shut down.
- All nodes at all other sites must have a Connection State of **Connected** (✅); however, these other nodes can have active alerts.
- You must understand that you will no longer be able to use StorageGRID to view or retrieve any object data that was stored at the site. When StorageGRID performs this procedure, it makes no attempt to preserve any data from the disconnected site.



If your ILM rules and policy were designed to protect against the loss of a single site, copies of your objects still exist on the remaining sites.

- You must understand that if the site contained the only copy of an object, the object is lost and can't be retrieved.

Considerations for consistency when you remove a site

The consistency for an S3 bucket determines whether StorageGRID fully replicates object metadata to all nodes and sites before telling a client that object ingest was successful. Consistency provides a balance between the availability of the objects and the consistency of those objects across different Storage Nodes and sites.

When StorageGRID removes a site, it needs to ensure that no data is written to the site being removed. As a result, it temporarily overrides the consistency for each bucket or container. After you start the site decommission process, StorageGRID temporarily uses strong-site consistency to prevent object metadata from being written to the site being removed.

As a result of this temporary override, be aware that any client write, update, and delete operations that occur during a site decommission can fail if multiple nodes become unavailable at the remaining sites.

Gather required materials

Before you decommission a site, you must obtain the following materials.

Item	Notes
Recovery Package .zip file	You must download the most recent Recovery Package .zip file (sgws-recovery-package-id-revision.zip). You can use the Recovery Package file to restore the system if a failure occurs. Download the Recovery Package
Passwords.txt file	This file contains the passwords required to access grid nodes on the command line and is included in the Recovery Package.
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the Passwords.txt file.
Description of StorageGRID system's topology before decommissioning	If available, obtain any documentation that describes the system's current topology.

Related information

[Web browser requirements](#)

Step 1: Select Site

To determine if a site can be decommissioned, start by accessing the Decommission Site wizard.

Before you begin

- You have obtained all required materials.
- You have reviewed the considerations for removing a site.
- You are signed in to the Grid Manager using a [supported web browser](#).
- You have the [Root access permission or the Maintenance and ILM permissions](#).

Steps

1. Select **MAINTENANCE > Tasks > Decommission**.
2. Select **Decommission Site**.

Step 1 (Select Site) of the Decommission Site wizard appears. This step includes an alphabetic list of the sites in your StorageGRID system.

Decommission Site

1

2

3

4

5

6

Select Site

View Details

Revise ILM Policy

Remove ILM References

Resolve Node Conflicts

Monitor Decommission

When you decommission a site, all nodes at the site and the site itself are permanently removed from the StorageGRID system.

Review the table for the site you want to remove. If Decommission Possible is Yes, select the site. Then, select **Next** to ensure that the site is not referred to by ILM and that all StorageGRID nodes are in the correct state.

You might not be able to remove certain sites. For example, you cannot decommission the site that contains the primary Admin Node or a site that contains an Archive Node.

Sites

	Site Name	Used Storage Capacity	Decommission Possible
☐	Raleigh	3.93 MB	
☐	Sunnyvale	3.97 MB	
	Vancouver	3.90 MB	No. This site contains the primary Admin Node.

Next

- View the values in the **Used Storage Capacity** column to determine how much storage is currently being used for object data at each site.

The Used Storage Capacity is an estimate. If nodes are offline, the Used Storage Capacity is the last known value for the site.

- For a connected site decommission, this value represents how much object data will need to be moved to other sites or deleted by ILM before you can safely decommission this site.
- For a disconnected site decommission, this value represents how much of your system's data storage will become inaccessible when you decommission this site.



If your ILM policy was designed to protect against the loss of a single site, copies of your object data should still exist on the remaining sites.

- Review the reasons in the **Decommission Possible** column to determine which sites can be decommissioned currently.



If there is more than one reason a site can't be decommissioned, the most critical reason is shown.

Decommission Possible reason	Description	Next step
Green check mark ()	You can decommission this site.	Go to the next step .
No. This site contains the primary Admin Node.	You can't decommission a site containing the primary Admin Node.	None. You can't perform this procedure.

Decommission Possible reason	Description	Next step
No. This site contains one or more Archive Nodes.	You can't decommission a site containing an Archive Node.	None. You can't perform this procedure.
No. All nodes at this site are disconnected. Contact your NetApp account representative.	You can't perform a connected site decommission unless every node in the site is connected ().	If you want to perform a disconnected site decommission, you must contact your NetApp account representative, who will review your requirements and enable the rest of the Decommission Site wizard. IMPORTANT: Never take online nodes offline so that you can remove a site. You will lose data.

The example shows a StorageGRID system with three sites. The green check mark () for the Raleigh and Sunnyvale sites indicates that you can decommission those sites. However, you can't decommission the Vancouver site because it contains the primary Admin Node.

5. If decommission is possible, select the radio button for the site.

The **Next** button is enabled.

6. Select **Next**.

Step 2 (View Details) appears.

Step 2: View Details

From Step 2 (View Details) of the Decommission Site wizard, you can review which nodes are included at the site, see how much space has been used on each Storage Node, and assess how much free space is available at the other sites in your grid.

Before you begin

Before decommissioning a site, you must review how much object data exists at the site.

- If you are performing a connected site decommission, you must understand how much object data currently exists at the site before updating ILM. Based on site capacities and your data protection needs, you can create new ILM rules to move data to other sites or to delete object data from the site.
- Perform any required Storage Node expansions before starting the decommission procedure if possible.
- If you are performing a disconnected site decommission, you must understand how much object data will become permanently inaccessible when you remove the site.




If you are performing a disconnected site decommission, ILM can't move or delete object data. Any data that remains at the site will be lost. However, if your ILM policy was designed to protect against the loss of a single site, copies of your object data still exist on the remaining sites. See [Enable site-loss protection](#).

Steps

1. From Step 2 (View Details), review any warnings related to the site you selected to remove.

Decommission Site



Data Center 2 Details

⚠ This site includes a Gateway Node. If clients are currently connecting to this node, you must configure an equivalent node at another site. Be sure clients can connect to the replacement node before starting the decommission procedure.

⚠ This site contains a mixture of connected and disconnected nodes. Before you can remove this site, you must bring all offline (blue or gray) nodes back online. Contact technical support if you need assistance.

A warning appears in these cases:

- The site includes a Gateway Node. If S3 clients are currently connecting to this node, you must configure an equivalent node at another site. Be sure clients can connect to the replacement node before continuing with the decommission procedure.
- The site contains a mixture of connected (✓) and disconnected nodes (☾ or ⚙). Before you can remove this site, you must bring all offline nodes back online.

2. Review details about the site you selected to remove.

Decommission Site



Raleigh Details

Number of Nodes: 3 Free Space: 475.38 GB
Used Space: 3.93 MB Site Capacity: 475.38 GB

Node Name	Node Type	Connection State	Details
RAL-S1-101-196	Storage Node	✓	1.30 MB used space
RAL-S2-101-197	Storage Node	✓	1.30 MB used space
RAL-S3-101-198	Storage Node	✓	1.34 MB used space

Details for Other Sites

Total Free Space for Other Sites: 950.76 GB
Total Capacity for Other Sites: 950.77 GB

Site Name	Free Space ?	Used Space ?	Site Capacity ?
Sunnyvale	475.38 GB	3.97 MB	475.38 GB
Vancouver	475.38 GB	3.90 MB	475.38 GB
Total	950.76 GB	7.87 MB	950.77 GB

[Previous](#)[Next](#)

The following information is included for the selected site:

- Number of nodes
- The total used space, free space, and capacity of all Storage Nodes in the site.
 - For a connected site decommission, the **Used Space** value represents how much object data must be moved to other sites or deleted with ILM.
 - For a disconnected site decommission, the **Used Space** value indicates how much object data will become inaccessible when you remove the site.
- Node names, types, and connection states:
 -  (Connected)
 -  (Administratively Down)
 -  (Unknown)
- Details about each node:
 - For each Storage Node, the amount of space that has been used for object data.

- For Admin Nodes and Gateway Nodes, whether the node is currently used in a high availability (HA) group. You can't decommission an Admin Node or a Gateway Node that is used in an HA group. Before you start the decommission, edit HA groups to remove all nodes at the site or remove the HA group if it only includes nodes from this site. For instructions, see [Manage high availability \(HA\) groups](#).

3. In the Details for Other Sites section of the page, assess how much space is available at the other sites in your grid.

Details for Other Sites

Total Free Space for Other Sites: 950.76 GB
Total Capacity for Other Sites: 950.77 GB

Site Name	Free Space ?	Used Space ?	Site Capacity ?
Sunnyvale	475.38 GB	3.97 MB	475.38 GB
Vancouver	475.38 GB	3.90 MB	475.38 GB
Total	950.76 GB	7.87 MB	950.77 GB

If you are performing a connected site decommission and you plan to use ILM to move object data from the selected site (instead of just deleting it), you must ensure that the other sites have enough capacity to accommodate the moved data and that adequate capacity remains for future growth.



A warning appears if the **Used Space** for the site you want to remove is greater than the **Total Free Space for Other Sites**. To ensure that adequate storage capacity is available after the site is removed, you might need to perform an expansion before performing this procedure.

4. Select **Next**.

Step 3 (Revise ILM Policy) appears.

Step 3: Revise ILM Policies

From Step 3 (Revise ILM Policies) of the Decommission Site wizard, you can determine if the site is referred to by any ILM policy.

Before you begin

You have a good understanding of how to [manage objects with ILM](#). You are familiar with creating storage pools and ILM rules and with simulating and activating an ILM policy.

About this task

StorageGRID can't decommission a site if any ILM rule in any policy (active or inactive) references that site.

If any ILM policy refers to the site you want to decommission, you must remove those policies or edit them so that they meet these requirements:

- Fully protect all object data.
- Don't refer to the site you're decommissioning.
- Don't use storage pools that refer to the site or use the All Sites option.

- Don't use erasure-coding profiles that refer to the site.
- Don't use the Make 2 Copies rule from StorageGRID 11.6 or earlier installations.



Never create a single-copy ILM rule to accommodate the removal of a site. An ILM rule that creates only one replicated copy for any time period puts data at risk of permanent loss. If only one replicated copy of an object exists, that object is lost if a Storage Node fails or has a significant error. You also temporarily lose access to the object during maintenance procedures such as upgrades.



If you're performing a *connected site decommission*, you must consider how StorageGRID should manage the object data currently at the site you want to remove. Depending on your data protection requirements, new rules can move existing object data to different sites or they can delete any extra object copies that are no longer needed.

Contact technical support if you need assistance designing a new policy.

Steps

1. From Step 3 (Revise ILM Policies), determine if any ILM policies refer to the site you selected to decommission.
2. If no policies are listed, select **Next** to go to [Step 4: Remove ILM References](#).
3. If one or more *active* ILM policies are listed, clone each existing policy or create new policies that don't reference the site being decommissioned:
 - a. Select the link for the policy in the Policy Name column.

The ILM policy details page for the policy appears in a new browser tab. The Decommission Site page will remain open on the other tab.

- b. Follow these guidelines and instructions as needed:

- Work with ILM rules:
 - [Create one or more storage pools](#) that don't refer to the site.
 - [Edit or replace rules](#) that refer to the site.



Don't select the **Make 2 Copies** rule because that rule uses the **All Storage Nodes** storage pool, which is not allowed.

- Work with ILM policies:
 - [Clone an existing ILM policy](#) or [create a new ILM policy](#).
 - Ensure the default rule and other rules don't refer to the site.



You must confirm that the ILM rules are in the correct order. When the policy is activated, new and existing objects are evaluated by the rules in the order listed, starting at the top.

- c. Ingest test objects and simulate the policy to ensure that the correct rules are applied.



Errors in an ILM policy can cause unrecoverable data loss. Carefully review and simulate the policy before activating it to confirm that it will work as intended.



When you activate a new ILM policy, StorageGRID uses it to manage all objects, including existing objects and newly ingested objects. Before activating a new ILM policy, review any changes to the placement of existing replicated and erasure-coded objects. Changing an existing object's location might result in temporary resource issues when the new placements are evaluated and implemented.

d. Activate the new policies and ensure the old policies are now inactive.

If you want to activate multiple policies, [follow the steps to create ILM policy tags](#).

If you are performing a connected site decommission, StorageGRID begins to remove object data from the selected site as soon as you activate the new ILM policy. Moving or deleting all object copies might take weeks. Although you can safely start a site decommission while object data still exists at the site, the decommission procedure will complete more quickly and with fewer disruptions and performance impacts if you allow data to be moved from the site before starting the actual decommission procedure (by selecting **Start Decommission** in Step 5 of the wizard).

4. For each *inactive* policy, edit or remove it by first selecting the link for each policy as described in the previous steps.
 - [Edit the policy](#) so it doesn't refer to the site to be decommissioned.
 - [Remove a policy](#).
5. When you finish making changes to ILM rules and policies, there should be no more policies listed in Step 3 (Revise ILM Policies). Select **Next**.

Step 4 (Remove ILM References) appears.

Step 4: Remove ILM References

From Step 4 (Remove ILM References) of the Decommission Site wizard, you must delete or edit any unused ILM rules that refer to the site, even if the rules are not used in any ILM policy.

Steps

1. Determine whether any unused ILM rules refer to the site.

If any ILM rules are listed, those rules still refer to the site but aren't used in any policy.



When StorageGRID decommissions the site, it will automatically deactivate any unused erasure-coding profiles that refer to the site, and it will automatically delete any unused storage pools that refer to the site. The All Storage Nodes storage pool (StorageGRID 11.6 and earlier) is removed because it uses the All Sites site.

2. Edit or delete each unused rule:
 - To edit a rule, go the ILM rules page and update all placements that use an erasure-coding profile or storage pool that refers to the site. Then, return to **Step 4 (Remove ILM References)**.
 - To delete a rule, select the trash can icon  and select **OK**.



You must delete the **Make 2 Copies** rule before you can decommission a site.

3. Confirm that no unused ILM rules refer to the site, and the **Next** button is enabled.

4. Select **Next**.



Any remaining storage pools and erasure-coding profiles that refer to the site will become invalid when the site is removed. When StorageGRID decommissions the site, it will automatically deactivate any unused erasure-coding profiles that refer to the site, and it will automatically delete any unused storage pools that refer to the site. The All Storage Nodes storage pool (StorageGRID 11.6 and earlier) is removed because it uses the All Sites site.

Step 5 (Resolve Node Conflicts) appears.

Step 5: Resolve Node Conflicts (and start decommission)

From Step 5 (Resolve Node Conflicts) of the Decommission Site wizard, you can determine if any nodes in your StorageGRID system are disconnected or if any nodes at the selected site belong to a high availability (HA) group. After any node conflicts are resolved, you start the decommission procedure from this page.

Before you begin

You must ensure that all nodes in your StorageGRID system are in the correct state, as follows:

- All nodes in your StorageGRID system must be connected (✓).



If you are performing a disconnected site decommission, all nodes at the site you are removing must be disconnected, and all nodes at all other sites must be connected.



The decommission will not start if one or more volumes are offline (unmounted) or if they are online (mounted) but in an error state.



If one or more volumes go offline while a decommission is in progress, the decommission process completes after these volumes have come back online.

- No node at the site you are removing can have an interface that belongs to a high availability (HA) group.

About this task

If any node is listed for Step 5 (Resolve Node Conflicts), you must correct the issue before you can start the decommission.

Before starting the site decommission procedure from this page, review the following considerations:

- You must allow adequate time for the decommission procedure to complete.



Moving or deleting object data from a site might take days, weeks, or even months, depending on the amount of data at the site, the load on your system, network latencies, and the nature of the required ILM changes.

- While the site decommission procedure is running:
 - You can't create ILM rules that refer to the site being decommissioned. You also can't edit an existing ILM rule to refer to the site.

- You can't perform other maintenance procedures, such as expansion or upgrade.




If you need to perform another maintenance procedure during a connected site decommission, you can pause the procedure while the Storage Nodes are being removed. The **Pause** button is enabled during the "Decommissioning Replicated and Erasure-Coded Data" stage.

- If you need to recover any node after starting the site decommission procedure, you must contact support.

Steps

1. Review the disconnected nodes section of Step 5 (Resolve Node Conflicts) to determine if any nodes in your StorageGRID system have a Connection State of Unknown (🔗) or Administratively Down (🌙).

Decommission Site



Before you can decommission the site, you must ensure the following:

- All nodes in your StorageGRID system are connected.
Note: If you are performing a disconnected site decommission, all nodes at the site you are removing must be disconnected.
- No node at the selected site belongs to a high availability (HA) group.

If a node is listed in either table, you must correct the issue before you can continue.

1 disconnected node in the grid

The following nodes have a Connection State of Unknown (blue) or Administratively Down (gray). You must bring these disconnected nodes back online.

For help bringing nodes back online, see the instructions for [monitoring and troubleshooting StorageGRID](#) and the [recovery and maintenance](#) instructions.

Node Name	Connection State	Site	Type
DC1-S3-99-193 🔗	🌙 Administratively Down	Data Center 1	Storage Node

1 node in the selected site belongs to an HA group

Passphrase

Provisioning Passphrase ?

Previous

Start Decommission

2. If any nodes are disconnected, bring them back online.

See the [Node procedures](#). Contact technical support if you need assistance.

- When all disconnected nodes have been brought back online, review the HA groups section of Step 5 (Resolve Node Conflicts).

This table lists any nodes at the selected site that belong to a high availability (HA) group.

Decommission Site



Before you can decommission the site, you must ensure the following:

- All nodes in your StorageGRID system are connected.
Note: If you are performing a disconnected site decommission, all nodes at the site you are removing must be disconnected.
- No node at the selected site belongs to a high availability (HA) group.

If a node is listed in either table, you must correct the issue before you can continue.

All grid nodes are connected

1 node in the selected site belongs to an HA group

The following nodes in the selected site belong to a high availability (HA) group. You must either edit the HA group to remove the node's interface or remove the entire HA group.

[Go to HA Groups page.](#)

For information about HA groups, see the instructions for [administering StorageGRID](#)

HA Group Name	Node Name	Node Type
HA group	DC1-GW1-99-190	API Gateway Node

Passphrase

Provisioning Passphrase

Previous

Start Decommission

- If any nodes are listed, do either of the following:

- Edit each affected HA group to remove the node interface.
- Remove an HA group that only includes nodes from this site.
See the instructions for administering StorageGRID.

If all nodes are connected and no nodes in the selected site are used in an HA group, the **Provisioning Passphrase** field is enabled.

- Enter the provisioning passphrase.

The **Start Decommission** button becomes enabled.

Decommission Site



Before you can decommission the site, you must ensure the following:

- All nodes in your StorageGRID system are connected.
Note: If you are performing a disconnected site decommission, all nodes at the site you are removing must be offline.
- No node at the selected site belongs to a high availability (HA) group.

If a node is listed in either table, you must correct the issue before you can continue.

All grid nodes are connected

No nodes in the selected site belong to an HA group

Passphrase

Provisioning Passphrase ?

.....

Previous

Start Decommission

6. If you are ready to start the site decommission procedure, select **Start Decommission**.

A warning lists the site and nodes that will be removed. You are reminded that it might take days, weeks, or even months to completely remove the site.

Warning

The following site and its nodes have been selected for decommissioning and will be permanently removed from the StorageGRID system:

Data Center 3

- DC3-S1
- DC3-S2
- DC3-S3

When StorageGRID removes a site, it temporarily uses strong-site consistency to prevent object metadata from being written to the site being removed. Client write and delete operations can fail if multiple nodes become unavailable at the remaining sites.

This procedure might take days, weeks, or even months to complete. Select **Maintenance > Decommission** to monitor the decommission progress.

Do you want to continue?

Cancel

OK

7. Review the warning. If you are ready to begin, select **OK**.

A message appears as the new grid configuration is generated. This process might take some time, depending on the type and number of decommissioned grid nodes.

Passphrase

Provisioning Passphrase 

 Generating grid configuration. This may take some time depending on the type and the number of decommissioned grid nodes.

Previous

Start Decommission 

When the new grid configuration has been generated, Step 6 (Monitor Decommission) appears.



The **Previous** button remains disabled until the decommission is complete.

Step 6: Monitor Decommission

From Step 6 (Monitor Decommission) of the Decommission Site page wizard, you can monitor the progress as the site is removed.

About this task

When StorageGRID removes a connected site, it removes nodes in this order:

1. Gateway Nodes

2. Admin Nodes
3. Storage Nodes

When StorageGRID removes a disconnected site, it removes nodes in this order:

1. Gateway Nodes
2. Storage Nodes
3. Admin Nodes

Each Gateway Node or Admin Node might only require a few minutes or an hour to remove; however, Storage Nodes might take days or weeks.

Steps

1. As soon as a new Recovery Package has been generated, download the file.

Decommission Site



i A new Recovery Package has been generated as a result of the configuration change. Go to the [Recovery Package](#) page to download it.



Download the Recovery Package as soon as possible to ensure you can recover your grid if something goes wrong during the decommission procedure.

- a. Select the link in the message, or select **MAINTENANCE > System > Recovery package**.
- b. Download the .zip file.

See the instructions for [downloading the Recovery Package](#).




The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

2. Using the Data Movement chart, monitor the movement of object data from this site to other sites.

Data movement started when you activated the new ILM policy in Step 3 (Revise ILM Policy). Data movement will occur throughout the decommission procedure.

Decommission Site Progress



3. In the Node Progress section of the page, monitor the progress of the decommission procedure as nodes are removed.

When a Storage Node is removed, each node goes through a series of stages. Although most of these stages occur quickly or even imperceptibly, you might need to wait days or even weeks for other stages to complete, based on how much data needs to be moved. Additional time is required to manage erasure-coded data and re-evaluate ILM.

Node Progress

 Depending on the number of objects stored, Storage Nodes might take significantly longer to decommission. Extra time is needed to manage erasure coded data and re-evaluate ILM.

The progress for each node is displayed while the decommission procedure is running. If you need to perform another maintenance procedure, select **Pause** to suspend the decommission (only allowed during certain stages).

Pause Resume

Search 

Name 	Type 	Progress 	Stage 
RAL-S1-101-196	Storage Node		Decommissioning Replicated and Erasure Coded Data
RAL-S2-101-197	Storage Node		Decommissioning Replicated and Erasure Coded Data
RAL-S3-101-198	Storage Node		Decommissioning Replicated and Erasure Coded Data

If you are monitoring the progress of a connected site decommission, refer to this table to understand the decommission stages for a Storage Node:

Stage	Estimated duration
Pending	Minute or less
Wait for Locks	Minutes
Prepare Task	Minute or less
Marking LDR Decommissioned	Minutes
Decommissioning Replicated and Erasure-Coded Data	Hours, days, or weeks based on the amount of data Note: If you need to perform other maintenance activities, you can pause the site decommission during this stage.
LDR Set State	Minutes
Flush Audit Queues	Minutes to hours, based on the number of messages and network latency.
Complete	Minutes

If you are monitoring the progress of a disconnected site decommission, refer to this table to understand the decommission stages for a Storage Node:

Stage	Estimated duration
Pending	Minute or less
Wait for Locks	Minutes
Prepare Task	Minute or less
Disable External Services	Minutes
Certificate Revocation	Minutes
Node Unregister	Minutes
Storage Grade Unregister	Minutes
Storage Group Removal	Minutes
Entity Removal	Minutes

Stage	Estimated duration
Complete	Minutes

4. After all nodes have reached the Complete stage, wait for the remaining site decommission operations to complete.
 - During the **Repair Cassandra** step, StorageGRID makes any necessary repairs to the Cassandra clusters that remain in your grid. These repairs might take several days or more, depending on how many Storage Nodes remain in your grid.

Decommission Site Progress

Decommission Nodes in Site	Completed
Repair Cassandra	In Progress 
StorageGRID is repairing the remaining Cassandra clusters after removing the site. This might take several days or more, depending on how many Storage Nodes remain in your grid.	
Overall Progress	0%
Deactivate EC Profiles & Delete Storage Pools	Pending
Remove Configurations	Pending

- During the **Deactivate EC Profiles & Delete Storage Pools** step, the following ILM changes are made:
 - Any erasure-coding profiles that referred to the site are deactivated.
 - Any Storage Pools that referred to the site are deleted.



The All Storage Nodes storage pool (StorageGRID 11.6 and earlier) is also removed because it uses the All Sites site.


- Finally, during the **Remove Configuration** step, any remaining references to the site and its nodes are removed from the rest of the grid.

Decommission Site Progress

Decommission Nodes in Site	Completed
Repair Cassandra	Completed
Deactivate EC Profiles & Delete Storage Pools	Completed
Remove Configurations	In Progress 
StorageGRID is removing the site and node configurations from the rest of the grid.	

5. When the decommission procedure has completed, the Decommission Site page shows a success message, and the removed site is no longer shown.

Decommission Site



The previous decommission procedure completed successfully at 2021-01-12 14:28:32 MST.

When you decommission a site, all nodes at the site and the site itself are permanently removed from the StorageGRID system.

Review the table for the site you want to remove. If Decommission Possible is Yes, select the site. Then, select **Next** to ensure that the site is not referred to by ILM and that all StorageGRID nodes are in the correct state.

You might not be able to remove certain sites. For example, you cannot decommission the site that contains the primary Admin Node or a site that contains an Archive Node.

Sites

	Site Name	Used Storage Capacity ?	Decommission Possible
☒	Sunnyvale	4.79 MB	✓
☐	Vancouver	4.90 MB	No. This site contains the primary Admin Node.

Next

After you finish

Complete these tasks after you complete the site decommission procedure:

- Ensure that the drives of all Storage Nodes in the decommissioned site are wiped clean. Use a commercially available data wiping tool or service to permanently and securely remove data from the drives.
- If the site included one or more Admin Nodes and single sign-on (SSO) is enabled for your StorageGRID system, remove all relying party trusts for the site from Active Directory Federation Services (AD FS).
- After the nodes have been gracefully powered off automatically as part of the connected site decommission procedure, remove the associated virtual machines.

Rename grid, site, or node

Use the rename procedure

As required, you can change the display names that are shown throughout the Grid Manager for the entire grid, each site, and each node. You can update display names safely and whenever you need.

What is the rename procedure?

When you install StorageGRID initially, you specify a name for the grid, each site, and each node. These initial

names are known as *system names*, and they are the names initially shown throughout StorageGRID.

System names are required for internal StorageGRID operations and can't be changed. However, you can use the rename procedure to define new *display names* for the grid, each site, and each node. These display names appear in various StorageGRID locations instead of (or in some cases, in addition to) the underlying system names.

Use the rename procedure to correct typos, to implement a different naming convention, or to indicate that a site and all of its nodes have been relocated. Unlike system names, display names can be updated whenever required and without impacting StorageGRID operations.

Where do system and display names appear?

The following table summarizes where system names and display names are shown in the StorageGRID user interface and in StorageGRID files.

Location	System name	Display name
Grid Manager pages	Shown unless the item is renamed	If an item is renamed, shown instead of the system name in these locations: • Dashboard • Nodes page • Configuration pages for high availability groups, load balancer endpoints, VLAN interfaces, key management servers, grid passwords, and firewall control • Alerts • Storage pool definitions • Object metadata lookup page • Pages related to maintenance procedures, including upgrade, hotfix, SANtricity OS upgrade, decommission, expansion, recovery, and object existence check • Support pages (logs and diagnostics) • Single sign-on page, next to the Admin Node hostname in the table for Admin Node details
NODES > Overview tab for a node	Always shown	Shown only if the item is renamed
Legacy pages in the Grid Manager (for example, SUPPORT > Grid Topology)	Shown	Not shown
node-health API	Always returned	Returned only if the item is renamed

Location	System name	Display name
Prompt when using SSH to access a node	Shown as the primary name unless the item has been renamed: <pre>admin@SYSTEM-NAME: ~ \$</pre> Included in parentheses when the item is renamed: <pre>admin@DISPLAY-NAME (SYSTEM-NAME) :~ \$</pre>	Shown as the primary name when the item is renamed: <pre>admin@DISPLAY-NAME (SYSTEM-NAME) :~ \$</pre>
Passwords.txt file in the Recovery Package	Shown as Server Name	Shown as Display Name
/etc/hosts file on all nodes For example: <pre>10.96.99.128 SYSTEM-NAME 28989c59-a2c3-4d30-bb09-6879adf2437f DISPLAY-NAME localhost-grid # storagegrid-gen-host</pre>	Always shown in the second column	When the item is renamed, shown in the fourth column
topology-display-names.json, included with AutoSupport data	Not included	Empty unless items have been renamed; otherwise, maps grid, site, and node IDs to their display names.

Display name requirements

Before using this procedure, review the requirements for display names.

Display names for nodes

Display names for nodes must follow these rules:

- Must be unique across your StorageGRID system.
- Can't be the same as the system name for any other item in your StorageGRID system.
- Must contain at least 1 and no more than 32 characters.
- Can contain numbers, hyphens (-), and uppercase and lowercase letters.
- Can start or end with a letter or number, but can't start or end with a hyphen.
- Can't be all numbers.
- Are case-insensitive. For example, DC1-ADM and dc1-adm are considered to be duplicates.

You can rename a node with a display name that was previously used by a different node, as long as the

rename doesn't result in a duplicate display name or system name.

Display names for grid and sites

Display names for the grid and sites follow the same rules with these exceptions:

- Can include spaces.
- Can include these special characters: = - _ : , . @ !
- Can start and end with the special characters, including hyphens.
- Can be all numbers or special characters.

Display name best practices

If you plan to rename multiple items, document your general naming scheme before using this procedure. Come up with a system that ensures that names are unique, consistent, and easy to understand at a glance.

You can use any naming convention that fits your organizational requirements. Consider these basic suggestions of what to include:

- **Site indicator:** If you have multiple sites, add a site code to each node name.
- **Node type:** Node names typically indicate the node's type. You can use abbreviations like `s`, `adm`, and `gw` (Storage Node, Admin Node, and Gateway Node).
- **Node number:** If a site contains more than one of a particular type of node, add a unique number to each node's name.

Think twice before adding specific details to the names that are likely to change over time. For example, don't include IP addresses in node names because these addresses can be changed. Similarly, rack locations or appliance model numbers can change if you move equipment or upgrade the hardware.

Example display names

Suppose your StorageGRID system has three data centers and has nodes of different types at each data center. Your display names might be as simple as these:

- **Grid:** StorageGRID Deployment
- **First site:** Data Center 1
 - `dc1-adm1`
 - `dc1-s1`
 - `dc1-s2`
 - `dc1-s3`
 - `dc1-gw1`
- **Second site:** Data Center 2
 - `dc2-adm2`
 - `dc2-s1`
 - `dc2-s2`

- dc2-s3

- **Third site:** Data Center 3

- dc3-s1

- dc3-s2

- dc3-s3

Add or update display names

You can use this procedure to add or update the display names used for your grid, sites, and nodes. You can rename a single item, multiple items, or even all items at the same time. Defining or updating a display name does not affect StorageGRID operations in any way.

Before you begin

- From the **primary Admin Node**, you are signed in to the Grid Manager using a [supported web browser](#).



You can add or update display names from a non-primary Admin Node, but you must be signed in to the primary Admin Node to download a Recovery Package.

- You have the [Maintenance or Root access permission](#).
- You have the provisioning passphrase.
- You understand the requirements and best practices for display names. See [Rename grid, sites, and nodes](#).

How to rename grid, sites, or nodes

You can rename your StorageGRID system, one or more sites, or one or more nodes.

You can use a display name that was previously used by a different node, as long as the rename doesn't result in a duplicate display name or system name.

Select items to rename

To start, select the items you want to rename.

Steps

1. Select **MAINTENANCE > Tasks > Rename grid, sites, and nodes**.
2. For the **Select names** step, select the items you want to rename.

Item to change	Instruction
Names of everything (or almost everything) in your system	1. Select Select all. 2. Optionally clear any items you don't want to rename.
Name of the grid	Select the checkbox for the grid.

Item to change	Instruction
Name of a site and some or all of its nodes	1. Select the checkbox in the table header for the site. 2. Optionally, clear any nodes you don't want to rename.
Name of a site	Select the checkbox for the site.
Name of a node	Select the checkbox for the node.

3. Select **Continue**.

4. Review the table, which includes the items you selected.

- The **Display name** column shows the current name for each item. If the item has never been renamed, its display name is the same as its system name.
- The **System name** column shows the name you entered for each item during installation. System names are used for internal StorageGRID operations and can't be changed. For example, the system name for a node might be its hostname.
- The **Type** column indicates the item's type: Grid, Site, or the specific type of node.

Propose new names

For the **Propose new names** step, you can enter a display name for each item individually, or you can rename items in bulk.

Rename items individually

Follow these steps to enter a display name for each item you want to rename.

Steps

1. In the **Display name** field, enter a proposed display name for each item in the list.

See [Rename grid, sites, and nodes](#) to learn the naming requirements.

2. To remove any items you don't want to rename, select  in the **Remove from list** column.

If you will not be proposing a new name for an item, you must remove it from the table.

3. When you have proposed new names for all items in the table, select **Rename**.

A success message appears. The new display names are now used throughout Grid Manager.

Rename items in bulk

Use the bulk rename tool if item names share a common string that you want to replace with a different string.

Steps

1. For the **Propose new names** step, select **Use bulk rename tool**.

The **Rename preview** includes all items that were shown for the **Propose new names** step. You can use the preview to see how display names will look after you replace a shared string.

2. In the **Existing string** field, enter the shared string you want to replace. For example, if the string you want to replace is `Data-Center-1`, enter **Data-Center-1**.

As you type, your text is highlighted wherever it is found in the names on the left.

3. Select  to remove any items that you don't want to rename with this tool.

For example, suppose you want to rename all nodes that contain the string `Data-Center-1`, but you don't want to rename the `Data-Center-1` site itself. Select  to remove the site from the rename preview.

Bulk rename tool

Rename preview ?

Data-Center-1 X

Data-Center-1-ADM1 X

Data-Center-1-ARC1 X

Data-Center-1-G1 X

Data-Center-1-S1 X

Data-Center-1-S2 X

Data-Center-1-S3 X

Data-Center-1-S4 ✓

Cancel

Add names

Enter the shared string you want to replace. Then, enter a new string to use instead. Optionally, remove any items that you do not want to rename with this tool.

Existing string

Data-Center-1

The string you want to replace. Represented by *italicized text* in the preview section.

New string

The string you want to use instead. Represented by **bolded text** in the preview section.

4. In the **New string** field, enter the replacement string you want to use instead. For example, enter **DC1**.

See [Rename grid, sites, and nodes](#) to learn the naming requirements.

As you enter the replacement string, the names on the left are updated, so you can verify that the new names will be correct.

Bulk rename tool

×

Rename preview ?

DC1-ADM1 ✕

DC1-ARC1 ✕

DC1-G1 ✕

DC1-S1 ✕

DC1-S2 ✕

DC1-S3 ✕

DC1-S4 ✕

Cancel

Add names

Enter the shared string you want to replace. Then, enter a new string to use instead. Optionally, remove any items that you do not want to rename with this tool.

Existing string

Data-Center-1

The string you want to replace. Represented by *italicized text* in the preview section.

New string

DC1

The string you want to use instead. Represented by **bolded text** in the preview section.

- When you are satisfied with the names shown in the preview, select **Add names** to add the names to the table for the **Propose new names** step.
- Make any additional changes required, or select ✕ to remove any items that you don't want to rename.
- When you are ready to rename all items in the table, select **Rename**.

A success message is shown. The new display names are now used throughout Grid Manager.

Download the Recovery Package

When you are done renaming items, download and save a new Recovery Package. The new display names for the items you renamed are included in the `Passwords.txt` file.

Steps

- Enter the provisioning passphrase.
- Select **Download Recovery Package**.

The download starts immediately.

- When the download completes, open the `Passwords.txt` file to see the server name for all nodes and the display names for any renamed nodes.
- Copy the `sgws-recovery-package-id-revision.zip` file to two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

5. Select **Finish** to return to the first step.

Revert display names back to system names

You can revert a renamed grid, site, or node back to its original system name. When you revert an item back to its system name, Grid Manager pages and other StorageGRID locations no longer show a **Display name** for that item. Only the item's system name is shown.

Steps

1. Select **MAINTENANCE > Tasks > Rename grid, sites, and nodes**.
2. For the **Select names** step, select any items you want to revert back to system names.
3. Select **Continue**.
4. For the **Propose new names** step, revert display names back to system names individually or in bulk.

Revert to system names individually

- a. Copy each item's original system name and paste it into the **Display name** field, or select  to remove any items you don't want to revert.

To revert a display name, the system name must appear in the **Display name** field, but the name is case insensitive.

- b. Select **Rename**.

A success message appears. The display names for these items are no longer used.

Revert to system names in bulk

- a. For the **Propose new names** step, select **Use bulk rename tool**.
- b. In the **Existing string** field, enter the display name string you want to replace.
- c. In the **New string** field, enter the system name string you want to use instead.
- d. Select **Add names** to add the names to the table for the **Propose new names** step.
- e. Confirm that each entry in the **Display name** field matches the name in the **System name** field. Make any changes or select  to remove any items that you don't want to revert.

To revert a display name, the system name must appear in the **Display name** field, but the name is case insensitive.

- f. Select **Rename**.

A success message is shown. The display names for these items are no longer used.

5. [Download and save a new Recovery Package](#).

Display names for the items you reverted are no longer included in the `Passwords.txt` file.

Node procedures

Node maintenance procedures

You might need to perform maintenance procedures related to specific grid nodes or node services.

Server Manager procedures

Server Manager runs on every grid node to supervise the starting and stopping of services and to ensure that services gracefully join and leave the StorageGRID system. Server Manager also monitors the services on every grid node and will automatically attempt to restart any services that report faults.

To perform Server Manager procedures, you typically need to access the node's command line.



You should access Server Manager only if technical support has directed you to do so.



You must close the current command shell session and log out after you are finished with Server Manager. Enter: `exit`

Node reboot, shut down, and power procedures

You use these procedures to reboot one or more nodes, to shut down and restart nodes, or to power nodes off and power them back on.

Port remap procedures

You can use the port remap procedures to remove the port remaps from a node, for example, if you want to configure a load balancer endpoint using a port that was previously remapped.

Server Manager procedures

View Server Manager status and version

For each grid node, you can view the current status and version of Server Manager running on that grid node. You can also obtain the current status of all services running on that grid node.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

2. View the current status of Server Manager running on the grid node: **service servermanager status**

The current status of Server Manager running on the grid node is reported (running or not). If Server Manager's status is running, the time it has been running since last it was started is listed. For example:

```
servermanager running for 1d, 13h, 0m, 30s
```

3. View the current version of Server Manager running on a grid node: **service servermanager version**

The current version is listed. For example:

```
11.1.1.0-20180425.1905.39c9493
```

4. Log out of the command shell: **exit**

View current status of all services

You can view the current status of all services running on a grid node at any time.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

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

For example, the output for the primary Admin Node shows the current status of the AMS, CMN, and NMS services as Running. This output is updated immediately if the status of a service changes.

Host Name	190-ADM1	
IP Address		
Operating System Kernel	4.9.0	Verified
Operating System Environment	Debian 9.4	Verified
StorageGRID Webscale Release	11.1.0	Verified
Networking		Verified
Storage Subsystem		Verified
Database Engine	5.5.9999+default	Running
Network Monitoring	11.1.0	Running
Time Synchronization	1:4.2.8p10+dfsg	Running
ams	11.1.0	Running
cmn	11.1.0	Running
nms	11.1.0	Running
ssm	11.1.0	Running
mi	11.1.0	Running
dynip	11.1.0	Running
nginx	1.10.3	Running
tomcat	8.5.14	Running
grafana	4.2.0	Running
mgmt api	11.1.0	Running
prometheus	1.5.2+ds	Running
persistence	11.1.0	Running
ade exporter	11.1.0	Running
attrDownPurge	11.1.0	Running
attrDownSampl	11.1.0	Running
attrDownSamp2	11.1.0	Running
node exporter	0.13.0+ds	Running

- Return to the command line, press **Ctrl+C**.
- Optionally, view a static report for all services running on the grid node:
`/usr/local/servermanager/reader.rb`

This report includes the same information as the continuously updated report, but it is not updated if the status of a service changes.

- Log out of the command shell: `exit`

Start Server Manager and all services

You might need to start Server Manager, which also starts all services on the grid node.

Before you begin

You have the `Passwords.txt` file.

About this task

Starting Server Manager on a grid node where it is already running results in a restart of Server Manager and all services on the grid node.

Steps

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

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

2. Start Server Manager: `service servermanager start`
3. Log out of the command shell: `exit`

Restart Server Manager and all services

You might need to restart server manager and all services running on a grid node.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
 - a. Enter the following command: `ssh admin@grid_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 #.

2. Restart Server Manager and all services on the grid node: `service servermanager restart`

Server Manager and all services on the grid node are stopped and then restarted.



Using the `restart` command is the same as using the `stop` command followed by the `start` command.

3. Log out of the command shell: `exit`

Stop Server Manager and all services

Server Manager is intended to run at all times, but you might need to stop Server Manager and all services running on a grid node.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
 - a. Enter the following command: `ssh admin@grid_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 #.

2. Stop Server manager and all services running on the grid node: `service servermanager stop`

Server Manager and all services running on the grid node are gracefully terminated. Services can take up to 15 minutes to shut down.

3. Log out of the command shell: `exit`

View current status of service

You can view the current status of a services running on a grid node at any time.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

2. View the current status of a service running on a grid node: **`service servicename status`**
The current status of the requested service running on the grid node is reported (running or not). For example:

```
cmn running for 1d, 14h, 21m, 2s
```

3. Log out of the command shell: **`exit`**

Stop service

Some maintenance procedures require you to stop a single service while keeping other services on the grid node running. Only stop individual services when directed to do so by a maintenance procedure.

Before you begin

You have the `Passwords.txt` file.

About this task

When you use these steps to "administratively stop" a service, Server Manager will not automatically restart the service. You must either start the single service manually or restart Server Manager.

If you need to stop the LDR service on a Storage Node, be aware that it might take a while to stop the service if there are active connections.

Steps

1. Log in to the grid node:

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

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

2. Stop an individual service: `service servicename stop`

For example:

```
service ldr stop
```



Services can take up to 11 minutes to stop.

3. Log out of the command shell: `exit`

Related information

[Force service to terminate](#)

Force service to terminate

If you need to stop a service immediately, you can use the `force-stop` command.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:

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

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

2. Manually force the service to terminate: `service servicename force-stop`

For example:

```
service ldr force-stop
```

The system waits 30 seconds before terminating the service.

3. Log out of the command shell: `exit`

Start or restart service

You might need to start a service that has been stopped, or you might need to stop and restart a service.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

2. Decide which command to issue, based on whether the service is currently running or stopped.
 - If the service is currently stopped, use the `start` command to start the service manually: `service servicename start`

For example:

```
service ldr start
```

- If the service is currently running, use the `restart` command to stop the service and then restart it: `service servicename restart`

For example:

```
service ldr restart
```



Using the `restart` command is the same as using the `stop` command followed by the `start` command. You can issue `restart` even if the service is currently stopped.

3. Log out of the command shell: `exit`

Use a DoNotStart file

If you are performing various maintenance or configuration procedures under the direction of technical support, you might be asked to use a `DoNotStart` file to prevent services from starting when Server Manager is started or restarted.



You should add or remove a DoNotStart file only if technical support has directed you to do so.

To prevent a service from starting, place a DoNotStart file in the directory of the service you want to prevent from starting. At start-up, Server Manager looks for the DoNotStart file. If the file is present, the service (and any services dependent on it) is prevented from starting. When the DoNotStart file is removed, the previously stopped service will start on the next start or restart of Server Manager. Services aren't automatically started when the DoNotStart file is removed.

The most efficient way to prevent all services from restarting is to prevent the NTP service from starting. All services are dependent on the NTP service and can't run if the NTP service is not running.

Add DoNotStart file for service

You can prevent an individual service from starting by adding a DoNotStart file to that service's directory on a grid node.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

2. Add a DoNotStart file: `touch /etc/sv/service/DoNotStart`

where `service` is the name of the service to be prevented from starting. For example,

```
touch /etc/sv/ldr/DoNotStart
```

A DoNotStart file is created. No file content is needed.

When Server Manager or the grid node is restarted, Server Manager restarts, but the service does not.

3. Log out of the command shell: `exit`

Remove DoNotStart file for service

When you remove a DoNotStart file that is preventing a service from starting, you must start that service.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:

- a. Enter the following command: `ssh admin@grid_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 `#`.

2. Remove the `DoNotStart` file from the service directory: `rm /etc/sv/service/DoNotStart`

where `service` is the name of the service. For example,

```
rm /etc/sv/ldr/DoNotStart
```

3. Start the service: `service servicename start`

4. Log out of the command shell: `exit`

Troubleshoot Server Manager

If a problem arises when using Server Manager, check its log file.

Error messages related to Server Manager are captured in the Server Manager log file, which is located at: `/var/local/log/servermanager.log`

Check this file for error messages regarding failures. Escalate the issue to technical support if required. You might be asked to forward log files to technical support.

Service with an error state

If you detect that a service has entered an error state, attempt to restart the service.

Before you begin

You have the `Passwords.txt` file.

About this task

Server Manager monitors services and restarts any that have stopped unexpectedly. If a service fails, Server Manager attempts to restart it. If there are three failed attempts to start a service within five minutes, the service enters an error state. Server Manager does not attempt another restart.

Steps

1. Log in to the grid node:

- a. Enter the following command: `ssh admin@grid_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 #.

2. Confirm the error state of the service: `service servicename status`

For example:

```
service ldr status
```

If the service is in an error state, the following message is returned: `servicename in error state`.
For example:

```
ldr in error state
```



If the service status is disabled, see the instructions for [removing a DoNotStart file for a service](#).

3. Attempt to remove the error state by restarting the service: `service servicename restart`

If the service fails to restart, contact technical support.

4. Log out of the command shell: `exit`

Reboot, shutdown, and power procedures

Perform a rolling reboot

You can perform a rolling reboot to reboot multiple grid nodes without causing a service disruption.

Before you begin

- You are signed in to the Grid Manager on the primary Admin Node, and you are using a [supported web browser](#).



You must be signed in to the primary Admin Node to perform this procedure.

- You have the [Maintenance or Root access permission](#).

About this task

Use this procedure if you need to reboot multiple nodes at the same time. For example, you can use this procedure after changing the FIPS mode for the grid's [TLS and SSH security policy](#). When the FIPS mode changes, you must reboot all nodes to put the change into effect.



If you only need to reboot one node, you can [reboot the node from the Tasks tab](#).

When StorageGRID reboots grid nodes, it issues the `reboot` command on each node, which causes the node to shut down and restart. All services are restarted automatically.

- Rebooting a VMware node reboots the virtual machine.

- Rebooting a Linux node reboots the container.
- Rebooting a StorageGRID Appliance node reboots the compute controller.

The rolling reboot procedure can reboot multiple nodes at the same time, with these exceptions:

- Two nodes of the same type will not be rebooted at the same time.
- Gateway Nodes and Admin Nodes will not be rebooted at the same time.

Instead, these nodes are rebooted sequentially to ensure that HA groups, object data, and critical node services always remain available.

When you reboot the primary Admin Node, your browser temporarily loses access to the Grid Manager, so you can no longer monitor the procedure. For this reason, the primary Admin Node is rebooted last.

Perform a rolling reboot

You select the nodes you want to reboot, review your selections, start the reboot procedure, and monitor the progress.

Select nodes

As your first step, access the Rolling reboot page and select the nodes you want to reboot.

Steps

1. Select **MAINTENANCE > Tasks > Rolling reboot**.
2. Review the connection state and alert icons in the **Node name** column.



You can't reboot a node if it is disconnected from the grid. The checkboxes are disabled for nodes with these icons:  or .

3. If any nodes have active alerts, review the list of alerts in the **Alert summary** column.



To see all current alerts for a node, you can also select the [Nodes > Overview tab](#).

4. Optionally, perform the recommended actions to resolve any current alerts.
5. Optionally, if all nodes are connected and you want to reboot all of them, select the checkbox in the table header and select **Select all**. Otherwise, select each node you want to reboot.

You can use the table's filter options to view subsets of nodes. For example, you can view and select only Storage Nodes or all nodes at a certain site.

6. Select **Review selection**.

Review selection

In this step, you can determine how long the total reboot procedure might take and confirm that you selected the correct nodes.

1. On the Review selection page, review the Summary, which indicates how many nodes will be rebooted and the estimated total time for all nodes to reboot.
2. Optionally, to remove a specific node from the reboot list, select **Remove**.

3. Optionally, to add more nodes, select **Previous step**, select the additional nodes, and select **Review selection**.
4. When you are ready to start the rolling reboot procedure for all selected nodes, select **Reboot nodes**.
5. If you selected to reboot the primary Admin Node, read the information message, and select **Yes**.



The primary Admin Node will be the last node to reboot. While this node is rebooting, your browser's connection will be lost. When the primary Admin Node is available again, you must reload the Rolling reboot page.

Monitor a rolling reboot

While the rolling reboot procedure is running, you can monitor it from the primary Admin Node.

Steps

1. Review the overall progress of the operation, which includes the following information:
 - Number of nodes rebooted
 - Number of nodes in process of being rebooted
 - Number of nodes that remain to be rebooted
2. Review the table for each type of node.

The tables provide a progress bar of the operation on each node and show the reboot stage for that node, which can be one of these:

- Waiting to Reboot
- Stopping services
- Rebooting system
- Starting services
- Reboot completed

Stop the rolling reboot procedure

You can stop the rolling reboot procedure from the primary Admin Node. When you stop the procedure, any nodes that have the status "Stopping services," "Rebooting system," or "Starting services" will complete the reboot operation. However, these nodes will no longer be tracked as part of the procedure.

Steps

1. Select **MAINTENANCE > Tasks > Rolling reboot**.
2. From the **Monitor reboot** step, select **Stop reboot procedure**.

Reboot grid node from Tasks tab

You can reboot an individual grid node from the Tasks tab on the Nodes page.

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 the provisioning passphrase.
- If you are rebooting the primary Admin Node or any Storage Node, you have reviewed the following considerations:
 - When you reboot the primary Admin Node, your browser temporarily loses access to the Grid Manager.
 - If you reboot two or more Storage Nodes at a given site, you might not be able to access certain objects for the duration of the reboot. This issue can occur if any ILM rule uses the **Dual commit** ingest option (or a rule specifies **Balanced** and it is not possible to immediately create all required copies). In this case, StorageGRID will commit newly ingested objects to two Storage Nodes on the same site and evaluate ILM later.
 - To ensure you can access all objects while a Storage Node is rebooting, stop ingesting objects at a site for approximately one hour before rebooting the node.

About this task

When StorageGRID reboots a grid node, it issues the `reboot` command on the node, which causes the node to shut down and restart. All services are restarted automatically.

- Rebooting a VMware node reboots the virtual machine.
- Rebooting a Linux node reboots the container.
- Rebooting a StorageGRID Appliance node reboots the compute controller.



If you need to reboot more than one node, you can use the [rolling reboot procedure](#).

Steps

1. Select **NODES**.
2. Select the grid node you want to reboot.
3. Select the **Tasks** tab.
4. Select **Reboot**.

A confirmation dialog box appears. If you are rebooting the primary Admin Node, the confirmation dialog box reminds you that your browser's connection to the Grid Manager will be lost temporarily when services are stopped.

5. Enter the provisioning passphrase, and select **OK**.
6. Wait for the node to reboot.

It might take some time for services to shut down.

When the node is rebooting, the gray (Administratively Down) icon appears for the node on the Nodes page. When all services have started again and the node is successfully connected to the grid, the Nodes page should display normal status (no icons to the left of the node name), indicating that no alerts are active and the node is connected to the grid.

Reboot grid node from command shell

If you need to monitor the reboot operation more closely or if you are unable to access the Grid Manager, you can log in to the grid node and run the Server Manager reboot command from the command shell.

Before you begin

You have the `Passwords.txt` file.

Steps

1. Log in to the grid node:

- a. Enter the following command: `ssh admin@grid_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 `#`.

2. Optionally, stop services: `service servermanager stop`

Stopping services is an optional, but recommended step. Services can take up to 15 minutes to shut down, and you might want to log in to the system remotely to monitor the shutdown process before you reboot the node in the next step.

3. Reboot the grid node: `reboot`

4. Log out of the command shell: `exit`

Shut down grid node

You can shut down a grid node from the node's command shell.

Before you begin

- You have the `Passwords.txt` file.

About this task

Before performing this procedure, review these considerations:

- In general, you should not shut down more than one node at a time to avoid disruptions.
- Don't shut down a node during a maintenance procedure unless explicitly instructed to do so by the documentation or by technical support.
- The shutdown process is based on where the node is installed, as follows:
 - Shutting down a VMware node shuts down the virtual machine.
 - Shutting down a Linux node shuts down the container.
 - Shutting down a StorageGRID appliance node shuts down the compute controller.
- If you plan to shut down more than one Storage Node at a site, stop ingesting objects at that site for approximately one hour before shutting down the nodes.

If any ILM rule uses the **Dual commit** ingest option (or if a rule uses the **Balanced** option and all required copies can't be created immediately), StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site and evaluates ILM later. If more than one Storage Node at a site is shut down, you might not be able to access newly ingested objects for the duration of the shutdown. Write operations might also fail if too few Storage Nodes remain available at the site. See [Manage objects with ILM](#).

Steps

1. Log in to the grid node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

2. Stop all services: `service servermanager stop`

Services can take up to 15 minutes to shut down, and you might want to log in to the system remotely to monitor the shutdown process.

3. If the node is running on a VMware virtual machine or it is an appliance node, issue the shutdown command: `shutdown -h now`

Perform this step regardless of the outcome of the `service servermanager stop` command.



After you issue the `shutdown -h now` command on an appliance node, you must power cycle the appliance to restart the node.

For the appliance, this command shuts down the controller, but the appliance is still powered on. You must complete the next step.

4. If you are powering down an appliance node, follow the steps for your appliance.

SG6160

- a. Turn off the power to the SG6100-CN storage controller.
- b. Wait for the blue power LED on the SG6100-CN storage controller to turn off.

SGF6112

- a. Turn off the power to the appliance.
- b. Wait for the blue power LED to turn off.

SG6000

- a. Wait for the green Cache Active LED on the back of the storage controllers to turn off.

This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

- b. Turn off the power to the appliance, and wait for the blue power LED to turn off.

SG5800

- a. Wait for the green Cache Active LED on the back of the storage controller to turn off.

This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

- b. From the home page of SANtricity System Manager, select **View Operations in Progress**.
- c. Confirm that all operations have completed before continuing with the next step.
- d. Turn off both power switches on the controller shelf, and wait for all LEDs on the controller shelf to turn off.

SG5700

- a. Wait for the green Cache Active LED on the back of the storage controller to turn off.

This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

- b. Turn off the power to the appliance, and wait for all LED and seven-segment display activity to stop.

SG100 or SG1000

- a. Turn off the power to the appliance.
- b. Wait for the blue power LED to turn off.

Power down host

Before you power down a host, you must stop services on all grid nodes on that host.

Steps

1. Log in to the grid node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

- 2. Stop all services running on the node: `service servermanager stop`

Services can take up to 15 minutes to shut down, and you might want to log in to the system remotely to monitor the shutdown process.

- 3. Repeat steps 1 and 2 for each node on the host.

- 4. If you have a Linux host:

- a. Log in to the host operating system.
- b. Stop the node: `storagegrid node stop`
- c. Shut down the host operating system.

- 5. If the node is running on a VMware virtual machine or it is an appliance node, issue the shutdown command: `shutdown -h now`

Perform this step regardless of the outcome of the `service servermanager stop` command.



After you issue the `shutdown -h now` command on an appliance node, you must power cycle the appliance to restart the node.

For the appliance, this command shuts down the controller, but the appliance is still powered on. You must complete the next step.

- 6. If you are powering down an appliance node, follow the steps for your appliance.

SG6160

- a. Turn off the power to the SG6100-CN storage controller.
- b. Wait for the blue power LED on the SG6100-CN storage controller to turn off.

SGF6112

- a. Turn off the power to the appliance.
- b. Wait for the blue power LED to turn off.

SG6000

- a. Wait for the green Cache Active LED on the back of the storage controllers to turn off.

This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

- b. Turn off the power to the appliance, and wait for the blue power LED to turn off.

SG5800

- a. Wait for the green Cache Active LED on the back of the storage controller to turn off.

This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

- b. From the home page of SANtricity System Manager, select **View Operations in Progress**.
- c. Confirm that all operations have completed before continuing with the next step.
- d. Turn off both power switches on the controller shelf, and wait for all LEDs on the controller shelf to turn off.

SG5700

- a. Wait for the green Cache Active LED on the back of the storage controller to turn off.

This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

- b. Turn off the power to the appliance, and wait for all LED and seven-segment display activity to stop.

SG110 or SG1100

- a. Turn off the power to the appliance.
- b. Wait for the blue power LED to turn off.

SG100 or SG1000

- a. Turn off the power to the appliance.
- b. Wait for the blue power LED to turn off.

7. Log out of the command shell: `exit`

Related information

- [SGF6112 and SG6160 storage appliances](#)

- [SG6000 storage appliances](#)
- [SG5700 storage appliances](#)
- [SG5800 storage appliances](#)
- [SG110 and SG1100 services appliances](#)
- [SG100 and SG1000 services appliances](#)

Power off and on all nodes in grid

You might need to shut down your entire StorageGRID system, for example, if you are moving a data center. These steps provide a high-level overview of the recommended sequence for performing a controlled shutdown and startup.

When you power off all nodes in a site or grid, you will not be able to access ingested objects while the Storage Nodes are offline.

Stop services and shut down grid nodes

Before you can power off a StorageGRID system, you must stop all services running on each grid node, and then shut down all VMware virtual machines, container engines, and StorageGRID appliances.

About this task

Stop services on Admin Nodes and Gateway Nodes first, and then stop services on Storage Nodes.

This approach allows you to use the primary Admin Node to monitor the status of the other grid nodes for as long as possible.



If a single host includes more than one grid node, don't shut down the host until you have stopped all of the nodes on that host. If the host includes the primary Admin Node, shut down that host last.



If required, you can [migrate nodes from one Linux host to another](#) to perform host maintenance without impacting the functionality or availability of your grid.

Steps

1. Stop all client applications from accessing the grid.
2. Log in to each Gateway Node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.
3. Stop all services running on the node: `service servermanager stop`

Services can take up to 15 minutes to shut down, and you might want to log in to the system remotely to monitor the shutdown process.

4. Repeat the previous two steps to stop the services on all Storage Nodes and non-primary Admin Nodes.

You can stop the services on these nodes in any order.



If you issue the `service servermanager stop` command to stop the services on an appliance Storage Node, you must power cycle the appliance to restart the node.

5. For the primary Admin Node, repeat the steps for [logging into the node](#) and [stopping all services on the node](#).
6. For nodes that are running on Linux hosts:
 - a. Log in to the host operating system.
 - b. Stop the node: `storagegrid node stop`
 - c. Shut down the host operating system.
7. For nodes that are running on VMware virtual machines and for appliance Storage Nodes, issue the shutdown command: `shutdown -h now`

Perform this step regardless of the outcome of the `service servermanager stop` command.

For the appliance, this command shuts down the compute controller, but the appliance is still powered on. You must complete the next step.

8. If you have appliance nodes, follow the steps for your appliance.

SG110 or SG1100

- a. Turn off the power to the appliance.
- b. Wait for the blue power LED to turn off.

SG100 or SG1000

- a. Turn off the power to the appliance.
- b. Wait for the blue power LED to turn off.

SG6160

- a. Turn off the power to the SG6100-CN storage controller.
- b. Wait for the blue power LED on the SG6100-CN storage controller to turn off.

SGF6112

- a. Turn off the power to the appliance.
- b. Wait for the blue power LED to turn off.

SG6000

- a. Wait for the green Cache Active LED on the back of the storage controllers to turn off.

This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

- b. Turn off the power to the appliance, and wait for the blue power LED to turn off.

SG5800

- a. Wait for the green Cache Active LED on the back of the storage controller to turn off.

This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

- b. From the home page of SANtricity System Manager, select **View Operations in Progress**.
- c. Confirm that all operations have completed before continuing with the next step.
- d. Turn off both power switches on the controller shelf, and wait for all LEDs on the controller shelf to turn off.

SG5700

- a. Wait for the green Cache Active LED on the back of the storage controller to turn off.

This LED is on when cached data needs to be written to the drives. You must wait for this LED to turn off before you turn off power.

- b. Turn off the power to the appliance, and wait for all LED and seven-segment display activity to stop.

9. If required, log out of the command shell: `exit`

The StorageGRID grid has now been shut down.

Start up grid nodes



If the entire grid has been shut down for more than 15 days, you must contact technical support before starting up any grid nodes. Don't attempt the recovery procedures that rebuild Cassandra data. Doing so might result in data loss.

If possible, power on the grid nodes in this order:

- Apply power to Admin Nodes first.
- Apply power to Gateway Nodes last.



If a host includes multiple grid nodes, the nodes will come back online automatically when you power on the host.

Steps

1. Power on the hosts for the primary Admin Node and any non-primary Admin Nodes.



You will not be able to log in to the Admin Nodes until the Storage Nodes have been restarted.

2. Power on the hosts for all Storage Nodes.

You can power on these nodes in any order.

3. Power on the hosts for all Gateway Nodes.
4. Sign in to the Grid Manager.
5. Select **NODES** and monitor the status of the grid nodes. Verify that there are no alert icons next to the node names.

Related information

- [SGF6112 and SG6160 storage appliances](#)
- [SG110 and SG1100 services appliances](#)
- [SG100 and SG1000 services appliances](#)
- [SG6000 storage appliances](#)
- [SG5800 storage appliances](#)
- [SG5700 storage appliances](#)

Port remap procedures

Remove port remaps

If you want to configure an endpoint for the Load Balancer service, and you want to use a port that has already been configured as the Mapped-To Port of a port remap, you must first remove the existing port remap, or the endpoint will not be effective. You must run a script on each Admin Node and Gateway Node that has conflicting remapped ports to remove all of the node's port remaps.

About this task

This procedure removes all port remaps. If you need to keep some of the remaps, contact technical support.

For information about configuring load balancer endpoints, see [Configuring load balancer endpoints](#).



If the port remap provides client access, reconfigure the client to use a different port as an load balancer endpoint to avoid loss of service. Otherwise, removing the port mapping will result in loss of client access and should be scheduled appropriately.



This procedure does not work for a StorageGRID system deployed as a container on bare metal hosts. See the instructions for [removing port remaps on bare metal hosts](#).

Steps

1. Log in to the node.

a. Enter the following command: `ssh -p 8022 admin@node_IP`

Port 8022 is the SSH port of the base OS, while port 22 is the SSH port of the container engine running StorageGRID.

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 `#`.

2. Run the following script: `remove-port-remap.sh`

3. Reboot the node: `reboot`

4. Log out of the command shell: `exit`

5. Repeat these steps on each Admin Node and Gateway Node that has conflicting remapped ports.

Remove port remaps on bare metal hosts

If you want to configure an endpoint for the Load Balancer service, and you want to use a port that has already been configured as the Mapped-To Port of a port remap, you must first remove the existing port remap, or the endpoint will not be effective.

About this task

If you are running StorageGRID on bare metal hosts, follow this procedure instead of the general procedure for removing port remaps. You must edit the node configuration file for each Admin Node and Gateway Node that has conflicting remapped ports to remove all of the node's port remaps and restart the node.



This procedure removes all port remaps. If you need to keep some of the remaps, contact technical support.

For information about configuring load balancer endpoints, see the instructions for administering StorageGRID.



This procedure can result in temporary loss of service as nodes are restarted.

Steps

1. Log in to the host supporting the node. Log in as root or with an account that has sudo permission.
2. Run the following command to temporarily disable the node: `sudo storagegrid node stop node-name`
3. Using a text editor such as vim or pico, edit the node configuration file for the node.

The node configuration file can be found at `/etc/storagegrid/nodes/node-name.conf`.

4. Locate the section of the node configuration file that contains the port remaps.

See the last two lines in the following example.

```
ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_ESL = 10.0.0.0/8, 172.19.0.0/16, 172.21.0.0/16
ADMIN_NETWORK_GATEWAY = 10.224.0.1
ADMIN_NETWORK_IP = 10.224.5.140
ADMIN_NETWORK_MASK = 255.255.248.0
ADMIN_NETWORK_MTU = 1400
ADMIN_NETWORK_TARGET = eth1
ADMIN_NETWORK_TARGET_TYPE = Interface
BLOCK_DEVICE_VAR_LOCAL = /dev/sda2
CLIENT_NETWORK_CONFIG = STATIC
CLIENT_NETWORK_GATEWAY = 47.47.0.1
CLIENT_NETWORK_IP = 47.47.5.140
CLIENT_NETWORK_MASK = 255.255.248.0
CLIENT_NETWORK_MTU = 1400
CLIENT_NETWORK_TARGET = eth2
CLIENT_NETWORK_TARGET_TYPE = Interface
GRID_NETWORK_CONFIG = STATIC
GRID_NETWORK_GATEWAY = 192.168.0.1
GRID_NETWORK_IP = 192.168.5.140
GRID_NETWORK_MASK = 255.255.248.0
GRID_NETWORK_MTU = 1400
GRID_NETWORK_TARGET = eth0
GRID_NETWORK_TARGET_TYPE = Interface
NODE_TYPE = VM_API_Gateway
PORT_REMAP = client/tcp/8082/443
PORT_REMAP_INBOUND = client/tcp/8082/443
```

5. Edit the `PORT_REMAP` and `PORT_REMAP_INBOUND` entries to remove port remaps.

```
PORT_REMAP =
PORT_REMAP_INBOUND =
```

6. Run the following command to validate your changes to the node configuration file for the node: `sudo storagegrid node validate node-name`

Address any errors or warnings before proceeding to the next step.

7. Run the following command to restart the node without port remaps: `sudo storagegrid node start node-name`
8. Log in to the node as admin using the password listed in the `Passwords.txt` file.
9. Verify that the services start correctly.
 - a. View a list of the statuses of all services on the server: `sudo storagegrid-status`

The status is updated automatically.
 - b. Wait until all services have a status of either Running or Verified.
 - c. Exit the status screen: `Ctrl+C`
10. Repeat these steps on each Admin Node and Gateway Node that has conflicting remapped ports.

Network procedures

Update subnets for Grid Network

StorageGRID maintains a list of the network subnets used to communicate between grid nodes on the Grid Network (eth0). These entries include the subnets used for the Grid Network by each site in your StorageGRID system as well as any subnets used for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway. When you add grid nodes or a new site in an expansion, you might need to update or add subnets to the Grid Network.

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 the provisioning passphrase.
- You have the network addresses, in CIDR notation, of the subnets you want to configure.

About this task

If you are performing an expansion activity that includes adding a new subnet, you must add a new subnet to the Grid Network subnet list before you start the expansion procedure. Otherwise, you will have to cancel the expansion, add the new subnet, and start the expansion again.

Do not use subnets that contain the following IPv4 addresses for the Grid Network, Admin Network, or Client Network of any node:

- 192.168.130.101
- 192.168.131.101
- 192.168.130.102
- 192.168.131.102
- 198.51.100.2
- 198.51.100.4



For example, do not use the following subnet ranges for the Grid Network, Admin Network, or Client Network of any node:

- 192.168.130.0/24 because this subnet range contains the IP addresses 192.168.130.101 and 192.168.130.102
- 192.168.131.0/24 because this subnet range contains the IP addresses 192.168.131.101 and 192.168.131.102
- 198.51.100.0/24 because this subnet range contains the IP addresses 198.51.100.2 and 198.51.100.4

Add a subnet

Steps

1. Select **MAINTENANCE > Network > Grid Network**.
2. Select **Add another subnet** to add a new subnet in CIDR notation.

For example, enter 10.96.104.0/22.

3. Enter the provisioning passphrase, and select **Save**.
4. Wait until the changes are applied, then download a new Recovery Package.
 - a. Select **MAINTENANCE > System > Recovery package**.
 - b. Enter the **Provisioning Passphrase**.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system. It is also used to recover the primary Admin Node.

The subnets you have specified are configured automatically for your StorageGRID system.

Edit a subnet

Steps

1. Select **MAINTENANCE > Network > Grid Network**.
2. Select the subnet you want to edit and make the necessary changes.
3. Enter the Provisioning passphrase, and select **Save**.
4. Select **Yes** in the confirmation dialog box.

5. Wait until the changes are applied, then download a new Recovery Package.
 - a. Select **MAINTENANCE > System > Recovery package**.
 - b. Enter the **Provisioning Passphrase**.

Delete a subnet

Steps

1. Select **MAINTENANCE > Network > Grid Network**.
2. Select the delete icon  next to the subnet.
3. Enter the Provisioning passphrase, and select **Save**.
4. Select **Yes** in the confirmation dialog box.
5. Wait until the changes are applied, then download a new Recovery Package.
 - a. Select **MAINTENANCE > System > Recovery package**.
 - b. Enter the **Provisioning Passphrase**.

Configure IP addresses

IP address guidelines

You can perform network configuration by configuring IP addresses for grid nodes using the Change IP tool.

You must use the Change IP tool to make most changes to the networking configuration that was initially set during grid deployment. Manual changes using standard Linux networking commands and files might not propagate to all StorageGRID services, and might not persist across upgrades, reboots, or node recovery procedures.



The IP change procedure can be a disruptive procedure. Parts of the grid might be unavailable until the new configuration is applied.



If you are making changes to the Grid Network Subnet List only, use the Grid Manager to add or change the network configuration. Otherwise, use the Change IP tool if the Grid Manager is inaccessible due to a network configuration issue, or you are performing both a Grid Network routing change and other network changes at the same time.



If you want to change the Grid Network IP address for all nodes in the grid, use the [special procedure for grid-wide changes](#).

Ethernet interfaces

The IP address assigned to eth0 is always the grid node's Grid Network IP address. The IP address assigned to eth1 is always the grid node's Admin Network IP address. The IP address assigned to eth2 is always the grid node's Client Network IP address.

Note that on some platforms, such as StorageGRID appliances, eth0, eth1, and eth2 might be aggregate interfaces composed of subordinate bridges or bonds of physical or VLAN interfaces. On these platforms, the **SSM > Resources** tab might show the Grid, Admin, and Client Network IP address assigned to other interfaces in addition to eth0, eth1, or eth2.

DHCP

You can only set up DHCP during the deployment phase. You can't set up DHCP during configuration. You must use the IP address change procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. Using the Change IP tool will cause DHCP addresses to become static.

High availability (HA) groups

- If a Client Network interface is contained in an HA group, you can't change the Client Network IP address for that interface to an address that is outside of the subnet configured for the HA group.
- You can't change the Client Network IP address to the value of an existing virtual IP address assigned to an HA group configured on the Client Network interface.
- If a Grid network interface is contained in an HA group, you can't change the Grid network IP address for that interface to an address that is outside of the subnet configured for the HA group.
- You can't change the Grid Network IP address to the value of an existing virtual IP address assigned to an HA group configured on the Grid Network interface.

Change node network configuration

You can change the network configuration of one or more nodes using the Change IP tool. You can change the configuration of the Grid Network, or add, change, or remove the Admin or Client Networks.

Before you begin

You have the `Passwords.txt` file.

About this task

Linux: If you are adding a grid node to the Admin Network or Client Network for the first time, and you did not previously configure `ADMIN_NETWORK_TARGET` or `CLIENT_NETWORK_TARGET` in the node configuration file, you must do so now.

See the StorageGRID installation instructions for your Linux operating system:

- [Install StorageGRID on Red Hat Enterprise Linux](#)
- [Install StorageGRID on Ubuntu or Debian](#)

Appliances: On StorageGRID appliances, if the Client or Admin Network was not configured in the StorageGRID Appliance Installer during the initial installation, the network can't be added by using only the Change IP tool. First, you must [place the appliance in maintenance mode](#), configure the links, return the appliance to normal operating mode, and then use the Change IP tool to modify the network configuration. See the [procedure for configuring network links](#).

You can change the IP address, subnet mask, gateway, or MTU value for one or more nodes on any network.

You can also add or remove a node from a Client Network or from an Admin Network:

- You can add a node to a Client Network or to an Admin Network by adding an IP address/subnet mask on that network to the node.
- You can remove a node from a Client Network or from an Admin Network by deleting the IP address/subnet mask for the node on that network.

Nodes can't be removed from the Grid Network.



IP address swaps aren't allowed. If you must exchange IP addresses between grid nodes, you must use a temporary intermediate IP address.



If single sign-on (SSO) is enabled for your StorageGRID system and you are changing the IP address of an Admin Node, be aware that any relying party trust that was configured using the Admin Node's IP address (instead of its fully qualified domain name, as recommended) will become invalid. You will no longer be able to sign in to the node. Immediately after changing the IP address, you must update or reconfigure the node's relying party trust in Active Directory Federation Services (AD FS) with the new IP address. See the instructions for [configuring SSO](#).



Any changes you make to the network using the Change IP tool are propagated to the installer firmware for the StorageGRID appliances. That way, if StorageGRID software is reinstalled on an appliance, or if an appliance is placed into maintenance mode, the networking configuration will be correct.

Steps

1. Log in to the primary Admin Node:

- a. Enter the following command: `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 `#`.

2. Start the Change IP tool by entering the following command: `change-ip`

3. Enter the provisioning passphrase at the prompt.

The main menu appears.

```
Welcome to the StorageGRID IP Change Tool.

Selected nodes: all

1:  SELECT NODES to edit
2:  EDIT IP/mask, gateway and MTU
3:  EDIT admin network subnet lists
4:  EDIT grid network subnet list
5:  SHOW changes
6:  SHOW full configuration, with changes highlighted
7:  VALIDATE changes
8:  SAVE changes, so you can resume later
9:  CLEAR all changes, to start fresh
10: APPLY changes to the grid
0:  Exit

Selection: █
```

4. Optionally select **1** to choose which nodes to update. Then select one of the following options:

- **1:** Single node — select by name
- **2:** Single node — select by site, then by name
- **3:** Single node — select by current IP

- **4:** All nodes at a site
- **5:** All nodes in the grid

Note: If you want to update all nodes, allow "all" to remain selected.

After you make your selection, the main menu appears, with the **Selected nodes** field updated to reflect your choice. All subsequent actions are performed only on the nodes displayed.

5. On the main menu, select option **2** to edit IP/mask, gateway, and MTU information for the selected nodes.

a. Select the network where you want to make changes:

- **1:** Grid network
- **2:** Admin network
- **3:** Client network
- **4:** All networks

After you make your selection, the prompt shows the node name, network name (Grid, Admin, or Client), data type (IP/mask, Gateway, or MTU), and current value.

Editing the IP address, prefix length, gateway, or MTU of a DHCP-configured interface will change the interface to static. When you select to change an interface configured by DHCP, a warning is displayed to inform you that the interface will change to static.

Interfaces configured as `fixed` can't be edited.

b. To set a new value, enter it in the format shown for the current value.

c. To leave the current value unchanged, press **Enter**.

d. If the data type is `IP/mask`, you can delete the Admin or Client Network from the node by entering **d** or **0.0.0.0/0**.

e. After editing all nodes you want to change, enter **q** to return to the main menu.

Your changes are held until cleared or applied.

6. Review your changes by selecting one of the following options:

- **5:** Shows edits in output that is isolated to show only the changed item. Changes are highlighted in green (additions) or red (deletions), as shown in the example output:

```

=====
Site: RTP
=====
username-x Grid IP [ 172.16.0.239/21 ]: 172.16.0.240/21
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Admin IP [ 10.224.0.244/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.245/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.240/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.241/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.242/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.243/21 ]: 0.0.0.0/0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
Press Enter to continue

```

- 6: Shows edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions).



Certain command line interfaces might show additions and deletions using strikethrough formatting. Proper display depends on your terminal client supporting the necessary VT100 escape sequences.

7. Select option 7 to validate all changes.

This validation ensures that the rules for the Grid, Admin, and Client Networks, such as not using overlapping subnets, aren't violated.

In this example, validation returned errors.

```

Validating new networking configuration... FAILED.

DK-10-224-S-20-G1: The admin subnet 172.18.0.0/16 overlaps the 172.18.0.0/21 grid network.
DK-10-224-S-22-S1: Duplicate Grid IP 172.16.5.18 (also in use by DK-10-224-S-21-ADM1)

You must correct these errors before you can apply any changes.
Checking for Grid Network IP address swaps... PASSED.

Press Enter to continue

```

In this example, validation passed.

```

Validating new networking configuration... PASSED.
Checking for Grid Network IP address swaps... PASSED.

Press Enter to continue

```

8. After validation passes, choose one of the following options:

- **8:** Save unapplied changes.

This option allows you to quit the Change IP tool and start it again later, without losing any unapplied changes.

- **10:** Apply the new network configuration.

9. If you selected option **10**, choose one of the following options:

- **apply:** Apply the changes immediately and automatically restart each node if necessary.

If the new network configuration does not require any physical networking changes, you can select **apply** to apply the changes immediately. Nodes will be restarted automatically, if necessary. Nodes that need to be restarted will be displayed.

- **stage:** Apply the changes the next time the nodes are restarted manually.

If you need to make physical or virtual networking configuration changes for the new network configuration to function, you must use the **stage** option, shut down the affected nodes, make the necessary physical networking changes, and restart the affected nodes. If you select **apply** without first making these networking changes, the changes will usually fail.



If you use the **stage** option, you must restart the node as soon as possible after staging to minimize disruptions.

- **cancel:** Don't make any network changes at this time.

If you were unaware that the proposed changes require nodes to be restarted, you can defer the changes to minimize user impact. Selecting **cancel** returns you to the main menu and preserves your changes so you can apply them later.

When you select **apply** or **stage**, a new network configuration file is generated, provisioning is performed, and nodes are updated with new working information.

During provisioning, the output displays the status as updates are applied.

```
Generating new grid networking description file...
```

```
Running provisioning...
```

```
Updating grid network configuration on Name
```

After you apply or stage changes, a new Recovery Package is generated as a result of the grid configuration change.

10. If you selected **stage**, follow these steps after provisioning is complete:

- a. Make the physical or virtual networking changes that are required.

Physical networking changes: Make the necessary physical networking changes, safely shutting down the node if necessary.

Linux: If you are adding the node to an Admin Network or Client Network for the first time, ensure that you have added the interface as described in [Linux: Add interfaces to existing node](#).

- b. Restart the affected nodes.
11. Select **0** to exit the Change IP tool after your changes are complete.
12. Download a new Recovery Package from the Grid Manager.
 - a. Select **MAINTENANCE > System > Recovery package**.
 - b. Enter the provisioning passphrase.

Add to or change subnet lists on Admin Network

You can add, delete, or change the subnets in the Admin Network Subnet List of one or more nodes.

Before you begin

- You have the `Passwords.txt` file.

You can add, delete, or change subnets to all nodes on the Admin Network Subnet List.

Do not use subnets that contain the following IPv4 addresses for the Grid Network, Admin Network, or Client Network of any node:

- 192.168.130.101
- 192.168.131.101
- 192.168.130.102
- 192.168.131.102
- 198.51.100.2
- 198.51.100.4



For example, do not use the following subnet ranges for the Grid Network, Admin Network, or Client Network of any node:

- 192.168.130.0/24 because this subnet range contains the IP addresses 192.168.130.101 and 192.168.130.102
- 192.168.131.0/24 because this subnet range contains the IP addresses 192.168.131.101 and 192.168.131.102
- 198.51.100.0/24 because this subnet range contains the IP addresses 198.51.100.2 and 198.51.100.4

Steps

1. Log in to the primary Admin Node:
 - a. Enter the following command: `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 #.

2. Start the Change IP tool by entering the following command: `change-ip`
3. Enter the provisioning passphrase at the prompt.

The main menu appears.

```
Welcome to the StorageGRID IP Change Tool.

Selected nodes: all

1:  SELECT NODES to edit
2:  EDIT IP/mask, gateway and MTU
3:  EDIT admin network subnet lists
4:  EDIT grid network subnet list
5:  SHOW changes
6:  SHOW full configuration, with changes highlighted
7:  VALIDATE changes
8:  SAVE changes, so you can resume later
9:  CLEAR all changes, to start fresh
10: APPLY changes to the grid
0:  Exit

Selection: █
```

4. Optionally, limit the networks/nodes on which operations are performed. Choose one of the following:
 - Select the nodes to edit by choosing **1**, if you want to filter on specific nodes on which to perform the operation. Select one of the following options:
 - **1**: Single node (select by name)
 - **2**: Single node (select by site, then by name)
 - **3**: Single node (select by current IP)
 - **4**: All nodes at a site
 - **5**: All nodes in the grid
 - **0**: Go back
 - Allow "all" to remain selected.

After the selection is made, the main menu screen appears. The Selected nodes field reflects your new selection, and now all operations selected will only be performed on this item.
5. On the main menu, select the option to edit subnets for the Admin Network (option **3**).
6. Choose one of the following:
 - Add a subnet by entering this command: `add CIDR`
 - Delete a subnet by entering this command: `del CIDR`
 - Set the list of subnets by entering this command: `set CIDR`



For all commands, you can enter multiple addresses using this format: `add CIDR, CIDR`

Example: `add 172.14.0.0/16, 172.15.0.0/16, 172.16.0.0/16`



You can reduce the amount of typing required by using "up arrow" to recall previously typed values to the current input prompt, and then edit them if necessary.

The example input below shows adding subnets to the Admin Network Subnet List:

```
Editing: Admin Network Subnet List for node DK-10-224-5-20-G1

Press <enter> to use the list as shown
Use up arrow to recall a previously typed value, which you can then edit
Use 'add <CIDR> [, <CIDR>]' to add subnets <CIDR> [, <CIDR>] to the list
Use 'del <CIDR> [, <CIDR>]' to delete subnets <CIDR> [, <CIDR>] from the list
Use 'set <CIDR> [, <CIDR>]' to set the list to the given list
Use q to complete the editing session early and return to the previous menu

DK-10-224-5-20-G1
 10.0.0.0/8
 172.19.0.0/16
 172.21.0.0/16
 172.20.0.0/16

[add/del/set/quit <CIDR>, ...]: add 172.14.0.0/16, 172.15.0.0/16
```

7. When ready, enter **q** to go back to the main menu screen. Your changes are held until cleared or applied.



If you selected any of the "all" node selection modes in step 2, press **Enter** (without **q**) to get to the next node in the list.

8. Choose one of the following:

- Select option **5** to show edits in output that is isolated to show only the changed item. Changes are highlighted in green (additions) or red (deletions), as shown in the example output below:

```
=====
Site: Data Center 1
=====
DC1-ADM1-105-154 Admin Subnets
                                     add 172.17.0.0/16
                                     del 172.16.0.0/16
                                     [ 172.14.0.0/16 ]
                                     [ 172.15.0.0/16 ]
                                     [ 172.17.0.0/16 ]
                                     [ 172.19.0.0/16 ]
                                     [ 172.20.0.0/16 ]
                                     [ 172.21.0.0/16 ]
Press Enter to continue
```

- Select option **6** to show edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions).

Note: Certain terminal emulators might show additions and deletions using strikethrough formatting.

When you attempt to change the subnet list, the following message is displayed:

CAUTION: The Admin Network subnet list on the node might contain /32 subnets derived from automatically applied routes that aren't persistent. Host routes (/32 subnets) are applied automatically if the IP addresses provided for external services such as NTP or DNS aren't reachable using default StorageGRID routing, but are reachable using a different interface and gateway. Making and applying changes to the subnet list will make all automatically applied subnets persistent. If you don't want that to happen, delete the unwanted subnets before applying changes. If you know that all /32 subnets in the list were added intentionally, you can ignore this caution.

If you did not specifically assign the NTP and DNS server subnets to a network, StorageGRID creates a host route (/32) for the connection automatically. If, for example, you would rather have a /16 or /24 route for outbound connection to a DNS or NTP server, you should delete the automatically created /32 route and add the routes you want. If you don't delete the automatically created host route, it will be persisted after you apply any changes to the subnet list.



Although you can use these automatically discovered host routes, in general you should manually configure the DNS and NTP routes to ensure connectivity.

9. Select option **7** to validate all staged changes.

This validation ensures that the rules for the Grid, Admin, and Client Networks are followed, such as using overlapping subnets.

10. Optionally, select option **8** to save all staged changes and return later to continue making changes.

This option allows you to quit the Change IP tool and start it again later, without losing any unapplied changes.

11. Do one of the following:

- Select option **9** if you want to clear all changes without saving or applying the new network configuration.
- Select option **10** if you are ready to apply changes and provision the new network configuration. During provisioning, the output displays the status as updates are applied as shown in the following example output:

```
Generating new grid networking description file...
```

```
Running provisioning...
```

```
Updating grid network configuration on Name
```

12. Download a new Recovery Package from the Grid Manager.
 - a. Select **MAINTENANCE > System > Recovery package**.
 - b. Enter the provisioning passphrase.

Add to or change subnet lists on Grid Network

You can use the Change IP tool to add or change subnets on the Grid Network.

Before you begin

- You have the `Passwords.txt` file.

You can add, delete, or change subnets in the Grid Network Subnet List. Changes will affect routing on all nodes in the grid.



If you are making changes to the Grid Network Subnet List only, use the Grid Manager to add or change the network configuration. Otherwise, use the Change IP tool if the Grid Manager is inaccessible due to a network configuration issue, or you are performing both a Grid Network routing change and other network changes at the same time.

Do not use subnets that contain the following IPv4 addresses for the Grid Network, Admin Network, or Client Network of any node:

- 192.168.130.101
- 192.168.131.101
- 192.168.130.102
- 192.168.131.102
- 198.51.100.2
- 198.51.100.4



For example, do not use the following subnet ranges for the Grid Network, Admin Network, or Client Network of any node:

- 192.168.130.0/24 because this subnet range contains the IP addresses 192.168.130.101 and 192.168.130.102
- 192.168.131.0/24 because this subnet range contains the IP addresses 192.168.131.101 and 192.168.131.102
- 198.51.100.0/24 because this subnet range contains the IP addresses 198.51.100.2 and 198.51.100.4

Steps

1. Log in to the primary Admin Node:
 - a. Enter the following command: `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 `#`.

2. Start the Change IP tool by entering the following command: `change-ip`
3. Enter the provisioning passphrase at the prompt.

The main menu appears.

```
Welcome to the StorageGRID IP Change Tool.

Selected nodes: all

1:  SELECT NODES to edit
2:  EDIT IP/mask, gateway and MTU
3:  EDIT admin network subnet lists
4:  EDIT grid network subnet list
5:  SHOW changes
6:  SHOW full configuration, with changes highlighted
7:  VALIDATE changes
8:  SAVE changes, so you can resume later
9:  CLEAR all changes, to start fresh
10: APPLY changes to the grid
0:  Exit

Selection: █
```

4. On the main menu, select the option to edit subnets for the Grid Network (option 4).



Changes to the Grid Network Subnet List are grid-wide.

5. Choose one of the following:

- Add a subnet by entering this command: `add CIDR`
- Delete a subnet by entering this command: `del CIDR`
- Set the list of subnets by entering this command: `set CIDR`



For all commands, you can enter multiple addresses using this format: `add CIDR, CIDR`

Example: `add 172.14.0.0/16, 172.15.0.0/16, 172.16.0.0/16`



You can reduce the amount of typing required by using "up arrow" to recall previously typed values to the current input prompt, and then edit them if necessary.

The example input below shows setting subnets for the Grid Network Subnet List:

```
Editing: Grid Network Subnet List

Press <enter> to use the list as shown
Use up arrow to recall a previously typed value, which you can then edit
Use 'add <CIDR> [, <CIDR>]' to add subnets <CIDR> [, <CIDR>] to the list
Use 'del <CIDR> [, <CIDR>]' to delete subnets <CIDR> [, <CIDR>] from the list
Use 'set <CIDR> [, <CIDR>]' to set the list to the given list
Use q to complete the editing session early and return to the previous menu

Grid Network Subnet List
 172.16.0.0/21
 172.17.0.0/21
 172.18.0.0/21
192.168.0.0/21

[add/del/set/quit <CIDR>, ...]: set 172.30.0.0/21, 172.31.0.0/21, 192.168.0.0/21 █
```

6. When ready, enter **q** to go back to the main menu screen. Your changes are held until cleared or applied.

7. Choose one of the following:

- Select option **5** to show edits in output that is isolated to show only the changed item. Changes are highlighted in green (additions) or red (deletions), as shown in the example output below:

```
=====
Grid Network Subnet List (GNSL)
=====
                                     add 172.30.0.0/21
                                     add 172.31.0.0/21
                                     del 172.16.0.0/21
                                     del 172.17.0.0/21
                                     del 172.18.0.0/21
[      172.30.0.0/21 ]
[      172.31.0.0/21 ]
[      192.168.0.0/21 ]
Press Enter to continue
```

- Select option **6** to show edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions).



Certain command line interfaces might show additions and deletions using strikethrough formatting.

8. Select option **7** to validate all staged changes.

This validation ensures that the rules for the Grid, Admin, and Client Networks are followed, such as using overlapping subnets.

9. Optionally, select option **8** to save all staged changes and return later to continue making changes.

This option allows you to quit the Change IP tool and start it again later, without losing any unapplied changes.

10. Do one of the following:

- Select option **9** if you want to clear all changes without saving or applying the new network configuration.
- Select option **10** if you are ready to apply changes and provision the new network configuration. During provisioning, the output displays the status as updates are applied as shown in the following example output:

```
Generating new grid networking description file...

Running provisioning...

Updating grid network configuration on Name
```

11. If you selected option **10** when making Grid Network changes, select one of the following options:

- **apply**: Apply the changes immediately and automatically restart each node if necessary.

If the new network configuration will function simultaneously with the old network configuration without

any external changes, you can use the **apply** option for a fully automated configuration change.

- **stage**: Apply the changes the next time the nodes are restarted.

If you need to make physical or virtual networking configuration changes for the new network configuration to function, you must use the **stage** option, shut down the affected nodes, make the necessary physical networking changes, and restart the affected nodes.



If you use the **stage** option, restart the node as soon as possible after staging to minimize disruptions.

- **cancel**: Don't make any network changes at this time.

If you were unaware that the proposed changes require nodes to be restarted, you can defer the changes to minimize user impact. Selecting **cancel** returns you to the main menu and preserves your changes so you can apply them later.

After you apply or stage changes, a new Recovery Package is generated as a result of the grid configuration change.

12. If configuration is stopped due to errors, the following options are available:

- To terminate the IP change procedure and return to the main menu, enter **a**.
- To retry the operation that failed, enter **r**.
- To continue to the next operation, enter **c**.

The failed operation can be retried later by selecting option **10** (Apply Changes) from the main menu. The IP change procedure will not be complete until all operations have completed successfully.

- If you had to manually intervene (to reboot a node, for example) and are confident that the action the tool thinks has failed was actually completed successfully, enter **f** to mark it as successful and move to the next operation.

13. Download a new Recovery Package from the Grid Manager.

- a. Select **MAINTENANCE > System > Recovery package**.
- b. Enter the provisioning passphrase.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

Change IP addresses for all nodes in grid

If you need to change the Grid Network IP address for all nodes in the grid, you must follow this special procedure. You can't do a grid-wide Grid Network IP change using the procedure to change individual nodes.

Before you begin

- You have the `Passwords.txt` file.

To ensure that the grid starts up successfully, you must make all the changes at the same time.



This procedure applies to the Grid Network only. You can't use this procedure to change IP addresses on the Admin or Client Networks.

If you want to change the IP addresses and MTU for the nodes at one site only, follow the [Change node network configuration](#) instructions.

Steps

1. Plan ahead for changes that you need to make outside of the Change IP tool, such as changes to DNS or NTP, and changes to the single sign-on (SSO) configuration, if used.



If the existing NTP servers will not be accessible to the grid on the new IP addresses, add the new NTP servers before you perform the change-ip procedure.



If the existing DNS servers will not be accessible to the grid on the new IP addresses, add the new DNS servers before you perform the change-ip procedure.



If SSO is enabled for your StorageGRID system and any relying party trusts were configured using Admin Node IP addresses (instead of fully qualified domain names, as recommended), be prepared to update or reconfigure these relying party trusts in Active Directory Federation Services (AD FS) immediately after you change IP addresses. See [Configure single sign-on](#).



If necessary, add the new subnet for the new IP addresses.

2. Log in to the primary Admin Node:

- a. Enter the following command: `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 `#`.

3. Start the Change IP tool by entering the following command: `change-ip`
4. Enter the provisioning passphrase at the prompt.

The main menu appears. By default, the `Selected nodes` field is set to `all`.

```
Welcome to the StorageGRID IP Change Tool.

Selected nodes: all

1:  SELECT NODES to edit
2:  EDIT IP/mask, gateway and MTU
3:  EDIT admin network subnet lists
4:  EDIT grid network subnet list
5:  SHOW changes
6:  SHOW full configuration, with changes highlighted
7:  VALIDATE changes
8:  SAVE changes, so you can resume later
9:  CLEAR all changes, to start fresh
10: APPLY changes to the grid
0:  Exit

Selection: █
```

5. On the main menu, select **2** to edit IP/subnet mask, gateway, and MTU information for all the nodes.
 - a. Select **1** to make changes to the Grid Network.

After you make your selection, the prompt shows the node names, Grid Network name, data type (IP/mask, Gateway, or MTU), and current values.

Editing the IP address, prefix length, gateway, or MTU of a DHCP-configured interface will change the interface to static. A warning is displayed before each interface configured by DHCP.

Interfaces configured as `fixed` can't be edited.

- b. To set a new value, enter it in the format shown for the current value.
 - c. After editing all nodes you want to change, enter **q** to return to the main menu.

Your changes are held until cleared or applied.

6. Review your changes by selecting one of the following options:
 - **5**: Shows edits in output that is isolated to show only the changed item. Changes are highlighted in green (additions) or red (deletions), as shown in the example output:

```

=====
Site: RTP
=====
username-x Grid IP [ 172.16.0.239/21 ]: 172.16.0.240/21
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Grid MTU [ 1400 ]: 9000
username-x Admin IP [ 10.224.0.244/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.245/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.240/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.241/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.242/21 ]: 0.0.0.0/0
username-x Admin IP [ 10.224.0.243/21 ]: 0.0.0.0/0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin Gateway [ 10.224.0.1 ]: 0.0.0.0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
username-x Admin MTU [ 1400 ]: 0
Press Enter to continue

```

- 6: Shows edits in output that displays the full configuration. Changes are highlighted in green (additions) or red (deletions).



Certain command line interfaces might show additions and deletions using strikethrough formatting. Proper display depends on your terminal client supporting the necessary VT100 escape sequences.

7. Select option 7 to validate all changes.

This validation ensures that the rules for the Grid Network, such as not using overlapping subnets, aren't violated.

In this example, validation returned errors.

```

Validating new networking configuration... FAILED.

DK-10-224-S-20-G1: The admin subnet 172.18.0.0/16 overlaps the 172.18.0.0/21 grid network.
DK-10-224-S-22-S1: Duplicate Grid IP 172.16.5.18 (also in use by DK-10-224-S-21-ADM1)

You must correct these errors before you can apply any changes.
Checking for Grid Network IP address swaps... PASSED.

Press Enter to continue

```

In this example, validation passed.

```

Validating new networking configuration... PASSED.
Checking for Grid Network IP address swaps... PASSED.

Press Enter to continue

```

8. After validation passes, select **10** to apply the new network configuration.
9. Select **stage** to apply the changes the next time the nodes are restarted.



You must select **stage**. Don't perform a rolling restart, either manually or by selecting **apply** instead of **stage**; the grid will not start up successfully.

10. After your changes are complete, select **0** to exit the Change IP tool.
11. Shut down all nodes simultaneously.



The entire grid must be shut down, so that all nodes are down at the same time.

12. Make the physical or virtual networking changes that are required.
13. Verify that all grid nodes are down.
14. Power on all nodes.
15. After the grid starts up successfully:
 - a. If you added new NTP servers, delete the old NTP server values.
 - b. If you added new DNS servers, delete the old DNS server values.
16. Download the new Recovery Package from the Grid Manager.
 - a. Select **MAINTENANCE > System > Recovery package**.
 - b. Enter the provisioning passphrase.

Related information

- [Add to or change subnet lists on Grid Network](#)
- [Shut down grid node](#)

Add interfaces to existing node

Linux: Add Admin or Client interfaces to an existing node

Use these steps to add an interface on the Admin Network or the Client Network to a Linux node after it has been installed.

If you did not configure `ADMIN_NETWORK_TARGET` or `CLIENT_NETWORK_TARGET` in the node configuration file on the Linux host during installation, use this procedure to add the interface. For more information about the node configuration file, see the instructions for your Linux operating system:

- [Install StorageGRID on Red Hat Enterprise Linux](#)
- [Install StorageGRID on Ubuntu or Debian](#)

You perform this procedure on the Linux server hosting the node that needs the new network assignment, not inside the node. This procedure only adds the interface to the node; a validation error occurs if you attempt to specify any other network parameters.

To provide addressing information, you must use the Change IP tool. See [Change node network configuration](#).

Steps

1. Log in to the Linux server hosting the node.

2. Edit the node configuration file: `/etc/storagegrid/nodes/node-name.conf`.



Don't specify any other network parameters, or a validation error will result.

a. Add an entry for the new network target. For example:

```
CLIENT_NETWORK_TARGET = bond0.3206
```

b. Optional: Add an entry for the MAC address. For example:

```
CLIENT_NETWORK_MAC = aa:57:61:07:ea:5c
```

3. Run the node validate command:

```
sudo storagegrid node validate node-name
```

4. Resolve all validation errors.

5. Run the node reload command:

```
sudo storagegrid node reload node-name
```

Linux: Add trunk or access interfaces to a node

You can add extra trunk or access interfaces to a Linux node after it has been installed. The interfaces you add are displayed on the VLAN interfaces page and the HA groups page.

Before you begin

- You have access to the instructions for installing StorageGRID on your Linux platform.
 - [Install StorageGRID on Red Hat Enterprise Linux](#)
 - [Install StorageGRID on Ubuntu or Debian](#)
- You have the `Passwords.txt` file.
- You have [specific access permissions](#).



Don't attempt to add interfaces to a node while a software upgrade, recovery procedure, or expansion procedure is active.

About this task

Use these steps to add one or more extra interfaces to a Linux node after the node has been installed. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use VLAN interfaces to segregate the traffic belonging to different applications or tenants. Or, you might want to add an access interface to use in a high availability (HA) group.

If you add a trunk interface, you must configure a VLAN interface in StorageGRID. If you add an access interface, you can add the interface directly to an HA group; you don't need to configure a VLAN interface.

The node is unavailable for a brief time when you add interfaces. You should perform this procedure on one node at a time.

Steps

1. Log in to the Linux server hosting the node.
2. Using a text editor such as vim or pico, edit the node configuration file:

```
/etc/storagegrid/nodes/node-name.conf
```

3. Add an entry to the file to specify the name and, optionally, the description of each extra interface you want to add to the node. Use this format.

```
INTERFACE_TARGET_nnnn=value
```

For *nnnn*, specify a unique number for each INTERFACE_TARGET entry you are adding.

For *value*, specify the name of the physical interface on the bare-metal host. Then, optionally, add a comma and provide a description of the interface, which is displayed on the VLAN interfaces page and the HA groups page.

For example:

```
INTERFACE_TARGET_0001=ens256, Trunk
```



Don't specify any other network parameters, or a validation error will result.

4. Run the following command to validate your changes to the node configuration file:

```
sudo storagegrid node validate node-name
```

Address any errors or warnings before proceeding to the next step.

5. Run the following command to update the node's configuration:

```
sudo storagegrid node reload node-name
```

After you finish

- If you added one or more trunk interfaces, go to [configure VLAN interfaces](#) to configure one or more VLAN interfaces for each new parent interface.
- If you added one or more access interfaces, go to [configure high availability groups](#) to add the new interfaces directly to HA groups.

VMware: Add trunk or access interfaces to a node

You can add a trunk or access interface to a VM node after the node has been installed. The interfaces you add are displayed on the VLAN interfaces page and the HA groups page.

Before you begin

- You have access to the instructions for [installing StorageGRID on your VMware platform](#).
- You have Admin Node and Gateway Node VMware virtual machines.
- You have a network subnet that is not being used as Grid, Admin, or Client Network.

- You have the `Passwords.txt` file.
- You have [specific access permissions](#).



Don't attempt to add interfaces to a node while a software upgrade, recovery procedure, or expansion procedure is active.

About this task

Use these steps to add one or more extra interfaces to a VMware node after the node has been installed. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use VLAN interfaces to segregate the traffic belonging to different applications or tenants. Or you might want to add an access interface to use in a high availability (HA) group.

If you add a trunk interface, you must configure a VLAN interface in StorageGRID. If you add an access interface, you can add the interface directly to an HA group; you don't need to configure a VLAN interface.

The node might be unavailable for a brief time when you add interfaces.

Steps

1. In vCenter, add a new network adapter (type VMXNET3) to an Admin Node and Gateway Node VM. Select **Connected** and **Connect At Power On** checkboxes.

Network adapter 4 *	CLIENT683_old_vlan	☒ Connected
Status	☒ Connect At Power On	
Adapter Type	VMXNET 3	
DirectPath I/O	☒ Enable	

2. Use SSH to log in to the Admin Node or Gateway Node.
3. Use `ip link show` to confirm the new network interface `ens256` is detected.

```
ip link show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode
DEFAULT group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1400 qdisc mq state UP
mode DEFAULT group default qlen 1000
    link/ether 00:50:56:a0:4e:5b brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN mode
DEFAULT group default qlen 1000
    link/ether 00:50:56:a0:fa:ce brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1400 qdisc mq state UP
mode DEFAULT group default qlen 1000
    link/ether 00:50:56:a0:d6:87 brd ff:ff:ff:ff:ff:ff
5: ens256: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq master
ens256vrf state UP mode DEFAULT group default qlen 1000
    link/ether 00:50:56:a0:ea:88 brd ff:ff:ff:ff:ff:ff
```

After you finish

- If you added one or more trunk interfaces, go to [configure VLAN interfaces](#) to configure one or more VLAN interfaces for each new parent interface.
- If you added one or more access interfaces, go to [configure high availability groups](#) to add the new interfaces directly to HA groups.

Configure DNS servers

You can add, update, and remove DNS servers, so that you can use fully qualified domain name (FQDN) hostnames rather than IP addresses.

To use fully qualified domain names (FQDNs) instead of IP addresses when specifying hostnames for external destinations, specify the IP address of each DNS server you will use. These entries are used for AutoSupport, alert emails, SNMP notifications, platform services endpoints, Cloud Storage Pools, and more.

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 the IP addresses of the DNS servers to configure.

About this task

To ensure proper operation, specify two or three DNS servers. If you specify more than three, it is possible that only three will be used because of known OS limitations on some platforms. If you have routing restrictions in your environment, you can [customize the DNS server list](#) for individual nodes (typically all nodes at a site) to use a different set of up to three DNS servers.

If possible, use DNS servers that each site can access locally to ensure that an islanded site can resolve the FQDNs for external destinations.

Add a DNS server

Follow these steps to add a DNS server.

Steps

1. Select **MAINTENANCE > Network > DNS servers**.
2. Select **Add another server** to add a DNS server.
3. Select **Save**.

Modify a DNS server

Follow these steps to modify a DNS server.

Steps

1. Select **MAINTENANCE > Network > DNS servers**.
2. Select the IP address of the server name you want to edit and make the necessary changes.
3. Select **Save**.

Delete a DNS server

Follow these steps to delete an IP address of a DNS server.

Steps

1. Select **MAINTENANCE > Network > DNS servers**.
2. Select the delete icon  next to the IP address.
3. Select **Save**.

Modify DNS configuration for single grid node

Rather than configure the DNS globally for the entire deployment, you can run a script to configure DNS differently for each grid node.

In general, you should use the **MAINTENANCE > Network > DNS servers** option on the Grid Manager to configure DNS servers. Only use the following script if you need to use different DNS servers for different grid nodes.

Steps

1. Log in to the primary Admin Node:
 - a. Enter the following command: `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. Add the SSH private key to the SSH agent. Enter: `ssh-add`
- f. Enter the SSH Access Password listed in the `Passwords.txt` file.

2. Log in to the node you want to update with a custom DNS configuration: `ssh node_IP_address`
3. Run the DNS setup script: `setup_resolv.rb`.

The script responds with the list of supported commands.

```
Tool to modify external name servers

available commands:
  add search <domain>
      add a specified domain to search list
      e.g.> add search netapp.com
  remove search <domain>
      remove a specified domain from list
      e.g.> remove search netapp.com
  add nameserver <ip>
      add a specified IP address to the name server list
      e.g.> add nameserver 192.0.2.65
  remove nameserver <ip>
      remove a specified IP address from list
      e.g.> remove nameserver 192.0.2.65
  remove nameserver all
      remove all nameservers from list
  save
      write configuration to disk and quit
  abort
      quit without saving changes
  help
      display this help message

Current list of name servers:
  192.0.2.64
Name servers inherited from global DNS configuration:
  192.0.2.126
  192.0.2.127
Current list of search entries:
  netapp.com

Enter command [`add search <domain>|remove search <domain>|add
nameserver <ip>`]
      [`remove nameserver <ip>|remove nameserver
all|save|abort|help`]
```

4. Add the IPv4 address of a server that provides domain name service for your network: `add <nameserver IP_address>`
5. Repeat the `add nameserver` command to add name servers.
6. Follow instructions as prompted for other commands.

7. Save your changes and exit the application: `save`
8. Close the command shell on the server: `exit`
9. For each grid node, repeat the steps from [logging into the node](#) through [closing the command shell](#).
10. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter: `ssh-add -D`

Manage NTP servers

You can add, update, or remove Network Time Protocol (NTP) servers to ensure that data is synchronized accurately between grid nodes in your StorageGRID system.

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 the provisioning passphrase.
- You have the IPv4 addresses of the NTP servers to configure.

How StorageGRID uses NTP

The StorageGRID system uses the Network Time Protocol (NTP) to synchronize time between all grid nodes in the grid.

At each site, at least two nodes in the StorageGRID system are assigned the primary NTP role. They synchronize to a suggested minimum of four, and a maximum of six, external time sources and with each other. Every node in the StorageGRID system that is not a primary NTP node acts as an NTP client and synchronizes with these primary NTP nodes.

The external NTP servers connect to the nodes to which you previously assigned primary NTP roles. For this reason, specifying at least two nodes with primary NTP roles is recommended.

NTP server guidelines

Follow these guidelines to protect against timing issues:

- The external NTP servers connect to the nodes to which you previously assigned primary NTP roles. For this reason, specifying at least two nodes with primary NTP roles is recommended.
- Make sure at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.
- The specified external NTP servers must use the NTP protocol. You must specify NTP server references of Stratum 3 or better to prevent issues with time drift.



When specifying the external NTP source for a production-level StorageGRID installation, don't use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, including StorageGRID. For details, see [Support boundary to configure the Windows Time service for high-accuracy environments](#).

Configure NTP servers

Follow these steps to add, update, or remove NTP servers.

Steps

1. Select **MAINTENANCE > Network > NTP servers**.
2. In the Servers section, add, update, or remove NTP server entries, as necessary.

You should include at least four NTP servers, and you can specify up to six servers.

3. Enter the provisioning passphrase for your StorageGRID system, then select **Save**.

The page is disabled until the configuration updates are complete.



If all of your NTP servers fail the connection test after you save the new NTP servers, don't proceed. Contact technical support.

Resolve NTP server issues

If you encounter problems with the stability or availability of the NTP servers originally specified during installation, you can update the list of external NTP sources that the StorageGRID system uses by adding additional servers, or updating or removing existing servers.

Restore network connectivity for isolated nodes

Under certain circumstances, one or more groups of nodes might not be able to contact the rest of the grid. For example, site- or grid-wide IP address changes can result in isolated nodes.

About this task

Node isolation is indicated by:

- Alerts, such as **Unable to communicate with node (Alerts > Current)**
- Connectivity-related diagnostics (**SUPPORT > Tools > Diagnostics**)

Some of the consequences of having isolated nodes include the following:

- If multiple nodes are isolated, you might not be able to sign in to or access the Grid Manager.
- If multiple nodes are isolated, the storage usage and quota values shown on the dashboard for the Tenant Manager might be out of date. The totals will be updated when network connectivity is restored.

To resolve the isolation issue, you run a command line utility on each isolated node or on one node in a group (all nodes in a subnet that does not contain the primary Admin Node) that is isolated from the grid. The utility provides the nodes with the IP address of a non-isolated node in the grid, which allows the isolated node or group of nodes to contact the entire grid again.



If the multicast domain name system (mDNS) is disabled in the networks, you might have to run the command line utility on each isolated node.

Steps

This procedure does not apply when only some services are offline or reporting communication errors.

1. Access the node and check `/var/local/log/dynip.log` for isolation messages.

For example:

```
[2018-01-09T19:11:00.545] UpdateQueue - WARNING -- Possible isolation,
no contact with other nodes.
If this warning persists, manual action might be required.
```

If you are using the VMware console, it will contain a message that the node might be isolated.

On Linux deployments, isolation messages would appear in `/var/log/storagegrid/node/<nodename>.log` files.

2. If the isolation messages are recurring and persistent, run the following command:

```
add_node_ip.py <address>
```

where `<address>` is the IP address of a remote node that is connected to the grid.

```
# /usr/sbin/add_node_ip.py 10.224.4.210

Retrieving local host information
Validating remote node at address 10.224.4.210
Sending node IP hint for 10.224.4.210 to local node
Local node found on remote node. Update complete.
```

3. Verify the following for each node that was previously isolated:
 - The node's services have started.
 - The status of the Dynamic IP service is "Running" after you run the `storagegrid-status` command.
 - On the Nodes page, the node no longer appears disconnected from the rest of the grid.



If running the `add_node_ip.py` command does not solve the problem, there could be other networking issues that need to be resolved.

Host and middleware procedures

Linux: Migrate grid node to new host

You can migrate one or more StorageGRID nodes from one Linux host (the *source host*) to another Linux host (the *target host*) to perform host maintenance without impacting the functionality or availability of your grid.

For example, you might want to migrate a node to perform OS patching and reboot.

Before you begin

- You planned your StorageGRID deployment to include migration support.
 - [Node container migration requirements for Red Hat Enterprise Linux](#)
 - [Node container migration requirements for Ubuntu or Debian](#)
- The target host is already prepared for StorageGRID use.
- Shared storage is used for all per-node storage volumes
- Network interfaces have consistent names across hosts.



In a production deployment, don't run more than one Storage Node on a single host. Using a dedicated host for each Storage Node provides an isolated failure domain.

Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same host. However, if you have multiple nodes of the same type (two Gateway Nodes, for example), don't install all instances on the same host.

Export node from source host

As a first step, shut down the grid node and export it from the source Linux host.

Run the following commands on the *source host*.

Steps

1. Obtain the status of all nodes currently running on the source host.

```
sudo storagegrid node status all
```

Example output:

```
Name Config-State Run-State
DC1-ADM1 Configured Running
DC1-ARC1 Configured Running
DC1-GW1 Configured Running
DC1-S1 Configured Running
DC1-S2 Configured Running
DC1-S3 Configured Running
```

2. Identify the name of the node you want to migrate, and stop it if its Run-State is Running.

```
sudo storagegrid node stop DC1-S3
```

Example output:

```
Stopping node DC1-S3
Waiting up to 630 seconds for node shutdown
```

3. Export the node from the source host.

```
sudo storagegrid node export DC1-S3
```

Example output:

```
Finished exporting node DC1-S3 to /dev/mapper/sgws-dc1-s3-var-local.  
Use 'storagegrid node import /dev/mapper/sgws-dc1-s3-var-local' if you  
want to import it again.
```

4. Make note of the `import` command suggested in the output.

You will run this command on the target host in the next step.

Import node on target host

After exporting the node from the source host, you import and validate the node on the target host. Validation confirms that the node has access to the same block storage and network interface devices as it had on the source host.

Run the following commands on the *target host*.

Steps

1. Import the node on the target host.

```
sudo storagegrid node import /dev/mapper/sgws-dc1-s3-var-local
```

Example output:

```
Finished importing node DC1-S3 from /dev/mapper/sgws-dc1-s3-var-local.  
You should run 'storagegrid node validate DC1-S3'
```

2. Validate the node configuration on the new host.

```
sudo storagegrid node validate DC1-S3
```

Example output:

```
Confirming existence of node DC1-S3... PASSED  
Checking configuration file /etc/storagegrid/nodes/DC1-S3.conf for node  
DC1-S3... PASSED  
Checking for duplication of unique values... PASSED
```

3. If any validation errors occur, address them before starting the migrated node.

For troubleshooting information, see the StorageGRID installation instructions for your Linux operating system.

- [Install StorageGRID on Red Hat Enterprise Linux](#)
- [Install StorageGRID on Ubuntu or Debian](#)

Start migrated node

After you validate the migrated node, you start the node by running a command on the *target host*.

Steps

1. Start the node on the new host.

```
sudo storagegrid node start DC1-S3
```

2. Sign in to the Grid Manager and verify that the status of the node is green with no alert.



Verifying that the status of the node is green ensures that the migrated node has fully restarted and rejoined the grid. If the status is not green, don't migrate any additional nodes so that you will not have more than one node out of service.

3. If you are unable to access the Grid Manager, wait for 10 minutes, then run the following command:

```
sudo storagegrid node status _node-name
```

Confirm that the migrated node has a Run-State of Running.

VMware: Configure virtual machine for automatic restart

If the virtual machine does not restart after VMware vSphere Hypervisor is restarted, you might need to configure the virtual machine for automatic restart.

You should perform this procedure if you notice that a virtual machine does not restart while you are recovering a grid node or performing another maintenance procedure.

Steps

1. In the VMware vSphere Client tree, select the virtual machine that is not started.
2. Right-click the virtual machine, and select **Power on**.
3. Configure VMware vSphere Hypervisor to restart the virtual machine automatically in future.

Recover or replace nodes

Warnings and considerations for grid node recovery

If a grid node fails, you must recover it as soon as possible. You must review all warnings and considerations for node recovery before you begin.



StorageGRID is a distributed system composed of multiple nodes working with each other. Don't use disk snapshots to restore grid nodes. Instead, refer to the recovery and maintenance procedures for each type of node.



If an entire StorageGRID site has failed, contact technical support. Technical support will work with you to develop and execute a site recovery plan that maximizes the amount of data that is recovered and meets your business objectives. See [How technical support recovers a site](#).

Some of the reasons for recovering a failed grid node as soon as possible include the following:

- A failed grid node can reduce the redundancy of system and object data, leaving you vulnerable to the risk of permanent data loss if another node fails.
- A failed grid node can impact the efficiency of day-to-day operations.
- A failed grid node can reduce your ability to monitor system operations.
- A failed grid node can cause a 500 internal server error if strict ILM rules are in place.
- If a grid node is not recovered promptly, recovery times might increase. For example, queues might develop that need to be cleared before recovery is complete.

Always follow the recovery procedure for the specific type of grid node you are recovering. Recovery procedures vary for primary or non-primary Admin Nodes, Gateway Nodes, appliance nodes, and Storage Nodes.

Preconditions for recovering grid nodes

All of the following conditions are assumed when recovering grid nodes:

- The failed physical or virtual hardware has been replaced and configured.
- The StorageGRID Appliance Installer version on the replacement appliance matches the software version of your StorageGRID system, as described in [Verify and upgrade StorageGRID Appliance Installer version](#).
- If you are recovering a grid node other than the primary Admin Node, there is connectivity between the grid node being recovered and the primary Admin Node.
- If you are recovering an appliance Storage Node, you must specify the same storage type as the original appliance (Combined, Metadata-only, or Data-only) during appliance installation. If you specify a different storage type, the recovery will fail and require reinstallation of the appliance with the correct storage type specified.

Order of node recovery if a server hosting more than one grid node fails

If a server that is hosting more than one grid node fails, you can recover the nodes in any order. However, if the failed server is hosting the primary Admin Node, you must recover that node first. Recovering the primary Admin Node first prevents other node recoveries from halting as they wait to contact the primary Admin Node.

IP addresses for recovered nodes

Don't attempt to recover a node using an IP address that is currently assigned to any other node. When you deploy the new node, use the failed node's current IP address or an unused IP address.

If you use a new IP address to deploy the new node and then recover the node, the new IP address will continue to be used for the recovered node. If you want to revert to the original IP address, use the Change IP tool after the recovery is complete.

Gather required materials for grid node recovery

Before performing maintenance procedures, you must ensure you have the necessary materials to recover a failed grid node.

Item	Notes
StorageGRID installation archive	If you need to recover a grid node, you need to download the StorageGRID installation files for your platform. Note: You don't need to download files if you are recovering failed storage volumes on a Storage Node.
Service laptop	The service laptop must have the following: • Network port • SSH client (for example, PuTTY) • Supported web browser
Recovery Package .zip file	Obtain a copy of the most recent Recovery Package .zip file: <code>sgws-recovery-package-id-revision.zip</code> The contents of the .zip file are updated each time the system is modified. You are directed to store the most recent version of the Recovery Package in a secure location after making such changes. Use the most recent copy to recover from grid failures. If the primary Admin Node is operating normally, you can download the Recovery Package from the Grid Manager. Select MAINTENANCE > System > Recovery package. If you can't access the Grid Manager, you can find encrypted copies of the Recovery Package on some Storage Nodes that contain the ADC service. On each Storage Node, examine this location for the Recovery Package: <code>/var/local/install/sgws-recovery-package-grid-id-revision.zip.gpg</code> Use the Recovery Package with the highest revision number.
Passwords.txt file	Contains the passwords required to access grid nodes on the command line. Included in the Recovery Package.

Item	Notes
Provisioning passphrase	The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the <code>Passwords.txt</code> file.
Current documentation for your platform	Go to the platform vendor's website for documentation. For the current supported versions of your platform, see the NetApp Interoperability Matrix Tool .

Download and extract StorageGRID installation files

Download the software and extract the files, unless you are [recovering failed storage volumes on a Storage Node](#).

You must use the version of StorageGRID that is currently running on the grid.

Steps

1. Determine which version of the software is currently installed. From the top of the Grid Manager, select the help icon and select **About**.
2. Go to the [NetApp Downloads page for StorageGRID](#).
3. Select the version of StorageGRID that is currently running on the grid.

StorageGRID software versions have this format: `11.x.y`.

4. Sign in with the username and password for your NetApp account.
5. Read the End User License Agreement, select the checkbox, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column of the download page, select the `.tgz` or `.zip` file for your platform.

The version shown in the installation archive file must match the version of the software that is currently installed.

Use the `.zip` file if you are running Windows.

Platform	Installation archive
Red Hat Enterprise Linux	<code>StorageGRID-Webscale-version-RPM-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-RPM-uniqueID.tgz</code>
Ubuntu or Debian or Appliances	<code>StorageGRID-Webscale-version-DEB-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-DEB-uniqueID.tgz</code>
VMware	<code>StorageGRID-Webscale-version-VMware-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-VMware-uniqueID.tgz</code>

7. Download and extract the archive file.

8. Follow the appropriate step for your platform to choose the files you need, based on your platform and which grid nodes you need to recover.

The paths listed in the step for each platform are relative to the top-level directory installed by the archive file.

9. If you are recovering a [Red Hat Enterprise Linux system](#), select the appropriate files.

Path and file name	Description
<code>./rpms/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./rpms/NLF000000.txt</code>	A free license that does not provide any support entitlement for the product.
<code>./rpms/StorageGRID-Webscale-Images-version-SHA.rpm</code>	RPM package for installing the StorageGRID node images on your RHEL hosts.
<code>./rpms/StorageGRID-Webscale-Service-version-SHA.rpm</code>	RPM package for installing the StorageGRID host service on your RHEL hosts.
Deployment scripting tool	Description
<code>./rpms/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.
<code>./rpms/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./rpms/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./rpms/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate integration.
<code>./rpms/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./rpms/extras/ansible</code>	Example Ansible role and playbook for configuring RHEL hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
<code>./rpms/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.

Path and file name	Description
<code>./rpms/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.
<code>./rpms/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

10. If you are recovering an [Ubuntu or Debian system](#), select the appropriate files.

Path and file name	Description
<code>./debs/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./debs/NLF000000.txt</code>	A non-production NetApp License File that you can use for testing and proof of concept deployments.
<code>./debs/storagegrid-webscale-images-version-SHA.deb</code>	DEB package for installing the StorageGRID node images on Ubuntu or Debian hosts.
<code>./debs/storagegrid-webscale-images-version-SHA.deb.md5</code>	MD5 checksum for the file <code>/debs/storagegrid-webscale-images-version-SHA.deb</code> .
<code>./debs/storagegrid-webscale-service-version-SHA.deb</code>	DEB package for installing the StorageGRID host service on Ubuntu or Debian hosts.
Deployment scripting tool	Description
<code>./debs/configure-storagegrid.py</code>	A Python script used to automate the configuration of a StorageGRID system.
<code>./debs/configure-sga.py</code>	A Python script used to automate the configuration of StorageGRID appliances.
<code>./debs/storagegrid-ssoauth.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled. You can also use this script for Ping Federate integration.
<code>./debs/configure-storagegrid.sample.json</code>	An example configuration file for use with the <code>configure-storagegrid.py</code> script.

Path and file name	Description
<code>./debs/configure-storagegrid.blank.json</code>	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
<code>./debs/extras/ansible</code>	Example Ansible role and playbook for configuring Ubuntu or Debian hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.
<code>./debs/storagegrid-ssoauth-azure.py</code>	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
<code>./debs/storagegrid-ssoauth-azure.js</code>	A helper script called by the companion <code>storagegrid-ssoauth-azure.py</code> Python script to perform SSO interactions with Azure.
<code>./debs/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

11. If you are recovering a [VMware system](#), select the appropriate files.

Path and file name	Description
<code>./vsphere/README</code>	A text file that describes all of the files contained in the StorageGRID download file.
<code>./vsphere/NLF000000.txt</code>	A free license that does not provide any support entitlement for the product.
<code>./vsphere/NetApp-SG-version-SHA.vmdk</code>	The virtual machine disk file that is used as a template for creating grid node virtual machines.
<code>./vsphere/vsphere-primary-admin.ovf</code> <code>./vsphere/vsphere-primary-admin.mf</code>	The Open Virtualization Format template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying the primary Admin Node.
<code>./vsphere/vsphere-non-primary-admin.ovf</code> <code>./vsphere/vsphere-non-primary-admin.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying non-primary Admin Nodes.

Path and file name	Description
./vsphere/vsphere-gateway.ovf ./vsphere/vsphere-gateway.mf	The template file (.ovf) and manifest file (.mf) for deploying Gateway Nodes.
./vsphere/vsphere-storage.ovf ./vsphere/vsphere-storage.mf	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
Deployment scripting tool	Description
./vsphere/deploy-vsphere-ovftool.sh	A Bash shell script used to automate the deployment of virtual grid nodes.
./vsphere/deploy-vsphere-ovftool-sample.ini	An example configuration file for use with the deploy-vsphere-ovftool.sh script.
./vsphere/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./vsphere/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./vsphere/storagegrid-ssoauth.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled. You can also use this script for Ping Federate integration.
./vsphere/configure-storagegrid.sample.json	An example configuration file for use with the configure-storagegrid.py script.
./vsphere/configure-storagegrid.blank.json	A blank configuration file for use with the configure-storagegrid.py script.
./vsphere/storagegrid-ssoauth-azure.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on (SSO) is enabled using Active Directory or Ping Federate.
./vsphere/storagegrid-ssoauth-azure.js	A helper script called by the companion storagegrid-ssoauth-azure.py Python script to perform SSO interactions with Azure.

Path and file name	Description
<code>./vsphere/extras/api-schemas</code>	API schemas for StorageGRID. Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you don't have a non-production StorageGRID environment for upgrade compatibility testing.

12. If you are recovering a StorageGRID appliance-based system, select the appropriate files.

Path and file name	Description
<code>./debs/storagegrid-webscale-images-version-SHA.deb</code>	DEB package for installing the StorageGRID node images on your appliances.
<code>./debs/storagegrid-webscale-images-version-SHA.deb.md5</code>	MD5 checksum for the file <code>/debs/storagegridwebscale-images-version-SHA.deb</code> .



For appliance installation, these files are only required if you need to avoid network traffic. The appliance can download the required files from the primary Admin Node.

Select node recovery procedure

You must select the correct recovery procedure for the type of node that has failed.

Grid node	Recovery procedure
More than one Storage Node	Contact technical support. If more than one Storage Node has failed, technical support must assist with recovery to prevent database inconsistencies that could lead to data loss. A site recovery procedure might be required. How technical support recovers a site
A single Storage Node	The Storage Node recovery procedure depends on the type and duration of the failure. Recover from Storage Node failures
Admin Node	The Admin Node procedure depends on whether you need to recover the primary Admin Node or a non-primary Admin Node. Recover from Admin Node failures
Gateway Node	Recover from Gateway Node failures

Grid node	Recovery procedure
Archive Node	Recover from Archive Node failures (StorageGRID 11.8 doc site)



If a server that is hosting more than one grid node fails, you can recover the nodes in any order. However, if the failed server is hosting the primary Admin Node, you must recover that node first. Recovering the primary Admin Node first prevents other node recoveries from halting as they wait to contact the primary Admin Node.

Recover from Storage Node failures

Recover from Storage Node failures

The procedure for recovering a failed Storage Node depends on the type of failure and the type of Storage Node that has failed.

Use this table to select the recovery procedure for a failed Storage Node.

Issue	Action	Notes
- More than one Storage Node has failed. - A second Storage Node has failed less than 15 days after a Storage Node failure or recovery. This includes the case where a Storage Node fails while recovery of another Storage Node is still in progress.	Contact technical support.	Recovering more than one Storage Node (or more than one Storage Node within 15 days) might affect the integrity of the Cassandra database, which can cause data loss. Technical support can determine when it is safe to begin recovery of a second Storage Node. Note: If more than one Storage Node that contains the ADC service fails at a site, you lose any pending platform service requests for that site.
More than one Storage Node at a site has failed or an entire site has failed.	Contact technical support. It might be necessary to perform a site recovery procedure.	Technical support will assess your situation and develop a recovery plan. See How technical support recovers a site .
An appliance Storage Node has failed.	Recover appliance Storage Node	The recovery procedure for appliance Storage Nodes is the same for all failures.
One or more storage volumes have failed, but the system drive is intact	Recover from storage volume failure where system drive is intact	This procedure is used for software-based Storage Nodes.

Issue	Action	Notes
The system drive has failed.	Recover from system drive failure	The node replacement procedure depends on the deployment platform and on whether any storage volumes have also failed.



Some StorageGRID recovery procedures use Reaper to handle Cassandra repairs. Repairs occur automatically as soon as the related or required services have started. You might notice script output that mentions "reaper" or "Cassandra repair." If you see an error message indicating the repair has failed, run the command indicated in the error message.

Recover appliance Storage Node

Warnings for recovering appliance Storage Nodes

The procedure for recovering a failed StorageGRID appliance Storage Node is the same whether you are recovering from the loss of the system drive or from the loss of storage volumes only.



If more than one Storage Node has failed (or is offline), contact technical support. Don't perform the following recovery procedure. Data loss could occur.



If this is the second Storage Node failure in less than 15 days after a Storage Node failure or recovery, contact technical support. Rebuilding Cassandra on two or more Storage Nodes within 15 days can result in data loss.



If more than one Storage Node at a site has failed, a site recovery procedure might be required. See [How technical support recovers a site](#).



If ILM rules are configured to store only one replicated copy and the copy exists on a storage volume that has failed, you will not be able to recover the object.



For hardware maintenance procedures, such as instructions for replacing a controller or reinstalling SANtricity OS, see the [maintenance instructions for your storage appliance](#).

Prepare appliance Storage Node for reinstallation

When recovering an appliance Storage Node, you must first prepare the appliance for reinstallation of StorageGRID software.

Steps

1. Log in to the failed Storage Node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

2. Prepare the appliance Storage Node for the installation of StorageGRID software. `sgareinstall`

3. When prompted to continue, enter: `y`

The appliance reboots, and your SSH session ends. It usually takes about 5 minutes for the StorageGRID Appliance Installer to become available, although in some cases you might need to wait up to 30 minutes.



Don't attempt to accelerate the reboot by cycling power or otherwise resetting the appliance. You might interrupt automatic BIOS, BMC, or other firmware upgrades.

The StorageGRID appliance Storage Node is reset, and data on the Storage Node is no longer accessible. IP addresses configured during the original installation process should remain intact; however, it is recommended that you confirm this when the procedure completes.

After executing the `sgareinstall` command, all StorageGRID-provisioned accounts, passwords, and SSH keys are removed, and new host keys are generated.

Start StorageGRID appliance installation

To install StorageGRID on an appliance Storage Node, you use the StorageGRID Appliance Installer, which is included on the appliance.

Before you begin

- The appliance has been installed in a rack, connected to your networks, and powered on.
- Network links and IP addresses have been configured for the appliance using the StorageGRID Appliance Installer.
- You know the IP address of the primary Admin Node for the StorageGRID grid.
- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.
- You have completed these prerequisite tasks by following the installation instructions for your storage appliance. See [Quick start for hardware installation](#).
- You are using a [supported web browser](#).
- You know one of the IP addresses assigned to the compute controller in the appliance. You can use the IP address for the Admin Network (management port 1 on the controller), the Grid Network, or the Client Network.

About this task

To install StorageGRID on an appliance Storage Node:

- You specify or confirm the IP address of the primary Admin Node and the hostname (system name) of the node.
- You start the installation and wait as volumes are configured and the software is installed.



When recovering an appliance Storage Node, reinstall it with the same storage type as the original appliance (Combined, Metadata-only, or Data-only). If you specify a different storage type, the recovery will fail and require reinstallation of the appliance with the correct storage type specified.

- Partway through the process, the installation pauses. To resume the installation, you must sign into the Grid Manager and configure the pending Storage Node as a replacement for the failed node.
- After you have configured the node, the appliance installation process completes, and the appliance is rebooted.

Steps

1. Open a browser and enter one of the IP addresses for the compute controller in the appliance.

`https://Controller_IP:8443`

The StorageGRID Appliance Installer Home page appears.

2. In the Primary Admin Node connection section, determine whether you need to specify the IP address for the primary Admin Node.

The StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

3. If this IP address is not shown or you need to change it, specify the address:

Option	Steps
Manual IP entry	a. Clear the Enable Admin Node discovery checkbox. b. Enter the IP address manually. c. Click Save. d. Wait while the connection state for the new IP address becomes "ready."
Automatic discovery of all connected primary Admin Nodes	a. Select the Enable Admin Node discovery checkbox. b. From the list of discovered IP addresses, select the primary Admin Node for the grid where this appliance Storage Node will be deployed. c. Click Save. d. Wait while the connection state for the new IP address becomes "ready."

4. In the **Node Name** field, enter the same hostname (system name) that was used for the node you are recovering, and click **Save**.
5. In the Installation section, confirm that the current state is "Ready to start installation of *node name* into grid with Primary Admin Node ``admin_ip``" and that the **Start Installation** button is enabled.

If the **Start Installation** button is not enabled, you might need to change the network configuration or port settings. For instructions, see the maintenance instructions for your appliance.

6. From the StorageGRID Appliance Installer home page, click **Start Installation**.

NetApp® StorageGRID® Appliance Installer

Home

Configure Networking ▾

Configure Hardware ▾

Monitor Installation

Advanced ▾

Home

The installation is ready to be started. Review the settings below, and then click Start Installation.

Primary Admin Node connection

Enable Admin Node discovery

☐

Primary Admin Node IP

172.16.4.210

Connection state

Connection to 172.16.4.210 ready

Cancel

Save

Node name

Node name

NetApp-SGA

Cancel

Save

Installation

Current state

Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

Start Installation

The Current state changes to "Installation is in progress," and the Monitor Installation page is displayed.



If you need to access the Monitor Installation page manually, click **Monitor Installation** from the menu bar. See [Monitor appliance installation](#).

Monitor StorageGRID appliance installation

The StorageGRID Appliance Installer provides status until installation is complete. When the software installation is complete, the appliance is rebooted.

Steps

- 1. To monitor the installation progress, click **Monitor Installation** from the menu bar.

The Monitor Installation page shows the installation progress.

Monitor Installation

1. Configure storage			Running
Step	Progress	Status	
Connect to storage controller		Complete	
Clear existing configuration		Complete	
Configure volumes		Creating volume StorageGRID-obj-00	
Configure host settings		Pending	
2. Install OS			Pending
3. Install StorageGRID			Pending
4. Finalize installation			Pending

The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.



The installer ensures that tasks completed in a previous install aren't re-run. If you are re-running an installation, any tasks that don't need to be re-run are shown with a green status bar and a status of "Skipped."

- 2. Review the progress of first two installation stages.

- **1. Configure storage**

During this stage, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity OS to configure volumes, and configures host settings.

- **2. Install OS**

During this stage, the installer copies the base operating system image for StorageGRID to the appliance.

- 3. Continue monitoring the installation progress until the **Install StorageGRID** stage pauses and a message appears on the embedded console prompting you to approve this node on the Admin Node using the Grid Manager.

Monitor Installation

1. Configure storage	Complete
2. Install OS	Complete
3. Install StorageGRID	Running
4. Finalize installation	Pending

Connected (unencrypted) to: QEMU

```

/platform.type: Device or resource busy
[2017-07-31T22:09:12.362566] INFO -- [INSG] NOTICE: seeding /var/local with c
ontainer data
[2017-07-31T22:09:12.366205] INFO -- [INSG] Fixing permissions
[2017-07-31T22:09:12.369633] INFO -- [INSG] Enabling syslog
[2017-07-31T22:09:12.511533] INFO -- [INSG] Stopping system logging: syslog-n
g.
[2017-07-31T22:09:12.570096] INFO -- [INSG] Starting system logging: syslog-n
g.
[2017-07-31T22:09:12.576360] INFO -- [INSG] Beginning negotiation for downloa
d of node configuration
[2017-07-31T22:09:12.581363] INFO -- [INSG]
[2017-07-31T22:09:12.585066] INFO -- [INSG]
[2017-07-31T22:09:12.588314] INFO -- [INSG]
[2017-07-31T22:09:12.591851] INFO -- [INSG]
[2017-07-31T22:09:12.594886] INFO -- [INSG]
[2017-07-31T22:09:12.598360] INFO -- [INSG]
[2017-07-31T22:09:12.601324] INFO -- [INSG]
[2017-07-31T22:09:12.604759] INFO -- [INSG]
[2017-07-31T22:09:12.607800] INFO -- [INSG]
[2017-07-31T22:09:12.610985] INFO -- [INSG]
[2017-07-31T22:09:12.614597] INFO -- [INSG]
[2017-07-31T22:09:12.618282] INFO -- [INSG] Please approve this node on the A
dmin Node GMI to proceed...

```

- Go to [Select Start Recovery](#) to configure appliance Storage Node.

Select Start Recovery to configure appliance Storage Node

You must select Start Recovery in the Grid Manager to configure an appliance Storage Node as a replacement for the failed node.

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 the provisioning passphrase.

- You have deployed a recovery appliance Storage Node.
- You have the start date of any repair jobs for erasure-coded data.
- You have verified that the Storage Node has not been rebuilt within the last 15 days.

Steps

1. From the Grid Manager, select **MAINTENANCE > Tasks > Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

Nodes appear in the list after they fail, but you can't select a node until it has been reinstalled and is ready for recovery.

3. Enter the **Provisioning Passphrase**.
4. Click **Start Recovery**.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

					Search 
	Name	IPv4 Address	State	Recoverable	
☒	104-217-S1	10.96.104.217	Unknown		

Passphrase

Provisioning Passphrase

Start Recovery

5. Monitor the progress of the recovery in the Recovering Grid Node table.

When the grid node reaches the "Waiting for Manual Steps" stage, go to the next topic and perform the manual steps to remount and reformat appliance storage volumes.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Recovering Grid Node

Name	Start Time	Progress	Stage
dc2-s3	2016-09-12 16:12:40 PDT		Waiting For Manual Steps

Reset



At any point during the recovery, you can click **Reset** to start a new recovery. A dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

Info

Reset Recovery

Resetting the recovery procedure leaves the deployed grid node in an indeterminate state. To retry a recovery after resetting the procedure, you must restore the node to a pre-installed state:

- For VMware nodes, delete the deployed VM and then redeploy it.
- For StorageGRID appliance nodes, run "sgareinstall" on the node.
- For Linux nodes, run "storagegrid node force-recovery *node-name*" on the Linux host.

Do you want to reset recovery?

Cancel

OK

If you want to retry the recovery after resetting the procedure, you must restore the appliance node to a pre-installed state by running `sgareinstall` on the node.

Remount and reformat appliance storage volumes (manual steps)

You must manually run two scripts to remount preserved storage volumes and reformat any failed storage volumes. The first script remounts volumes that are properly formatted as StorageGRID storage volumes. The second script reformats any unmounted volumes, rebuilds the Cassandra database, if needed, and starts services.

Before you begin

- You have already replaced the hardware for any failed storage volumes that you know require replacement.

Running the `sn-remount-volumes` script might help you identify additional failed storage volumes.

- You have checked that a Storage Node decommissioning is not in progress, or you have paused the node decommission procedure. (In the Grid Manager, select **MAINTENANCE > Tasks > Decommission.**)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **MAINTENANCE > Tasks > Expansion.**)



Contact technical support if more than one Storage Node is offline or if a Storage Node in this grid has been rebuilt in the last 15 days. Don't run the `sn-recovery-postinstall.sh` script. Rebuilding Cassandra on two or more Storage Nodes within 15 days of each other might result in data loss.

About this task

To complete this procedure, you perform these high-level tasks:

- Log in to the recovered Storage Node.
- Run the `sn-remount-volumes` script to remount properly formatted storage volumes. When this script runs, it does the following:

- Mounts and unmounts each storage volume to replay the XFS journal.
- Performs an XFS file consistency check.
- If the file system is consistent, determines if the storage volume is a properly formatted StorageGRID storage volume.
- If the storage volume is properly formatted, remounts the storage volume. Any existing data on the volume remains intact.
- Review the script output and resolve any issues.
- Run the `sn-recovery-postinstall.sh` script. When this script runs, it does the following.



Don't reboot a Storage Node during recovery before running `sn-recovery-postinstall.sh` (step 4) to reformat the failed storage volumes and restore object metadata. Rebooting the Storage Node before `sn-recovery-postinstall.sh` completes causes errors for services that attempt to start and causes StorageGRID appliance nodes to exit maintenance mode.

- Reformats any storage volumes that the `sn-remount-volumes` script could not mount or that were found to be improperly formatted.



If a storage volume is reformatted, any data on that volume is lost. You must perform an additional procedure to restore object data from other locations in the grid, assuming that ILM rules were configured to store more than one object copy.

- Rebuilds the Cassandra database on the node, if needed.
- Starts the services on the Storage Node.

Steps

1. Log in to the recovered Storage Node:

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

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

2. Run the first script to remount any properly formatted storage volumes.



If all storage volumes are new and need to be formatted, or if all storage volumes have failed, you can skip this step and run the second script to reformat all unmounted storage volumes.

- Run the script: `sn-remount-volumes`

This script might take hours to run on storage volumes that contain data.

- As the script runs, review the output and answer any prompts.



As required, you can use the `tail -f` command to monitor the contents of the script's log file (`/var/local/log/sn-remount-volumes.log`). The log file contains more detailed information than the command line output.

```
root@SG:~ # sn-remount-volumes
The configured LDR noid is 12632740

===== Device /dev/sdb =====
Mount and unmount device /dev/sdb and checking file system
consistency:
The device is consistent.
Check rangedb structure on device /dev/sdb:
Mount device /dev/sdb to /tmp/sdb-654321 with rangedb mount options
This device has all rangedb directories.
Found LDR node id 12632740, volume number 0 in the volID file
Attempting to remount /dev/sdb
Device /dev/sdb remounted successfully

===== Device /dev/sdc =====
Mount and unmount device /dev/sdc and checking file system
consistency:
Error: File system consistency check retry failed on device /dev/sdc.
You can see the diagnosis information in the /var/local/log/sn-
remount-volumes.log.

This volume could be new or damaged. If you run sn-recovery-
postinstall.sh, this volume and any data on this volume will be
deleted. If you only had two copies of object data, you will
temporarily have only a single copy.
StorageGRID will attempt to restore data redundancy by making
additional replicated copies or EC fragments, according to the rules
in the active ILM policies.

Don't continue to the next step if you believe that the data
remaining on this volume can't be rebuilt from elsewhere in the grid
(for example, if your ILM policy uses a rule that makes only one copy
or if volumes have failed on multiple nodes). Instead, contact
support to determine how to recover your data.

===== Device /dev/sdd =====
Mount and unmount device /dev/sdd and checking file system
consistency:
Failed to mount device /dev/sdd
This device could be an uninitialized disk or has corrupted
superblock.
File system check might take a long time. Do you want to continue? (y
```

```
or n) [y/N]? y
```

```
Error: File system consistency check retry failed on device /dev/sdd.  
You can see the diagnosis information in the /var/local/log/sn-  
remount-volumes.log.
```

```
This volume could be new or damaged. If you run sn-recovery-  
postinstall.sh, this volume and any data on this volume will be  
deleted. If you only had two copies of object data, you will  
temporarily have only a single copy.
```

```
StorageGRID will attempt to restore data redundancy by making  
additional replicated copies or EC fragments, according to the rules  
in the active ILM policies.
```

```
Don't continue to the next step if you believe that the data  
remaining on this volume can't be rebuilt from elsewhere in the grid  
(for example, if your ILM policy uses a rule that makes only one copy  
or if volumes have failed on multiple nodes). Instead, contact  
support to determine how to recover your data.
```

```
===== Device /dev/sde =====
```

```
Mount and unmount device /dev/sde and checking file system  
consistency:
```

```
The device is consistent.
```

```
Check rangedb structure on device /dev/sde:
```

```
Mount device /dev/sde to /tmp/sde-654321 with rangedb mount options
```

```
This device has all rangedb directories.
```

```
Found LDR node id 12000078, volume number 9 in the volID file
```

```
Error: This volume does not belong to this node. Fix the attached  
volume and re-run this script.
```

In the example output, one storage volume was remounted successfully and three storage volumes had errors.

- /dev/sdb passed the XFS file system consistency check and had a valid volume structure, so it was remounted successfully. Data on devices that are remounted by the script is preserved.
- /dev/sdc failed the XFS file system consistency check because the storage volume was new or corrupt.
- /dev/sdd could not be mounted because the disk was not initialized or the disk's superblock was corrupted. When the script can't mount a storage volume, it asks if you want to run the file system consistency check.
 - If the storage volume is attached to a new disk, answer **N** to the prompt. You don't need check the file system on a new disk.
 - If the storage volume is attached to an existing disk, answer **Y** to the prompt. You can use the results of the file system check to determine the source of the corruption. The results are saved in the /var/local/log/sn-remount-volumes.log log file.

- `/dev/sde` passed the XFS file system consistency check and had a valid volume structure; however, the LDR node ID in the `volID` file did not match the ID for this Storage Node (the configured LDR noid displayed at the top). This message indicates that this volume belongs to another Storage Node.

3. Review the script output and resolve any issues.



If a storage volume failed the XFS file system consistency check or could not be mounted, carefully review the error messages in the output. You must understand the implications of running the `sn-recovery-postinstall.sh` script on these volumes.

- a. Check to make sure that the results include an entry for all of the volumes you expected. If any volumes aren't listed, rerun the script.
- b. Review the messages for all mounted devices. Make sure there are no errors indicating that a storage volume does not belong to this Storage Node.

In the example, the output for `/dev/sde` includes the following error message:

```
Error: This volume does not belong to this node. Fix the attached
volume and re-run this script.
```



If a storage volume is reported as belonging to another Storage Node, contact technical support. If you run the `sn-recovery-postinstall.sh` script, the storage volume will be reformatted, which might cause data loss.

- c. If any storage devices could not be mounted, make a note of the device name, and repair or replace the device.



You must repair or replace any storage devices that could not be mounted.

You will use the device name to look up the volume ID, which is required input when you run the `repair-data` script to restore object data to the volume (the next procedure).

- d. After repairing or replacing all unmountable devices, run the `sn-remount-volumes` script again to confirm that all storage volumes that can be remounted have been remounted.



If a storage volume can't be mounted or is improperly formatted, and you continue to the next step, the volume and any data on the volume will be deleted. If you had two copies of object data, you will have only a single copy until you complete the next procedure (restoring object data).



Don't run the `sn-recovery-postinstall.sh` script if you believe that the data remaining on a failed storage volume can't be rebuilt from elsewhere in the grid (for example, if your ILM policy uses a rule that makes only one copy or if volumes have failed on multiple nodes). Instead, contact technical support to determine how to recover your data.

4. Run the `sn-recovery-postinstall.sh` script: `sn-recovery-postinstall.sh`

This script reformats any storage volumes that could not be mounted or that were found to be improperly

formatted; rebuilds the Cassandra database on the node, if needed; and starts the services on the Storage Node.

Be aware of the following:

- The script might take hours to run.
- In general, you should leave the SSH session alone while the script is running.
- Don't press **Ctrl+C** while the SSH session is active.
- The script will run in the background if a network disruption occurs and terminates the SSH session, but you can view the progress from the Recovery page.
- If the Storage Node uses the RSM service, the script might appear to stall for 5 minutes as node services are restarted. This 5-minute delay is expected whenever the RSM service boots for the first time.



The RSM service is present on Storage Nodes that include the ADC service.



Some StorageGRID recovery procedures use Reaper to handle Cassandra repairs. Repairs occur automatically as soon as the related or required services have started. You might notice script output that mentions "reaper" or "Cassandra repair." If you see an error message indicating the repair has failed, run the command indicated in the error message.

5. As the `sn-recovery-postinstall.sh` script runs, monitor the Recovery page in the Grid Manager.

The Progress bar and the Stage column on the Recovery page provide a high-level status of the `sn-recovery-postinstall.sh` script.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

Search				
Name	IPv4 Address	State	Recoverable	
No results found.				

Recovering Grid Node

Name	Start Time	Progress	Stage
DC1-S3	2016-06-02 14:03:35 PDT		Recovering Cassandra

6. After the `sn-recovery-postinstall.sh` script has started services on the node, you can restore object data to any storage volumes that were formatted by the script.

The script asks if you want to use the Grid Manager volume restoration process.

- In most cases, you should [restore object data using Grid Manager](#). Answer `y` to use the Grid Manager.
- In rare cases, such as when instructed by technical support, or when you know that the replacement node has fewer volumes available for object storage than the original node, you must [restore object data manually](#) using the `repair-data` script. If one of these cases applies, answer `n`.



If you answer `n` to using the Grid Manager volume restoration process (restore object data manually):

- You aren't able to restore object data using Grid Manager.
- You can monitor the progress of manual restoration jobs using Grid Manager.

After making your selection, the script completes and the next steps to recover object data are shown. After reviewing these steps, press any key to return to the command line.

Restore object data to storage volume for appliance

After recovering storage volumes for the appliance Storage Node, you can restore the replicated or erasure-coded object data that was lost when the Storage Node failed.

Which procedure should I use?

Whenever possible, restore object data using the **Volume restoration** page in the Grid Manager.

- If the volumes are listed at **MAINTENANCE > Volume restoration > Nodes to restore**, restore object data using the [Volume restoration page in the Grid Manager](#).
- If the volumes aren't listed at **MAINTENANCE > Volume restoration > Nodes to restore**, follow the steps below for using the `repair-data` script to restore object data.

If the recovered Storage Node contains fewer volumes than the node it is replacing, you must use the `repair-data` script.



The `repair-data` script is deprecated and will be removed in a future release. When possible, use the [Volume restoration procedure in the Grid Manager](#).

Use the `repair-data` script to restore object data

Before you begin

- You have confirmed that the recovered Storage Node has a Connection State of **Connected**  on the **NODES > Overview** tab in the Grid Manager.

About this task

Object data can be restored from other Storage Nodes or a Cloud Storage Pool, assuming that the grid's ILM rules were configured such that object copies are available.

Note the following:

- If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.
- If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.

About the `repair-data` script

To restore object data, you run the `repair-data` script. This script begins the process of restoring object data and works with ILM scanning to ensure that ILM rules are met.

Select **Replicated data** or **Erasure-coded (EC) data** below to learn the different options for the `repair-data` script, based on whether you are restoring replicated data or erasure-coded data. If you need to restore both types of data, you must run both sets of commands.



For more information about the `repair-data` script, enter `repair-data --help` from the command line of the primary Admin Node.



The `repair-data` script is deprecated and will be removed in a future release. When possible, use the [Volume restoration procedure in the Grid Manager](#).

Replicated data

Two commands are available for restoring replicated data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-replicated-node-repair
```

```
repair-data start-replicated-volume-repair
```

You can track repairs of replicated data with this command:

```
repair-data show-replicated-repair-status
```

Erasure-coded (EC) data

Two commands are available for restoring erasure-coded data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-ec-node-repair
```

```
repair-data start-ec-volume-repair
```

You can track repairs of erasure-coded data with this command:

```
repair-data show-ec-repair-status
```



Repairs of erasure-coded data can begin while some Storage Nodes are offline. However, if all erasure-coded data can't be accounted for, the repair can't be completed. Repair will complete after all nodes are available.



The EC repair job temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the repair is complete. If there is not enough storage for the reservation, the EC repair job will fail. Storage reservations are released when the EC repair job completes, whether the job failed or succeeded.

Find hostname for Storage Node

1. Log in to the primary Admin Node:
 - a. Enter the following command: `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 `#`.

2. Use the `/etc/hosts` file to find the hostname of the Storage Node for the restored storage volumes. To see a list of all nodes in the grid, enter the following: `cat /etc/hosts`.

Repair data if all volumes have failed

If all storage volumes have failed, repair the entire node. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If only some volumes have failed, go to [Repair data if only some volumes have failed](#).



You can't run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

Replicated data

If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the replicated data on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-node-repair --nodes SG-DC-SN3
```



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See [Investigate lost objects](#).

Erasure-coded (EC) data

If your grid contains erasure-coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the erasure-coded data on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-node-repair --nodes SG-DC-SN3
```

The operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.

Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Repair data if only some volumes have failed

If only some of the volumes have failed, repair the affected volumes. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If all volumes have failed, go to [Repair data if all volumes have failed](#).

Enter the volume IDs in hexadecimal. For example, `0000` is the first volume and `000F` is the sixteenth volume. You can specify one volume, a range of volumes, or multiple volumes that aren't in a sequence.

All the volumes must be on the same Storage Node. If you need to restore volumes for more than one Storage Node, contact technical support.

Replicated data

If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores replicated data to volume 0002 on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes 0002
```

Range of volumes: This command restores replicated data to all volumes in the range 0003 to 0009 on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volume-range 0003,0009
```

Multiple volumes not in a sequence: This command restores replicated data to volumes 0001, 0005, and 0008 on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes 0001,0005,0008
```



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. Note the alert description and recommended actions to determine the cause of the loss and if recovery is possible.

Erasure-coded (EC) data

If your grid contains erasure-coded data, use the `start-ec-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores erasure-coded data to volume 0007 on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 0007
```

Range of volumes: This command restores erasure-coded data to all volumes in the range 0004 to 0006 on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volume-range 0004,0006
```

Multiple volumes not in a sequence: This command restores erasure-coded data to volumes 000A, 000C, and 000E on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 000A,000C,000E
```

The `repair-data` operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.



Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Monitor repairs

Monitor the status of the repair jobs, based on whether you use **replicated data**, **erasure-coded (EC) data**, or both.

You can also monitor the status of volume restoration jobs in process and view a history of restoration jobs completed in [Grid Manager](#).

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the `repair-data` command.

```
repair-data show-replicated-repair-status
```

- To determine if repairs are complete:
 1. Select **NODES > Storage Node being repaired > ILM**.
 2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
 1. Select **SUPPORT > Tools > Grid topology**.
 2. Select **grid > Storage Node being repaired > LDR > Data Store**.
 3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.



Cassandra inconsistencies might be present, and failed repairs aren't tracked.

- **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.



High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

- **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure-coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
 - Select **SUPPORT > Tools > Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.

- Use this command to see the status of a specific `repair-data` operation:

```
repair-data show-ec-repair-status --repair-id repair ID
```

- Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

This command retries a failed node repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-node-repair --repair-id 6949309319275667690
```

This command retries a failed volume repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-volume-repair --repair-id 6949309319275667690
```

Check storage state after recovering appliance Storage Node

After recovering an appliance Storage Node, you must verify that the desired state of the appliance Storage Node is set to online and ensure that the state will be online by default whenever the Storage Node server is restarted.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- The Storage Node has been recovered, and data recovery is complete.

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Check the values of **Recovered Storage Node > LDR > Storage > Storage State — Desired** and **Storage State — Current**.

The value of both attributes should be Online.

3. If the Storage State — Desired is set to Read-only, complete the following steps:
 - a. Click the **Configuration** tab.
 - b. From the **Storage State — Desired** drop-down list, select **Online**.
 - c. Click **Apply Changes**.
 - d. Click the **Overview** tab and confirm that the values of **Storage State — Desired** and **Storage State — Current** are updated to Online.

Recover from storage volume failure where system drive is intact

Recover from storage volume failure where system drive is intact

You must complete a series of tasks to recover a software-based Storage Node where one or more storage volumes on the Storage Node have failed, but the system drive is intact. If only storage volumes have failed, the Storage Node is still available to the StorageGRID system.



This recovery procedure applies to software-based Storage Nodes only. If storage volumes have failed on an appliance Storage Node, use the appliance procedure instead: [Recover appliance Storage Node](#).

This recovery procedure includes the following tasks:

- [Review warnings for storage volume recovery](#)
- [Identify and unmount failed storage volumes](#)
- [Recover the volumes and rebuild the Cassandra database](#)
- [Restore object data](#)
- [Check the storage state](#)

Warnings for storage volume recovery

Before recovering failed storage volumes for a Storage Node, review the following warnings.

The storage volumes (or rangedbs) in a Storage Node are identified by a hexadecimal number, which is known as the volume ID. For example, 0000 is the first volume and 000F is the sixteenth volume. The first object store (volume 0) on each Storage Node uses up to 4 TB of space for object metadata and Cassandra database operations; any remaining space on that volume is used for object data. All other storage volumes are used exclusively for object data.

If volume 0 fails and needs to be recovered, the Cassandra database might be rebuilt as part of the volume recovery procedure. Cassandra might also be rebuilt in the following circumstances:

- A Storage Node is brought back online after having been offline for more than 15 days.
- The system drive and one or more storage volumes fails and is recovered.

When Cassandra is rebuilt, the system uses information from other Storage Nodes. If too many Storage Nodes are offline, some Cassandra data might not be available. If Cassandra has been rebuilt recently, Cassandra data might not yet be consistent across the grid. Data loss can occur if Cassandra is rebuilt when too many Storage Nodes are offline or if two or more Storage Nodes are rebuilt within 15 days of each other.



If more than one Storage Node has failed (or is offline), contact technical support. Don't perform the following recovery procedure. Data loss could occur.



If this is the second Storage Node failure in less than 15 days after a Storage Node failure or recovery, contact technical support. Rebuilding Cassandra on two or more Storage Nodes within 15 days can result in data loss.



If more than one Storage Node at a site has failed, a site recovery procedure might be required. See [How technical support recovers a site](#).



If ILM rules are configured to store only one replicated copy and the copy exists on a storage volume that has failed, you will not be able to recover the object.

Related information

[Warnings and considerations for grid node recovery](#)

Identify and unmount failed storage volumes

When recovering a Storage Node with failed storage volumes, you must identify and unmount the failed volumes. You must verify that only the failed storage volumes are reformatted as part of the recovery procedure.

Before you begin

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

About this task

You should recover failed storage volumes as soon as possible.

The first step of the recovery process is to detect volumes that have become detached, need to be unmounted, or have I/O errors. If failed volumes are still attached but have a randomly corrupted file system, the system might not detect any corruption in unused or unallocated parts of the disk.



You must finish this procedure before performing manual steps to recover the volumes, such as adding or re-attaching the disks, stopping the node, starting the node, or rebooting. Otherwise, when you run the `reformat_storage_block_devices.rb` script, you might encounter a file system error that causes the script to hang or fail.



Repair the hardware and properly attach the disks before running the `reboot` command.




Identify failed storage volumes carefully. You will use this information to verify which volumes must be reformatted. After a volume has been reformatted, data on the volume can't be recovered.

To correctly recover failed storage volumes, you need to know both the device names of the failed storage volumes and their volume IDs.

At installation, each storage device is assigned a file system universal unique identifier (UUID) and is mounted to a `rangedb` directory on the Storage Node using that assigned file system UUID. The file system UUID and the `rangedb` directory are listed in the `/etc/fstab` file. The device name, `rangedb` directory, and the size of the mounted volume are displayed in the Grid Manager.

In the following example, device `/dev/sdc` has a volume size of 4 TB, is mounted to `/var/local/rangedb/0`, using the device name `/dev/disk/by-uuid/822b0547-3b2b-472e-ad5e-e1cf1809faba` in the `/etc/fstab` file:



Steps

1. Complete the following steps to record the failed storage volumes and their device names:
 - a. Select **SUPPORT > Tools > Grid topology**.
 - b. Select **site > failed Storage Node > LDR > Storage > Overview > Main**, and look for object stores with alarms.

Object Stores

ID	Total	Available	Stored Data	Stored (%)	Health
0000	96.6 GB	96.6 GB	823 KB	0.001 %	Error
0001	107 GB	107 GB	0 B	0 %	No Errors
0002	107 GB	107 GB	0 B	0 %	No Errors

- c. Select **site > failed Storage Node > SSM > Resources > Overview > Main**. Determine the mount point and volume size of each failed storage volume identified in the previous step.

Object stores are numbered in hex notation. For example, 0000 is the first volume and 000F is the sixteenth volume. In the example, the object store with an ID of 0000 corresponds to `/var/local/rangedb/0` with device name `sdc` and a size of 107 GB.

Volumes

Mount Point	Device	Status	Size	Space Available	Total Entries	Entries Available	Write Cache
/	croot	Online	10.4 GB	4.17 GB	655,360	554,806	Unknown
/var/local	cvloc	Online	96.6 GB	96.1 GB	94,369,792	94,369,423	Unknown
/var/local/rangedb/0	sdc	Online	107 GB	107 GB	104,857,600	104,856,202	Enabled
/var/local/rangedb/1	sdd	Online	107 GB	107 GB	104,857,600	104,856,536	Enabled
/var/local/rangedb/2	sde	Online	107 GB	107 GB	104,857,600	104,856,536	Enabled

2. Log in to the failed Storage Node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

3. Run the following script to unmount a failed storage volume:

```
sn-unmount-volume object_store_ID
```

The `object_store_ID` is the ID of the failed storage volume. For example, specify `0` in the command for an object store with ID `0000`.

4. If prompted, press **y** to stop the Cassandra service depending on storage volume `0`.



If the Cassandra service is already stopped, you aren't prompted. The Cassandra service is stopped only for volume `0`.

```
root@Storage-180:~/var/local/tmp/storage~ # sn-unmount-volume 0
Services depending on storage volume 0 (cassandra) aren't down.
Services depending on storage volume 0 must be stopped before running
this script.
Stop services that require storage volume 0 [y/N]? y
Shutting down services that require storage volume 0.
Services requiring storage volume 0 stopped.
Unmounting /var/local/rangedb/0
/var/local/rangedb/0 is unmounted.
```

In a few seconds, the volume is unmounted. Messages appear indicating each step of the process. The final message indicates that the volume is unmounted.

5. If the unmount fails because the volume is busy, you can force an unmount using the `--use-umountof` option:



Forcing an unmount using the `--use-umountof` option might cause processes or services using the volume to behave unexpectedly or crash.

```
root@Storage-180:~ # sn-unmount-volume --use-umountof
/var/local/rangedb/2
Unmounting /var/local/rangedb/2 using umountof
/var/local/rangedb/2 is unmounted.
Informing LDR service of changes to storage volumes
```

Recover failed storage volumes and rebuild Cassandra database

You must run a script that reformats and remounts storage on failed storage volumes, and rebuilds the Cassandra database on the Storage Node if the system determines that

it is necessary.

Before you begin

- You have the `Passwords.txt` file.
- The system drives on the server are intact.
- The cause of the failure has been identified and, if necessary, replacement storage hardware has already been acquired.
- The total size of the replacement storage is the same as the original.
- You have checked that a Storage Node decommissioning is not in progress, or you have paused the node decommission procedure. (In the Grid Manager, select **MAINTENANCE > Tasks > Decommission.**)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **MAINTENANCE > Tasks > Expansion.**)
- You have [reviewed the warnings about storage volume recovery](#).

Steps

1. As needed, replace failed physical or virtual storage associated with the failed storage volumes that you identified and unmounted earlier.

Don't remount the volumes in this step. The storage is remounted and added to `/etc/fstab` in a later step.

2. In the Grid Manager, go to **NODES > appliance Storage Node > Hardware**. In the StorageGRID Appliance section of the page, verify that the Storage RAID mode is healthy.
3. Log in to the failed Storage Node:
 - a. Enter the following command: `ssh admin@grid_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 `#`.

4. Use a text editor (`vi` or `vim`) to delete failed volumes from the `/etc/fstab` file and then save the file.



Commenting out a failed volume in the `/etc/fstab` file is insufficient. The volume must be deleted from `fstab` as the recovery process verifies that all lines in the `fstab` file match the mounted file systems.

5. Reformat any failed storage volumes and rebuild the Cassandra database if it is necessary. Enter:
`reformat_storage_block_devices.rb`
 - When storage volume 0 is unmounted, prompts and messages will indicate that the Cassandra service is being stopped.
 - You will be prompted to rebuild the Cassandra database if it is necessary.
 - Review the warnings. If none of them apply, rebuild the Cassandra database. Enter: **y**
 - If more than one Storage Node is offline or if another Storage Node has been rebuilt in the last 15 days. Enter: **n**

The script will exit without rebuilding Cassandra. Contact technical support.

- For each rangedb drive on the Storage Node, when you are asked: Reformat the rangedb drive <name> (device <major number>:<minor number>)? [y/n]?, enter one of the following responses:

- **y** to reformat a drive that had errors. This reformats the storage volume and adds the reformatted storage volume to the `/etc/fstab` file.
- **n** if the drive contains no errors, and you don't want to reformat it.



Selecting **n** exits the script. Either mount the drive (if you think the data on the drive should be retained and the drive was unmounted in error) or remove the drive. Then, run the `reformat_storage_block_devices.rb` command again.



Some StorageGRID recovery procedures use Reaper to handle Cassandra repairs. Repairs occur automatically as soon as the related or required services have started. You might notice script output that mentions "reaper" or "Cassandra repair." If you see an error message indicating the repair has failed, run the command indicated in the error message.

In the following example output, the drive `/dev/sdf` must be reformatted, and Cassandra did not need to be rebuilt:

```
root@DC1-S1:~ # reformat_storage_block_devices.rb
Formatting devices that are not in use...
Skipping in use device /dev/sdc
Skipping in use device /dev/sdd
Skipping in use device /dev/sde
Reformat the rangedb drive /dev/sdf (device 8:64)? [Y/n]? y
Successfully formatted /dev/sdf with UUID b951bfcb-4804-41ad-b490-805dfd8df16c
All devices processed
Running: /usr/local/ldr/setup_rangedb.sh 12368435
Cassandra does not need rebuilding.
Starting services.
Informing storage services of new volume

Reformatting done. Now do manual steps to
restore copies of data.
```

After the storage volumes are reformatted and remounted and necessary Cassandra operations are complete, you can [restore object data using Grid Manager](#).

Restore object data to storage volume where system drive is intact

After recovering a storage volume on a Storage Node where the system drive is intact, you can restore the replicated or erasure-coded object data that was lost when the

storage volume failed.

Which procedure should I use?

Whenever possible, restore object data using the **Volume restoration** page in the Grid Manager.

- If the volumes are listed at **MAINTENANCE > Volume restoration > Nodes to restore**, restore object data using the [Volume restoration page in the Grid Manager](#).
- If the volumes aren't listed at **MAINTENANCE > Volume restoration > Nodes to restore**, follow the steps below for using the `repair-data` script to restore object data.

If the recovered Storage Node contains fewer volumes than the node it is replacing, you must use the `repair-data` script.



The `repair-data` script is deprecated and will be removed in a future release. When possible, use the [Volume restoration procedure in the Grid Manager](#).

Use the `repair-data` script to restore object data

Before you begin

- You have confirmed that the recovered Storage Node has a Connection State of **Connected**  on the **NODES > Overview** tab in the Grid Manager.

About this task

Object data can be restored from other Storage Nodes or a Cloud Storage Pool, assuming that the grid's ILM rules were configured such that object copies are available.

Note the following:

- If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.
- If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.

About the `repair-data` script

To restore object data, you run the `repair-data` script. This script begins the process of restoring object data and works with ILM scanning to ensure that ILM rules are met.

Select **Replicated data** or **Erasure-coded (EC) data** below to learn the different options for the `repair-data` script, based on whether you are restoring replicated data or erasure-coded data. If you need to restore both types of data, you must run both sets of commands.



For more information about the `repair-data` script, enter `repair-data --help` from the command line of the primary Admin Node.



The `repair-data` script is deprecated and will be removed in a future release. When possible, use the [Volume restoration procedure in the Grid Manager](#).

Replicated data

Two commands are available for restoring replicated data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-replicated-node-repair
```

```
repair-data start-replicated-volume-repair
```

You can track repairs of replicated data with this command:

```
repair-data show-replicated-repair-status
```

Erasure-coded (EC) data

Two commands are available for restoring erasure-coded data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-ec-node-repair
```

```
repair-data start-ec-volume-repair
```

You can track repairs of erasure-coded data with this command:

```
repair-data show-ec-repair-status
```



Repairs of erasure-coded data can begin while some Storage Nodes are offline. However, if all erasure-coded data can't be accounted for, the repair can't be completed. Repair will complete after all nodes are available.



The EC repair job temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the repair is complete. If there is not enough storage for the reservation, the EC repair job will fail. Storage reservations are released when the EC repair job completes, whether the job failed or succeeded.

Find hostname for Storage Node

1. Log in to the primary Admin Node:

- a. Enter the following command: `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 `#`.

2. Use the `/etc/hosts` file to find the hostname of the Storage Node for the restored storage volumes. To see a list of all nodes in the grid, enter the following: `cat /etc/hosts`.

Repair data if all volumes have failed

If all storage volumes have failed, repair the entire node. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If only some volumes have failed, go to [Repair data if only some volumes have failed](#).



You can't run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

Replicated data

If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the replicated data on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-node-repair --nodes SG-DC-SN3
```



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See [Investigate lost objects](#).

Erasure-coded (EC) data

If your grid contains erasure-coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the erasure-coded data on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-node-repair --nodes SG-DC-SN3
```

The operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.

Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Repair data if only some volumes have failed

If only some of the volumes have failed, repair the affected volumes. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If all volumes have failed, go to [Repair data if all volumes have failed](#).

Enter the volume IDs in hexadecimal. For example, 0000 is the first volume and 000F is the sixteenth volume. You can specify one volume, a range of volumes, or multiple volumes that aren't in a sequence.

All the volumes must be on the same Storage Node. If you need to restore volumes for more than one Storage Node, contact technical support.

Replicated data

If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores replicated data to volume 0002 on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes 0002
```

Range of volumes: This command restores replicated data to all volumes in the range 0003 to 0009 on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volume-range 0003,0009
```

Multiple volumes not in a sequence: This command restores replicated data to volumes 0001, 0005, and 0008 on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes 0001,0005,0008
```



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. Note the alert description and recommended actions to determine the cause of the loss and if recovery is possible.

Erasure-coded (EC) data

If your grid contains erasure-coded data, use the `start-ec-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores erasure-coded data to volume 0007 on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 0007
```

Range of volumes: This command restores erasure-coded data to all volumes in the range 0004 to 0006 on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volume-range 0004,0006
```

Multiple volumes not in a sequence: This command restores erasure-coded data to volumes 000A, 000C, and 000E on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 000A,000C,000E
```

The `repair-data` operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.



Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Monitor repairs

Monitor the status of the repair jobs, based on whether you use **replicated data**, **erasure-coded (EC) data**, or both.

You can also monitor the status of volume restoration jobs in process and view a history of restoration jobs completed in [Grid Manager](#).

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the `repair-data` command.

```
repair-data show-replicated-repair-status
```

- To determine if repairs are complete:
 1. Select **NODES > Storage Node being repaired > ILM**.
 2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
 1. Select **SUPPORT > Tools > Grid topology**.
 2. Select **grid > Storage Node being repaired > LDR > Data Store**.
 3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.



Cassandra inconsistencies might be present, and failed repairs aren't tracked.

- **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.



High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

- **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure-coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
 - Select **SUPPORT > Tools > Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.

- Use this command to see the status of a specific `repair-data` operation:

```
repair-data show-ec-repair-status --repair-id repair ID
```

- Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

This command retries a failed node repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-node-repair --repair-id 6949309319275667690
```

This command retries a failed volume repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-volume-repair --repair-id 6949309319275667690
```

Check storage state after recovering storage volumes

After recovering storage volumes, you must verify that the desired state of the Storage Node is set to online and ensure that the state will be online by default whenever the Storage Node server is restarted.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- The Storage Node has been recovered, and data recovery is complete.

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Check the values of **Recovered Storage Node > LDR > Storage > Storage State — Desired** and **Storage State — Current**.

The value of both attributes should be Online.

3. If the Storage State — Desired is set to Read-only, complete the following steps:
 - a. Click the **Configuration** tab.
 - b. From the **Storage State — Desired** drop-down list, select **Online**.
 - c. Click **Apply Changes**.
 - d. Click the **Overview** tab and confirm that the values of **Storage State — Desired** and **Storage State — Current** are updated to Online.

Recover from system drive failure

Warnings for Storage Node system drive recovery

Before recovering a failed system drive of a Storage Node, review the general [warnings and considerations for grid node recovery](#) and the following specific warnings.

Storage Nodes have a Cassandra database that includes object metadata. The Cassandra database might be rebuilt in the following circumstances:

- A Storage Node is brought back online after having been offline for more than 15 days.
- A storage volume has failed and been recovered.
- The system drive and one or more storage volumes fails and is recovered.

When Cassandra is rebuilt, the system uses information from other Storage Nodes. If too many Storage Nodes are offline, some Cassandra data might not be available. If Cassandra has been rebuilt recently, Cassandra data might not yet be consistent across the grid. Data loss can occur if Cassandra is rebuilt when too many Storage Nodes are offline or if two or more Storage Nodes are rebuilt within 15 days of each other.



If more than one Storage Node has failed (or is offline), contact technical support. Don't perform the following recovery procedure. Data loss could occur.



If this is the second Storage Node failure in less than 15 days after a Storage Node failure or recovery, contact technical support. Rebuilding Cassandra on two or more Storage Nodes within 15 days can result in data loss.



If more than one Storage Node at a site has failed, a site recovery procedure might be required. See [How technical support recovers a site](#).



If this Storage Node is in read-only maintenance mode to allow for the retrieval of objects by another Storage Node with failed storage volumes, recover volumes on the Storage Node with failed storage volumes before recovering this failed Storage Node. See the instructions to [recover from storage volume failure where system drive is intact](#).



If ILM rules are configured to store only one replicated copy and the copy exists on a storage volume that has failed, you will not be able to recover the object.

Replace the Storage Node

If the system drive has failed, you must first replace the Storage Node.

You must select the node replacement procedure for your platform. The steps to replace a node are the same for all types of grid nodes.



This procedure applies to software-based Storage Nodes only. You must follow a different procedure to [recover an appliance Storage Node](#).

Linux: If you aren't sure if your system drive has failed, follow the instructions to replace the node to determine which recovery steps are required.

Platform	Procedure
VMware	Replace a VMware node
Linux	Replace a Linux node

Platform	Procedure
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node .

Select Start Recovery to configure Storage Node

After replacing a Storage Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

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 the provisioning passphrase.
- You have deployed and configured the replacement node.
- You have the start date of any repair jobs for erasure-coded data.
- You have verified that the Storage Node has not been rebuilt within the last 15 days.

About this task

If the Storage Node is installed as a container on a Linux host, you must perform this step only if one of these is true:

- You had to use the `--force` flag to import the node, or you issued `storagegrid node force-recovery node-name`
- You had to do a full node reinstall, or you needed to restore `/var/local`.

Steps

1. From the Grid Manager, select **MAINTENANCE > Tasks > Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

Nodes appear in the list after they fail, but you can't select a node until it has been reinstalled and is ready for recovery.

3. Enter the **Provisioning Passphrase**.
4. Click **Start Recovery**.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

Search 				
	Name	IPv4 Address	State	Recoverable
☒	104-217-S1	10.96.104.217	Unknown	

Passphrase

Provisioning Passphrase

Start Recovery

5. Monitor the progress of the recovery in the Recovering Grid Node table.



While the recovery procedure is running, you can click **Reset** to start a new recovery. A dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

Info

Reset Recovery

Resetting the recovery procedure leaves the deployed grid node in an indeterminate state. To retry a recovery after resetting the procedure, you must restore the node to a pre-installed state:

- For VMware nodes, delete the deployed VM and then redeploy it.
- For StorageGRID appliance nodes, run "sgareinstall" on the node.
- For Linux nodes, run "storagegrid node force-recovery *node-name*" on the Linux host.

Do you want to reset recovery?

Cancel

OK

If you want to retry the recovery after resetting the procedure, you must restore the node to a pre-installed state, as follows:

- **VMware:** Delete the deployed virtual grid node. Then, when you are ready to restart the recovery, redeploy the node.
- **Linux:** Restart the node by running this command on the Linux host: `storagegrid node force-recovery node-name`

6. When the Storage Node reaches the "Waiting for Manual Steps" stage, go to [Remount and reformat storage volumes \(manual steps\)](#).

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Recovering Grid Node

Name	Start Time	Progress	Stage
dc2-s3	2016-09-12 16:12:40 PDT		Waiting For Manual Steps

Reset

Remount and reformat storage volumes (manual steps)

You must manually run two scripts to remount preserved storage volumes and to reformat any failed storage volumes. The first script remounts volumes that are properly formatted as StorageGRID storage volumes. The second script reformats any unmounted volumes, rebuilds Cassandra, if needed, and starts services.

Before you begin

- You have already replaced the hardware for any failed storage volumes that you know require replacement.

Running the `sn-remount-volumes` script might help you identify additional failed storage volumes.

- You have checked that a Storage Node decommissioning is not in progress, or you have paused the node decommission procedure. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **MAINTENANCE** > **Tasks** > **Expansion**.)
- You have [reviewed the warnings for Storage Node system drive recovery](#).



Contact technical support if more than one Storage Node is offline or if a Storage Node in this grid has been rebuilt in the last 15 days. Don't run the `sn-recovery-postinstall.sh` script. Rebuilding Cassandra on two or more Storage Nodes within 15 days of each other might result in data loss.

About this task

To complete this procedure, you perform these high-level tasks:

- Log in to the recovered Storage Node.
- Run the `sn-remount-volumes` script to remount properly formatted storage volumes. When this script runs, it does the following:
 - Mounts and unmounts each storage volume to replay the XFS journal.
 - Performs an XFS file consistency check.
 - If the file system is consistent, determines if the storage volume is a properly formatted StorageGRID storage volume.
 - If the storage volume is properly formatted, remounts the storage volume. Any existing data on the volume remains intact.
- Review the script output and resolve any issues.

- Run the `sn-recovery-postinstall.sh` script. When this script runs, it does the following.



Don't reboot a Storage Node during recovery before running `sn-recovery-postinstall.sh` to reformat the failed storage volumes and restore object metadata. Rebooting the Storage Node before `sn-recovery-postinstall.sh` completes causes errors for services that attempt to start and causes StorageGRID appliance nodes to exit maintenance mode. See the step for [post-install script](#).

- Reformats any storage volumes that the `sn-remount-volumes` script could not mount or that were found to be improperly formatted.



If a storage volume is reformatted, any data on that volume is lost. You must perform an additional procedure to restore object data from other locations in the grid, assuming that ILM rules were configured to store more than one object copy.

- Rebuilds the Cassandra database on the node, if needed.
- Starts the services on the Storage Node.

Steps

1. Log in to the recovered Storage Node:

- a. Enter the following command: `ssh admin@grid_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 `#`.

2. Run the first script to remount any properly formatted storage volumes.



If all storage volumes are new and need to be formatted, or if all storage volumes have failed, you can skip this step and run the second script to reformat all unmounted storage volumes.

- a. Run the script: `sn-remount-volumes`

This script might take hours to run on storage volumes that contain data.

- b. As the script runs, review the output and answer any prompts.



As required, you can use the `tail -f` command to monitor the contents of the script's log file (`/var/local/log/sn-remount-volumes.log`). The log file contains more detailed information than the command line output.

```
root@SG:~ # sn-remount-volumes
The configured LDR noid is 12632740

===== Device /dev/sdb =====
```

```

Mount and unmount device /dev/sdb and checking file system
consistency:
The device is consistent.
Check rangedb structure on device /dev/sdb:
Mount device /dev/sdb to /tmp/sdb-654321 with rangedb mount options
This device has all rangedb directories.
Found LDR node id 12632740, volume number 0 in the volID file
Attempting to remount /dev/sdb
Device /dev/sdb remounted successfully

===== Device /dev/sdc =====
Mount and unmount device /dev/sdc and checking file system
consistency:
Error: File system consistency check retry failed on device /dev/sdc.
You can see the diagnosis information in the /var/local/log/sn-
remount-volumes.log.

This volume could be new or damaged. If you run sn-recovery-
postinstall.sh,
this volume and any data on this volume will be deleted. If you only
had two
copies of object data, you will temporarily have only a single copy.
StorageGRID will attempt to restore data redundancy by making
additional replicated copies or EC fragments, according to the rules
in
the active ILM policies.

Don't continue to the next step if you believe that the data
remaining on
this volume can't be rebuilt from elsewhere in the grid (for example,
if
your ILM policy uses a rule that makes only one copy or if volumes
have
failed on multiple nodes). Instead, contact support to determine how
to
recover your data.

===== Device /dev/sdd =====
Mount and unmount device /dev/sdd and checking file system
consistency:
Failed to mount device /dev/sdd
This device could be an uninitialized disk or has corrupted
superblock.
File system check might take a long time. Do you want to continue? (y
or n) [y/N]? y

```

```
Error: File system consistency check retry failed on device /dev/sdd.  
You can see the diagnosis information in the /var/local/log/sn-  
remount-volumes.log.
```

```
This volume could be new or damaged. If you run sn-recovery-  
postinstall.sh,  
this volume and any data on this volume will be deleted. If you only  
had two  
copies of object data, you will temporarily have only a single copy.  
StorageGRID will attempt to restore data redundancy by making  
additional replicated copies or EC fragments, according to the rules  
in  
the active ILM policies.
```

```
Don't continue to the next step if you believe that the data  
remaining on  
this volume can't be rebuilt from elsewhere in the grid (for example,  
if  
your ILM policy uses a rule that makes only one copy or if volumes  
have  
failed on multiple nodes). Instead, contact support to determine how  
to  
recover your data.
```

```
===== Device /dev/sde =====
```

```
Mount and unmount device /dev/sde and checking file system  
consistency:
```

```
The device is consistent.
```

```
Check rangedb structure on device /dev/sde:
```

```
Mount device /dev/sde to /tmp/sde-654321 with rangedb mount options
```

```
This device has all rangedb directories.
```

```
Found LDR node id 12000078, volume number 9 in the volID file
```

```
Error: This volume does not belong to this node. Fix the attached  
volume and re-run this script.
```

In the example output, one storage volume was remounted successfully and three storage volumes had errors.

- /dev/sdb passed the XFS file system consistency check and had a valid volume structure, so it was remounted successfully. Data on devices that are remounted by the script is preserved.
- /dev/sdc failed the XFS file system consistency check because the storage volume was new or corrupt.
- /dev/sdd could not be mounted because the disk was not initialized or the disk's superblock was corrupted. When the script can't mount a storage volume, it asks if you want to run the file system consistency check.
 - If the storage volume is attached to a new disk, answer **N** to the prompt. You don't need check

the file system on a new disk.

- If the storage volume is attached to an existing disk, answer **Y** to the prompt. You can use the results of the file system check to determine the source of the corruption. The results are saved in the `/var/local/log/sn-remount-volumes.log` log file.
- `/dev/sde` passed the XFS file system consistency check and had a valid volume structure; however, the LDR node ID in the `volID` file did not match the ID for this Storage Node (the configured LDR `noid` displayed at the top). This message indicates that this volume belongs to another Storage Node.

3. Review the script output and resolve any issues.



If a storage volume failed the XFS file system consistency check or could not be mounted, carefully review the error messages in the output. You must understand the implications of running the `sn-recovery-postinstall.sh` script on these volumes.

- a. Check to make sure that the results include an entry for all of the volumes you expected. If any volumes aren't listed, rerun the script.
- b. Review the messages for all mounted devices. Make sure there are no errors indicating that a storage volume does not belong to this Storage Node.

In the example, the output for `/dev/sde` includes the following error message:

```
Error: This volume does not belong to this node. Fix the attached
volume and re-run this script.
```



If a storage volume is reported as belonging to another Storage Node, contact technical support. If you run the `sn-recovery-postinstall.sh` script, the storage volume will be reformatted, which might cause data loss.

- c. If any storage devices could not be mounted, make a note of the device name, and repair or replace the device.



You must repair or replace any storage devices that could not be mounted.

You will use the device name to look up the volume ID, which is required input when you run the `repair-data` script to restore object data to the volume (the next procedure).

- d. After repairing or replacing all unmountable devices, run the `sn-remount-volumes` script again to confirm that all storage volumes that can be remounted have been remounted.



If a storage volume can't be mounted or is improperly formatted, and you continue to the next step, the volume and any data on the volume will be deleted. If you had two copies of object data, you will have only a single copy until you complete the next procedure (restoring object data).



Don't run the `sn-recovery-postinstall.sh` script if you believe that the data remaining on a failed storage volume can't be rebuilt from elsewhere in the grid (for example, if your ILM policy uses a rule that makes only one copy or if volumes have failed on multiple nodes). Instead, contact technical support to determine how to recover your data.

4. Run the `sn-recovery-postinstall.sh` script: `sn-recovery-postinstall.sh`

This script reformats any storage volumes that could not be mounted or that were found to be improperly formatted; rebuilds the Cassandra database on the node, if needed; and starts the services on the Storage Node.

Be aware of the following:

- The script might take hours to run.
- In general, you should leave the SSH session alone while the script is running.
- Don't press **Ctrl+C** while the SSH session is active.
- The script will run in the background if a network disruption occurs and terminates the SSH session, but you can view the progress from the Recovery page.
- If the Storage Node uses the RSM service, the script might appear to stall for 5 minutes as node services are restarted. This 5-minute delay is expected whenever the RSM service boots for the first time.



The RSM service is present on Storage Nodes that include the ADC service.



Some StorageGRID recovery procedures use Reaper to handle Cassandra repairs. Repairs occur automatically as soon as the related or required services have started. You might notice script output that mentions "reaper" or "Cassandra repair." If you see an error message indicating the repair has failed, run the command indicated in the error message.

5. As the `sn-recovery-postinstall.sh` script runs, monitor the Recovery page in the Grid Manager.

The Progress bar and the Stage column on the Recovery page provide a high-level status of the `sn-recovery-postinstall.sh` script.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

Search			
Name	IPv4 Address	State	Recoverable
No results found.			

Recovering Grid Node

Name	Start Time	Progress	Stage
DC1-S3	2016-06-02 14:03:35 PDT		Recovering Cassandra

6. After the `sn-recovery-postinstall.sh` script has started services on the node, you can restore object data to any storage volumes that were formatted by the script.

The script asks if you want to use the Grid Manager volume restoration process.

- In most cases, you should [restore object data using Grid Manager](#). Answer `y` to use the Grid Manager.
- In rare cases, such as when instructed by technical support, or when you know that the replacement node has fewer volumes available for object storage than the original node, you must [restore object data manually](#) using the `repair-data` script. If one of these cases applies, answer `n`.



If you answer `n` to using the Grid Manager volume restoration process (restore object data manually):

- You aren't able to restore object data using Grid Manager.
- You can monitor the progress of manual restoration jobs using Grid Manager.

After making your selection, the script completes and the next steps to recover object data are shown. After reviewing these steps, press any key to return to the command line.

Restore object data to storage volume (system drive failure)

After recovering storage volumes for a non-appliance Storage Node, you can restore the replicated or erasure-coded object data that was lost when the Storage Node failed.

Which procedure should I use?

Whenever possible, restore object data using the **Volume restoration** page in the Grid Manager.

- If the volumes are listed at **MAINTENANCE > Volume restoration > Nodes to restore**, restore object data using the [Volume restoration page in the Grid Manager](#).
- If the volumes aren't listed at **MAINTENANCE > Volume restoration > Nodes to restore**, follow the steps below for using the `repair-data` script to restore object data.

If the recovered Storage Node contains fewer volumes than the node it is replacing, you must use the `repair-data` script.



The `repair-data` script is deprecated and will be removed in a future release. When possible, use the [Volume restoration procedure in the Grid Manager](#).

Use the `repair-data` script to restore object data

Before you begin

- You have confirmed that the recovered Storage Node has a Connection State of **Connected**  on the **NODES > Overview** tab in the Grid Manager.

About this task

Object data can be restored from other Storage Nodes or a Cloud Storage Pool, assuming that the grid's ILM rules were configured such that object copies are available.

Note the following:

- If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.
- If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.

About the `repair-data` script

To restore object data, you run the `repair-data` script. This script begins the process of restoring object data and works with ILM scanning to ensure that ILM rules are met.

Select **Replicated data** or **Erasure-coded (EC) data** below to learn the different options for the `repair-data` script, based on whether you are restoring replicated data or erasure-coded data. If you need to restore both types of data, you must run both sets of commands.



For more information about the `repair-data` script, enter `repair-data --help` from the command line of the primary Admin Node.



The `repair-data` script is deprecated and will be removed in a future release. When possible, use the [Volume restoration procedure in the Grid Manager](#).

Replicated data

Two commands are available for restoring replicated data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-replicated-node-repair
```

```
repair-data start-replicated-volume-repair
```

You can track repairs of replicated data with this command:

```
repair-data show-replicated-repair-status
```

Erasure-coded (EC) data

Two commands are available for restoring erasure-coded data, based on whether you need to repair the entire node or only certain volumes on the node:

```
repair-data start-ec-node-repair
```

```
repair-data start-ec-volume-repair
```

You can track repairs of erasure-coded data with this command:

```
repair-data show-ec-repair-status
```



Repairs of erasure-coded data can begin while some Storage Nodes are offline. However, if all erasure-coded data can't be accounted for, the repair can't be completed. Repair will complete after all nodes are available.



The EC repair job temporarily reserves a large amount of storage. Storage alerts might be triggered, but will resolve when the repair is complete. If there is not enough storage for the reservation, the EC repair job will fail. Storage reservations are released when the EC repair job completes, whether the job failed or succeeded.

Find hostname for Storage Node

1. Log in to the primary Admin Node:

- a. Enter the following command: `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 `#`.

2. Use the `/etc/hosts` file to find the hostname of the Storage Node for the restored storage volumes. To see a list of all nodes in the grid, enter the following: `cat /etc/hosts`.

Repair data if all volumes have failed

If all storage volumes have failed, repair the entire node. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If only some volumes have failed, go to [Repair data if only some volumes have failed](#).



You can't run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

Replicated data

If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the replicated data on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-node-repair --nodes SG-DC-SN3
```



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. You should determine the cause of the loss and if recovery is possible. See [Investigate lost objects](#).

Erasure-coded (EC) data

If your grid contains erasure-coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option, where `--nodes` is the hostname (system name), to repair the entire Storage Node.

This command repairs the erasure-coded data on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-node-repair --nodes SG-DC-SN3
```

The operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.

Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Repair data if only some volumes have failed

If only some of the volumes have failed, repair the affected volumes. Follow the instructions for **replicated data**, **erasure-coded (EC) data**, or both, based on whether you use replicated data, erasure-coded (EC) data, or both.

If all volumes have failed, go to [Repair data if all volumes have failed](#).

Enter the volume IDs in hexadecimal. For example, `0000` is the first volume and `000F` is the sixteenth volume. You can specify one volume, a range of volumes, or multiple volumes that aren't in a sequence.

All the volumes must be on the same Storage Node. If you need to restore volumes for more than one Storage Node, contact technical support.

Replicated data

If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores replicated data to volume 0002 on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes 0002
```

Range of volumes: This command restores replicated data to all volumes in the range 0003 to 0009 on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volume-range 0003,0009
```

Multiple volumes not in a sequence: This command restores replicated data to volumes 0001, 0005, and 0008 on a Storage Node named SG-DC-SN3:

```
repair-data start-replicated-volume-repair --nodes SG-DC-SN3 --volumes 0001,0005,0008
```



As object data is restored, the **Objects Lost** alert is triggered if the StorageGRID system can't locate replicated object data. Alerts might be triggered on Storage Nodes throughout the system. Note the alert description and recommended actions to determine the cause of the loss and if recovery is possible.

Erasure-coded (EC) data

If your grid contains erasure-coded data, use the `start-ec-volume-repair` command with the `--nodes` option to identify the node (where `--nodes` is the hostname of the node). Then add either the `--volumes` or `--volume-range` option, as shown in the following examples.

Single volume: This command restores erasure-coded data to volume 0007 on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 0007
```

Range of volumes: This command restores erasure-coded data to all volumes in the range 0004 to 0006 on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volume-range 0004,0006
```

Multiple volumes not in a sequence: This command restores erasure-coded data to volumes 000A, 000C, and 000E on a Storage Node named SG-DC-SN3:

```
repair-data start-ec-volume-repair --nodes SG-DC-SN3 --volumes 000A,000C,000E
```

The `repair-data` operation returns a unique `repair ID` that identifies this `repair_data` operation. Use this `repair ID` to track the progress and result of the `repair_data` operation. No other feedback is returned as the recovery process completes.



Repairs of erasure-coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available.

Monitor repairs

Monitor the status of the repair jobs, based on whether you use **replicated data**, **erasure-coded (EC) data**, or both.

You can also monitor the status of volume restoration jobs in process and view a history of restoration jobs completed in [Grid Manager](#).

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the `repair-data` command.

```
repair-data show-replicated-repair-status
```

- To determine if repairs are complete:
 1. Select **NODES > Storage Node being repaired > ILM**.
 2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
 1. Select **SUPPORT > Tools > Grid topology**.
 2. Select **grid > Storage Node being repaired > LDR > Data Store**.
 3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.



Cassandra inconsistencies might be present, and failed repairs aren't tracked.

- **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.



High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

- **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure-coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
 - Select **SUPPORT > Tools > Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.

- Use this command to see the status of a specific `repair-data` operation:

```
repair-data show-ec-repair-status --repair-id repair ID
```

- Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

This command retries a failed node repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-node-repair --repair-id 6949309319275667690
```

This command retries a failed volume repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-volume-repair --repair-id 6949309319275667690
```

Check storage state after recovering Storage Node system drive

After recovering the system drive for a Storage Node, you must verify that the desired state of the Storage Node is set to online and ensure that the state will be online by default whenever the Storage Node server is restarted.

Before you begin

- You are signed in to the Grid Manager using a [supported web browser](#).
- The Storage Node has been recovered, and data recovery is complete.

Steps

1. Select **SUPPORT > Tools > Grid topology**.
2. Check the values of **Recovered Storage Node > LDR > Storage > Storage State — Desired** and **Storage State — Current**.

The value of both attributes should be Online.

3. If the Storage State — Desired is set to Read-only, complete the following steps:
 - a. Click the **Configuration** tab.
 - b. From the **Storage State — Desired** drop-down list, select **Online**.
 - c. Click **Apply Changes**.
 - d. Click the **Overview** tab and confirm that the values of **Storage State — Desired** and **Storage State — Current** are updated to Online.

Restore object data using Grid Manager

You can restore object data for a failed storage volume or Storage Node using Grid Manager. You can also use Grid Manager to monitor restoration processes in progress and display a restoration history.

Before you begin

- You have completed either of these procedures to format failed volumes:
 - [Remount and reformat appliance storage volumes \(manual steps\)](#)

- [Remount and reformat storage volumes \(manual steps\)](#)
- You have confirmed that the Storage Node where you are restoring objects has a Connection State of **Connected**  on the **NODES > Overview** tab in the Grid Manager.
- You have confirmed the following:
 - A grid expansion to add a Storage Node is not in process.
 - A Storage Node decommission is not in process or failed.
 - A recovery of a failed storage volume is not in process.
 - A recovery of a Storage Node with a failed system drive is not in process.
 - An EC rebalance job is not in process.
 - Appliance node cloning is not in process.

About this task

After you have replaced the drives and performed the manual steps to format the volumes, Grid Manager displays the volumes as candidates for restoration on the **MAINTENANCE > Volume restoration > Nodes to restore** tab.

Whenever possible, restore object data using the Volume restoration page in the Grid Manager. You can either [enable automatic restore mode](#) to automatically start volume restoration when the volumes are ready to be restored or [manually perform volume restoration](#). Follow these guidelines:

- If the volumes are listed at **MAINTENANCE > Volume restoration > Nodes to restore**, restore object data as described in the steps below. The volumes will be listed if:
 - Some, but not all, storage volumes in a node have failed
 - All storage volumes in a node have failed and are being replaced with the same number of volumes or more volumes

The Volume restoration page in the Grid Manager also allows you to [monitor the volume restoration process](#) and [view restoration history](#).

- If the volumes aren't listed in the Grid Manager as candidates for restoration, follow the appropriate steps for using the `repair-data` script to restore object data:
 - [Restoring object data to storage volume \(system drive failure\)](#)
 - [Restore object data to storage volume where system drive is intact](#)
 - [Restore object data to storage volume for appliance](#)



The `repair-data` script is deprecated and will be removed in a future release.

If the recovered Storage Node contains fewer volumes than the node it is replacing, you must use the `repair-data` script.

You can restore two types of object data:

- Replicated data objects are restored from other locations, assuming that the grid's ILM rules were configured to make object copies available.
 - If an ILM rule was configured to store only one replicated copy and that copy existed on a storage volume that failed, you will not be able to recover the object.

- If the only remaining copy of an object is in a Cloud Storage Pool, StorageGRID must issue multiple requests to the Cloud Storage Pool endpoint to restore object data.
- Erasure-coded (EC) data objects are restored by reassembling the stored fragments. Corrupt or lost fragments are recreated by the erasure-coding algorithm from the remaining data and parity fragments.

Repairs of erasure-coded data can begin while some Storage Nodes are offline. However, if all erasure-coded data cannot be accounted for, the repair can't be completed. Repair will complete after all nodes are available.



Volume restoration is dependent on the availability of resources where object copies are stored. Progress of volume restoration is nonlinear and might take days or weeks to complete.

Enable automatic restore mode

When you enable Automatic restore mode, volume restoration automatically starts when the volumes are ready to be restored.

Steps

1. In Grid Manager go to **MAINTENANCE > Volume restoration**.
2. Select the **Nodes to restore** tab, then slide the toggle for **Automatic restore mode** to the enabled position.
3. When the confirmation dialog box appears, review the details.



- You will not be able to start volume restoration jobs manually on any nodes.
- Volume restorations will begin automatically only when no other maintenance procedures are in progress.
- You can monitor the status of the job from the progress monitoring page.
- StorageGRID automatically retries volume restorations that fail to start.

4. When you understand the results of enabling Automatic restore mode, select **Yes** in the confirmation dialog box.

You can disable Automatic restore mode at any time.

Manually restore failed volume or node

Follow these steps to restore a failed volume or node.

Steps

1. In Grid Manager go to **MAINTENANCE > Volume restoration**.
2. Select the **Nodes to restore** tab, then slide the toggle for **Automatic restore mode** to the disabled position.

The number on the tab indicates the number of nodes with volumes requiring restoration.

3. Expand each node to see the volumes in it that need restoration and their status.
4. Correct any issues preventing restoration of each volume. Issues will be indicated when you select **Waiting for manual steps**, if it displays as the volume status.

5. Select a node to restore where all the volumes indicate a Ready to restore status.

You can only restore the volumes for one node at a time.

Each volume in the node must indicate that it is ready to restore.

6. Select **Start restore**.

7. Address any warnings that might appear or select **Start anyway** to ignore the warnings and start the restoration.

Nodes are moved from the **Nodes to restore** tab to the **Restoration progress** tab when the restoration starts.

If a volume restoration can't be started, the node returns to the **Nodes to restore** tab.

View restoration progress

The **Restoration progress** tab shows the status of the volume restoration process and information about the volumes for a node being restored.

Data repair rates for replicated and erasure-coded objects in all volumes are averages summarizing all restorations in process, including those restorations initiated using the `repair-data` script. The percentage of objects in those volumes that are intact and don't require restoration is also indicated.



Replicated data restoration is dependent on the availability of resources where the replicated copies are stored. Progress of replicated data restoration is nonlinear and might take days or weeks to complete.

The Restoration jobs section displays information about volume restorations started from Grid Manager.

- The number in the Restoration jobs section heading indicates the number of volumes that are either being restored or queued for restoration.
- The table displays information about each volume in a node being restored and its progress.
 - The progress for each node displays the percentage for each job.
 - Expand the Details column to display the restoration start time and job ID.
- If a volume restoration fails:
 - The Status column indicates `failed (attempting retry)`, and will be retried automatically.
 - If multiple restoration jobs have failed, the most recent job will be retried automatically first.
 - The **EC repair failure** alert is triggered if the retries continue to fail. Follow the steps in the alert to resolve the issue.

View restoration history

The **Restoration history** tab shows information about all volume restorations that have successfully completed.



Sizes aren't applicable for replicated objects and appear only for restorations that contain erasure-coded (EC) data objects.

Monitor repair-data jobs

You can monitor the status of repair jobs by using the `repair-data` script from the command line.

These include jobs that you initiated manually, or jobs that StorageGRID initiated automatically as part of a decommission procedure.



If you are running volume restoration jobs, [monitor the progress and view a history of those jobs in the Grid Manager](#) instead.

Monitor the status of `repair-data` jobs based on whether you use **replicated data**, **erasure-coded (EC) data**, or both.

Replicated data

- To get an estimated percent completion for the replicated repair, add the `show-replicated-repair-status` option to the `repair-data` command.

```
repair-data show-replicated-repair-status
```

- To determine if repairs are complete:
 1. Select **NODES > Storage Node being repaired > ILM**.
 2. Review the attributes in the Evaluation section. When repairs are complete, the **Awaiting - All** attribute indicates 0 objects.
- To monitor the repair in more detail:
 1. Select **SUPPORT > Tools > Grid topology**.
 2. Select **grid > Storage Node being repaired > LDR > Data Store**.
 3. Use a combination of the following attributes to determine, as well as possible, if replicated repairs are complete.



Cassandra inconsistencies might be present, and failed repairs aren't tracked.

- **Repairs Attempted (XRPA)**: Use this attribute to track the progress of replicated repairs. This attribute increases each time a Storage Node tries to repair a high-risk object. When this attribute does not increase for a period longer than the current scan period (provided by the **Scan Period — Estimated** attribute), it means that ILM scanning found no high-risk objects that need to be repaired on any nodes.



High-risk objects are objects that are at risk of being completely lost. This does not include objects that don't satisfy their ILM configuration.

- **Scan Period — Estimated (XSCM)**: Use this attribute to estimate when a policy change will be applied to previously ingested objects. If the **Repairs Attempted** attribute does not increase for a period longer than the current scan period, it is probable that replicated repairs are done. Note that the scan period can change. The **Scan Period — Estimated (XSCM)** attribute applies to the entire grid and is the maximum of all node scan periods. You can query the **Scan Period — Estimated** attribute history for the grid to determine an appropriate time frame.

Erasure-coded (EC) data

To monitor the repair of erasure-coded data and retry any requests that might have failed:

1. Determine the status of erasure-coded data repairs:
 - Select **SUPPORT > Tools > Metrics** to view the estimated time to completion and the completion percentage for the current job. Then, select **EC Overview** in the Grafana section. Look at the **Grid EC Job Estimated Time to Completion** and **Grid EC Job Percentage Completed** dashboards.

- Use this command to see the status of a specific `repair-data` operation:

```
repair-data show-ec-repair-status --repair-id repair ID
```

- Use this command to list all repairs:

```
repair-data show-ec-repair-status
```

The output lists information, including `repair ID`, for all previously and currently running repairs.

2. If the output shows that the repair operation failed, use the `--repair-id` option to retry the repair.

This command retries a failed node repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-node-repair --repair-id 6949309319275667690
```

This command retries a failed volume repair, using the repair ID 6949309319275667690:

```
repair-data start-ec-volume-repair --repair-id 6949309319275667690
```

Recover from Admin Node failures

Primary or non-primary Admin Node recovery

The recovery process for an Admin Node depends on whether it is the primary Admin Node or a non-primary Admin Node.

The high-level steps for recovering a primary or non-primary Admin Node are the same, although the details of the steps differ.

Always follow the correct recovery procedure for the Admin Node you are recovering. The procedures look the same at a high level, but differ in the details.

Choices

- [Recover from primary Admin Node failures](#)
- [Recover from non-primary Admin Node failures](#)

Recover from primary Admin Node failures

Recover from primary Admin Node failures

You must complete a specific set of tasks to recover from a primary Admin Node failure. The primary Admin Node hosts the Configuration Management Node (CMN) service for the grid.



You must repair or replace a failed primary Admin Node promptly or the grid might lose its ability to ingest new objects. The exact time period depends on your rate of object ingest: if you need a more accurate assessment of the time frame for your grid, contact technical support.

The Configuration Management Node (CMN) service on the primary Admin Node is responsible for issuing blocks of object identifiers for the grid. These identifiers are assigned to objects as they are ingested. New objects can't be ingested unless there are identifiers available. Object ingest can continue while the CMN is unavailable because approximately one month's supply of identifiers is cached in the grid. However, after cached identifiers are exhausted, no new objects can be added.

Follow these high-level steps to recover a primary Admin Node:

1. [Copy audit logs from failed primary Admin Node](#)
2. [Replace the primary Admin Node](#)
3. [Configure the replacement primary Admin Node](#)
4. [Determine if there is a hotfix requirement for the recovered primary Admin Node](#)
5. [Restore the audit log on the recovered primary Admin Node](#)
6. [Restore the Admin Node database when recovering a primary Admin Node](#)
7. [Restore Prometheus metrics when recovering a primary Admin Node](#)

Copy audit logs from failed primary Admin Node

If you are able to copy audit logs from the failed primary Admin Node, you should preserve them to maintain the grid's record of system activity and usage. You can restore the preserved audit logs to the recovered primary Admin Node after it is up and running.

About this task

This procedure copies the audit log files from the failed Admin Node to a temporary location on a separate grid node. These preserved audit logs can then be copied to the replacement Admin Node. Audit logs aren't automatically copied to the new Admin Node.

Depending on the type of failure, you might not be able to copy audit logs from a failed Admin Node. If the deployment has only one Admin Node, the recovered Admin Node starts recording events to the audit log in a new empty file and previously recorded data is lost. If the deployment includes more than one Admin Node, you can recover the audit logs from another Admin Node.



If the audit logs aren't accessible on the failed Admin Node now, you might be able to access them later, for example, after host recovery.

Steps

1. Log in to the failed Admin Node if possible. Otherwise, log in to the primary Admin Node or another Admin Node, if available.
 - a. Enter the following command: `ssh admin@grid_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 `#`.

2. Stop the AMS service to prevent it from creating a new log file: `service ams stop`
3. Navigate to the audit export directory:

```
cd /var/local/log
```

4. Rename the source `audit.log` file to a unique numbered file name. For example, rename the `audit.log` file to `2023-10-25.txt.1`.

```
ls -l
mv audit.log 2023-10-25.txt.1
```

5. Restart the AMS service: `service ams start`
6. Create the directory to copy all audit log files to a temporary location on a separate grid node: `ssh admin@grid_node_IP mkdir -p /var/local/tmp/saved-audit-logs`

When prompted, enter the password for admin.

7. Copy all audit log files to the temporary location: `scp -p * admin@grid_node_IP:/var/local/tmp/saved-audit-logs`

When prompted, enter the password for admin.

8. Log out as root: `exit`

Replace primary Admin Node

To recover a primary Admin Node, you must first replace the physical or virtual hardware.

You can replace a failed primary Admin Node with a primary Admin Node running on the same platform, or you can replace a primary Admin Node running on VMware or a Linux host with a primary Admin Node hosted on a services appliance.

Use the procedure that matches the replacement platform you select for the node. After you complete the node replacement procedure (which is suitable for all node types), that procedure will direct you to the next step for primary Admin Node recovery.

Replacement platform	Procedure
VMware	Replace a VMware node
Linux	Replace a Linux node
Services appliances	Replace a services appliance
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node .

Configure replacement primary Admin Node


The replacement node must be configured as the primary Admin Node for your StorageGRID system.

Before you begin

- For primary Admin Nodes hosted on virtual machines, the virtual machine has been deployed, powered on, and initialized.
- For primary Admin Nodes hosted on a services appliance, you have replaced the appliance and have installed software. See the [installation instructions for your appliance](#).
- You have the latest backup of the Recovery Package file (`sgws-recovery-package-id-revision.zip`).
- You have the provisioning passphrase.

Steps

1. Open your web browser and navigate to `https://primary_admin_node_ip`.
2. Manage a temporary installer password as needed:
 - If a password has already been set using one of these methods, enter the password to proceed.
 - A user set the password while accessing the installer previously
 - For bare metal systems, the password was automatically imported from the node config file at `/etc/storagegrid/nodes/<node_name>.conf`
 - For VMs, the SSH/console password was automatically imported from the OVF properties
 - If a password has not been set, optionally set a password to secure the StorageGRID installer.
3. Click **Recover a failed primary Admin Node**.



4. Upload the most recent backup of the Recovery Package:
 - a. Click **Browse**.
 - b. Locate the most recent Recovery Package file for your StorageGRID system, and click **Open**.
5. Enter the provisioning passphrase.

6. Click **Start Recovery**.

The recovery process begins. The Grid Manager might become unavailable for a few minutes as the required services start. When the recovery is complete, the sign in page is displayed.

7. If single sign-on (SSO) is enabled for your StorageGRID system and the relying party trust for the Admin Node you recovered was configured to use the default management interface certificate, update (or delete and recreate) the node's relying party trust in Active Directory Federation Services (AD FS). Use the new default server certificate that was generated during the Admin Node recovery process.



To configure a relying party trust, see [Configure single sign-on](#). To access the default server certificate, log in to the command shell of the Admin Node. Go to the `/var/local/mgmt-api` directory, and select the `server.crt` file.



After recovering a primary admin node, [determine if you need to apply a hotfix](#).

Determine hotfix requirement for primary Admin Node

After recovering a primary admin node, determine if you need to apply a hotfix.

Before you begin

Primary admin node recovery is complete.

Steps

1. Sign in to the Grid Manager using a [supported web browser](#).
2. Select **NODES**.
3. From the list on the left, select the primary Admin Node.
4. On the Overview tab, note the version displayed in the **Software Version** field.
5. Select any other grid node.
6. On the Overview tab, note the version displayed in the **Software Version** field.
 - If the versions displayed in the **Software Version** fields are the same, you don't need to apply a hotfix.
 - If the versions displayed in the **Software Version** fields are different, you must [apply a hotfix](#) to update the recovered primary Admin Node to the same version.

Restore audit log on recovered primary Admin Node

If you were able to preserve the audit log from the failed primary Admin Node, you can copy it to the primary Admin Node you are recovering.

Before you begin

- The recovered Admin Node is installed and running.
- You have copied the audit logs to another location after the original Admin Node failed.

About this task

If an Admin Node fails, audit logs saved to that Admin Node are potentially lost. It might be possible to preserve data from loss by copying audit logs from the failed Admin Node and then restoring these audit logs to the recovered Admin Node. Depending on the failure, it might not be possible to copy audit logs from the failed Admin Node. In that case, if the deployment has more than one Admin Node, you can recover audit logs

from another Admin Node as audit logs are replicated to all Admin Nodes.

If there is only one Admin Node and the audit log can't be copied from the failed node, the recovered Admin Node starts recording events to the audit log as if the installation is new.

You must recover an Admin Node as soon as possible to restore logging functionality.



By default, audit information is sent to the audit log on Admin Nodes. You can skip these steps if either of the following applies:

- You configured an external syslog server and audit logs are now being sent to the syslog server instead of to Admin Nodes.
- You explicitly specified that audit messages should be saved only on the local nodes that generated them.

See [Configure audit messages and log destinations](#) for details.

Steps

1. Log in to the recovered Admin Node:

- a. Enter the following command: `ssh admin@recovery_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.

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

2. Check which audit files have been preserved: `cd /var/local/log`

3. Copy the preserved audit log files to the recovered Admin Node: `scp admin@grid_node_IP:/var/local/tmp/saved-audit-logs/YYYY* .`

When prompted, enter the password for admin.

4. For security, delete the audit logs from the failed grid node after verifying that they have been copied successfully to the recovered Admin Node.

5. Update the user and group settings of the audit log files on the recovered Admin Node: `chown ams-user: bycast *`

6. Log out as root: `exit`

Restore Admin Node database when recovering primary Admin Node

If you want to retain the historical information about attributes and alerts on a primary Admin Node that has failed, you can restore the Admin Node database. You can only restore this database if your StorageGRID system includes another Admin Node.

Before you begin

- The recovered Admin Node is installed and running.
- The StorageGRID system includes at least two Admin Nodes.

- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

If an Admin Node fails, the historical information stored in its Admin Node database is lost. This database includes the following information:

- Alert history
- Historical attribute data, which is used in legacy-style charts on Nodes page

When you recover an Admin Node, the software installation process creates an empty Admin Node database on the recovered node. However, the new database only includes information for servers and services that are currently part of the system or added later.

If you restored a primary Admin Node and your StorageGRID system has another Admin Node, you can restore the historical information by copying the Admin Node database from a non-primary Admin Node (the *source Admin Node*) to the recovered primary Admin Node. If your system has only a primary Admin Node, you can't restore the Admin Node database.



Copying the Admin Node database might take several hours. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

Steps

1. Log in to the source Admin Node:
 - a. Enter the following command: `ssh admin@grid_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.
2. From the source Admin Node, stop the MI service: `service mi stop`
3. From the source Admin Node, stop the Management Application Program Interface (mgmt-api) service: `service mgmt-api stop`
4. Complete the following steps on the recovered Admin Node:
 - a. Log in to the recovered Admin 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.
 - b. Stop the MI service: `service mi stop`
 - c. Stop the mgmt-api service: `service mgmt-api stop`
 - d. Add the SSH private key to the SSH agent. Enter: `ssh-add`
 - e. Enter the SSH Access Password listed in the `Passwords.txt` file.
 - f. Copy the database from the source Admin Node to the recovered Admin Node:

```
/usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP
```

- g. When prompted, confirm that you want to overwrite the MI database on the recovered Admin Node.

The database and its historical data are copied to the recovered Admin Node. When the copy operation is done, the script starts the recovered Admin Node.

- h. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter: `ssh-add -D`

5. Restart the services on the source Admin Node: `service servermanager start`

Restore Prometheus metrics when recovering primary Admin Node

Optionally, you can retain the historical metrics maintained by Prometheus on a primary Admin Node that has failed. The Prometheus metrics can only be restored if your StorageGRID system includes another Admin Node.

Before you begin

- The recovered Admin Node is installed and running.
- The StorageGRID system includes at least two Admin Nodes.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

If an Admin Node fails, the metrics maintained in the Prometheus database on the Admin Node are lost. When you recover the Admin Node, the software installation process creates a new Prometheus database. After the recovered Admin Node is started, it records metrics as if you had performed a new installation of the StorageGRID system.

If you restored a primary Admin Node and your StorageGRID system has another Admin Node, you can restore the historical metrics by copying the Prometheus database from a non-primary Admin Node (the *source Admin Node*) to the recovered primary Admin Node. If your system has only a primary Admin Node, you can't restore the Prometheus database.



Copying the Prometheus database might take an hour or more. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

Steps

1. Log in to the source Admin Node:
 - a. Enter the following command: `ssh admin@grid_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.
2. From the source Admin Node, stop the Prometheus service: `service prometheus stop`
3. Complete the following steps on the recovered Admin Node:
 - a. Log in to the recovered Admin 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.
- b. Stop the Prometheus service: `service prometheus stop`
- c. Add the SSH private key to the SSH agent. Enter: `ssh-add`
- d. Enter the SSH Access Password listed in the `Passwords.txt` file.
- e. Copy the Prometheus database from the source Admin Node to the recovered Admin Node:
`/usr/local/prometheus/bin/prometheus-clone-db.sh Source_Admin_Node_IP`
- f. When prompted, press **Enter** to confirm that you want to destroy the new Prometheus database on the recovered Admin Node.

The original Prometheus database and its historical data are copied to the recovered Admin Node. When the copy operation is done, the script starts the recovered Admin Node. The following status appears:

Database cloned, starting services

- g. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter: `ssh-add -D`
4. Restart the Prometheus service on the source Admin Node. `service prometheus start`

Recover from non-primary Admin Node failures

Recover from non-primary Admin Node failures

You must complete the following tasks to recover from a non-primary Admin Node failure. One Admin Node hosts the Configuration Management Node (CMN) service and is known as the primary Admin Node. Although you can have multiple Admin Nodes, each StorageGRID system includes only one primary Admin Node. All other Admin Nodes are non-primary Admin Nodes.

Follow these high-level steps to recover a non-primary Admin Node:

1. [Copy audit logs from the failed non-primary Admin Node](#)
2. [Replace the non-primary Admin Node](#)
3. [Select Start Recovery to configure the non-primary Admin Node](#)
4. [Restore the audit log on a recovered non-primary Admin Node](#)
5. [Restore the Admin Node database when recovering a non-primary Admin Node](#)
6. [Restore Prometheus metrics when recovering a non-primary Admin Node](#)

Copy audit logs from failed non-primary Admin Node

If you are able to copy audit logs from the failed Admin Node, you should preserve them to maintain the grid's record of system activity and usage. You can restore the preserved

audit logs to the recovered non-primary Admin Node after it is up and running.

This procedure copies the audit log files from the failed Admin Node to a temporary location on a separate grid node. These preserved audit logs can then be copied to the replacement Admin Node. Audit logs aren't automatically copied to the new Admin Node.

Depending on the type of failure, you might not be able to copy audit logs from a failed Admin Node. If the deployment has only one Admin Node, the recovered Admin Node starts recording events to the audit log in a new empty file and previously recorded data is lost. If the deployment includes more than one Admin Node, you can recover the audit logs from another Admin Node.



If the audit logs aren't accessible on the failed Admin Node now, you might be able to access them later, for example, after host recovery.

1. Log in to the failed Admin Node if possible. Otherwise, log in to the primary Admin Node or another Admin Node, if available.

- a. Enter the following command: `ssh admin@grid_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 `#`.

2. Stop the AMS service to prevent it from creating a new log file: `service ams stop`
3. Navigate to the audit export directory:

```
cd /var/local/log
```

4. Rename the source `audit.log` file to a unique numbered file name. For example, rename the `audit.log` file to `2023-10-25.txt.1`.

```
ls -l
mv audit.log 2023-10-25.txt.1
```

5. Restart the AMS service: `service ams start`
6. Create the directory to copy all audit log files to a temporary location on a separate grid node: `ssh admin@grid_node_IP mkdir -p /var/local/tmp/saved-audit-logs`

When prompted, enter the password for admin.

7. Copy all audit log files to the temporary location: `scp -p * admin@grid_node_IP:/var/local/tmp/saved-audit-logs`

When prompted, enter the password for admin.

8. Log out as root: `exit`

Replace non-primary Admin Node

To recover a non-primary Admin Node, you first must replace the physical or virtual hardware.

You can replace a failed non-primary Admin Node with a non-primary Admin Node running on the same platform, or you can replace a non-primary Admin Node running on VMware or a Linux host with a non-primary Admin Node hosted on a services appliance.

Use the procedure that matches the replacement platform you select for the node. After you complete the node replacement procedure (which is suitable for all node types), that procedure will direct you to the next step for non-primary Admin Node recovery.

Replacement platform	Procedure
VMware	Replace a VMware node
Linux	Replace a Linux node
Services appliances	Replace a services appliance
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node .

Select Start Recovery to configure non-primary Admin Node

After replacing a non-primary Admin Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

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 the provisioning passphrase.
- You have deployed and configured the replacement node.

Steps

1. From the Grid Manager, select **MAINTENANCE > Tasks > Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

Nodes appear in the list after they fail, but you can't select a node until it has been reinstalled and is ready for recovery.

3. Enter the **Provisioning Passphrase**.
4. Click **Start Recovery**.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

Search 				
	Name	IPv4 Address	State	Recoverable
☒	104-217-S1	10.96.104.217	Unknown	

Passphrase

Provisioning Passphrase

Start Recovery

5. Monitor the progress of the recovery in the Recovering Grid Node table.



While the recovery procedure is running, you can click **Reset** to start a new recovery. A dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

Info

Reset Recovery

Resetting the recovery procedure leaves the deployed grid node in an indeterminate state. To retry a recovery after resetting the procedure, you must restore the node to a pre-installed state:

- For VMware nodes, delete the deployed VM and then redeploy it.
- For StorageGRID appliance nodes, run "sgareinstall" on the node.
- For Linux nodes, run "storagegrid node force-recovery *node-name*" on the Linux host.

Do you want to reset recovery?

Cancel

OK

If you want to retry the recovery after resetting the procedure, you must restore the node to a pre-installed state, as follows:

- **VMware:** Delete the deployed virtual grid node. Then, when you are ready to restart the recovery, redeploy the node.
- **Linux:** Restart the node by running this command on the Linux host: `storagegrid node force-recovery node-name`
- **Appliance:** If you want to retry the recovery after resetting the procedure, you must restore the appliance node to a pre-installed state by running `sgareinstall` on the node. See [Prepare appliance for reinstallation \(platform replacement only\)](#).

6. If single sign-on (SSO) is enabled for your StorageGRID system and the relying party trust for the Admin Node you recovered was configured to use the default management interface certificate, update (or delete and recreate) the node's relying party trust in Active Directory Federation Services (AD FS). Use the new default server certificate that was generated during the Admin Node recovery process.



To configure a relying party trust, see [Configure single sign-on](#). To access the default server certificate, log in to the command shell of the Admin Node. Go to the `/var/local/mgmt-api` directory, and select the `server.crt` file.

Restore audit log on recovered non-primary Admin Node

If you were able to preserve the audit log from the failed non-primary Admin Node, so that historical audit log information is retained, you can copy it to the non-primary Admin Node you are recovering.

Before you begin

- The recovered Admin Node is installed and running.
- You have copied the audit logs to another location after the original Admin Node failed.

About this task

If an Admin Node fails, audit logs saved to that Admin Node are potentially lost. It might be possible to preserve data from loss by copying audit logs from the failed Admin Node and then restoring these audit logs to the recovered Admin Node. Depending on the failure, it might not be possible to copy audit logs from the failed Admin Node. In that case, if the deployment has more than one Admin Node, you can recover audit logs from another Admin Node as audit logs are replicated to all Admin Nodes.

If there is only one Admin Node and the audit log can't be copied from the failed node, the recovered Admin Node starts recording events to the audit log as if the installation is new.

You must recover an Admin Node as soon as possible to restore logging functionality.

By default, audit information is sent to the audit log on Admin Nodes. You can skip these steps if either of the following applies:



- You configured an external syslog server and audit logs are now being sent to the syslog server instead of to Admin Nodes.
- You explicitly specified that audit messages should be saved only on the local nodes that generated them.

See [Configure audit messages and log destinations](#) for details.

Steps

1. Log in to the recovered Admin Node:
 - a. Enter the following command:
+
`ssh admin@recovery_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.

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

2. Check which audit files have been preserved:

```
cd /var/local/log
```

3. Copy the preserved audit log files to the recovered Admin Node:

```
scp admin@grid_node_IP:/var/local/tmp/saved-audit-logs/YYYY*
```

When prompted, enter the password for admin.

4. For security, delete the audit logs from the failed grid node after verifying that they have been copied successfully to the recovered Admin Node.

5. Update the user and group settings of the audit log files on the recovered Admin Node:

```
chown ams-user:bycast *
```

6. Log out as root: `exit`

Restore Admin Node database when recovering non-primary Admin Node

If you want to retain the historical information about attributes and alerts on a non-primary Admin Node that has failed, you can restore the Admin Node database from the primary Admin Node.

Before you begin

- The recovered Admin Node is installed and running.
- The StorageGRID system includes at least two Admin Nodes.
- You have the `Passwords.txt` file.
- You have the provisioning passphrase.

About this task

If an Admin Node fails, the historical information stored in its Admin Node database is lost. This database includes the following information:

- Alert history
- Historical attribute data, which is used in legacy-style charts on the Nodes page

When you recover an Admin Node, the software installation process creates an empty Admin Node database on the recovered node. However, the new database only includes information for servers and services that are currently part of the system or added later.

If you restored a non-primary Admin Node, you can restore the historical information by copying the Admin Node database from the primary Admin Node (the *source Admin Node*) to the recovered node.



Copying the Admin Node database might take several hours. Some Grid Manager features will be unavailable while services are stopped on the source node.

Steps

1. Log in to the source Admin Node:
 - a. Enter the following command: `ssh admin@grid_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.
2. Run the following command from the source Admin Node. Then, enter the provisioning passphrase if prompted. `recover-access-points`
3. From the source Admin Node, stop the MI service: `service mi stop`
4. From the source Admin Node, stop the Management Application Program Interface (mgmt-api) service: `service mgmt-api stop`
5. Complete the following steps on the recovered Admin Node:
 - a. Log in to the recovered Admin 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.
 - b. Stop the MI service: `service mi stop`
 - c. Stop the mgmt-api service: `service mgmt-api stop`
 - d. Add the SSH private key to the SSH agent. Enter: `ssh-add`
 - e. Enter the SSH Access Password listed in the `Passwords.txt` file.
 - f. Copy the database from the source Admin Node to the recovered Admin Node:
`/usr/local/mi/bin/mi-clone-db.sh Source_Admin_Node_IP`
 - g. When prompted, confirm that you want to overwrite the MI database on the recovered Admin Node.

The database and its historical data are copied to the recovered Admin Node. When the copy operation is done, the script starts the recovered Admin Node.
 - h. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter: `ssh-add -D`
6. Restart the services on the source Admin Node: `service servermanager start`

Restore Prometheus metrics when recovering non-primary Admin Node

Optionally, you can retain the historical metrics maintained by Prometheus on a non-primary Admin Node that has failed.

Before you begin

- The recovered Admin Node is installed and running.
- The StorageGRID system includes at least two Admin Nodes.
- You have the `Passwords.txt` file.

- You have the provisioning passphrase.

About this task

If an Admin Node fails, the metrics maintained in the Prometheus database on the Admin Node are lost. When you recover the Admin Node, the software installation process creates a new Prometheus database. After the recovered Admin Node is started, it records metrics as if you had performed a new installation of the StorageGRID system.

If you restored a non-primary Admin Node, you can restore the historical metrics by copying the Prometheus database from the primary Admin Node (the *source Admin Node*) to the recovered Admin Node.



Copying the Prometheus database might take an hour or more. Some Grid Manager features will be unavailable while services are stopped on the source Admin Node.

Steps

1. Log in to the source Admin Node:
 - a. Enter the following command: `ssh admin@grid_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.
2. From the source Admin Node, stop the Prometheus service: `service prometheus stop`
3. Complete the following steps on the recovered Admin Node:
 - a. Log in to the recovered Admin 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.
 - b. Stop the Prometheus service: `service prometheus stop`
 - c. Add the SSH private key to the SSH agent. Enter: `ssh-add`
 - d. Enter the SSH Access Password listed in the `Passwords.txt` file.
 - e. Copy the Prometheus database from the source Admin Node to the recovered Admin Node:
`/usr/local/prometheus/bin/prometheus-clone-db.sh Source_Admin_Node_IP`
 - f. When prompted, press **Enter** to confirm that you want to destroy the new Prometheus database on the recovered Admin Node.

The original Prometheus database and its historical data are copied to the recovered Admin Node. When the copy operation is done, the script starts the recovered Admin Node. The following status appears:

Database cloned, starting services

- g. When you no longer require passwordless access to other servers, remove the private key from the SSH agent. Enter: `ssh-add -D`

4. Restart the Prometheus service on the source Admin Node.`service prometheus start`

Recover from Gateway Node failures

Replace Gateway Node

You can replace a failed Gateway Node with a Gateway Node running on the same physical or virtual hardware, or you can replace a Gateway Node running on VMware or a Linux host with a Gateway Node hosted on a services appliance.

The node replacement procedure you must follow depends on which platform will be used by the replacement node. After you complete the node replacement procedure (which is suitable for all node types), that procedure will direct you to the next step for Gateway Node recovery.

Replacement platform	Procedure
VMware	Replace a VMware node
Linux	Replace a Linux node
Services appliances	Replace a services appliance
OpenStack	NetApp-provided virtual machine disk files and scripts for OpenStack are no longer supported for recovery operations. If you need to recover a node running in an OpenStack deployment, download the files for your Linux operating system. Then, follow the procedure for replacing a Linux node .

Select Start Recovery to configure Gateway Node

After replacing a Gateway Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

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 the provisioning passphrase.
- You have deployed and configured the replacement node.

Steps

1. From the Grid Manager, select **MAINTENANCE > Tasks > Recovery**.
2. Select the grid node you want to recover in the Pending Nodes list.

Nodes appear in the list after they fail, but you can't select a node until it has been reinstalled and is ready for recovery.

3. Enter the **Provisioning Passphrase**.

4. Click **Start Recovery**.

Recovery

Select the failed grid node to recover, enter your provisioning passphrase, and then click Start Recovery to begin the recovery procedure.

Pending Nodes

					Search 
	Name	IPv4 Address	State	Recoverable	
☒	104-217-S1	10.96.104.217	Unknown		

Passphrase

Provisioning Passphrase

Start Recovery

5. Monitor the progress of the recovery in the Recovering Grid Node table.



While the recovery procedure is running, you can click **Reset** to start a new recovery. A dialog box appears, indicating that the node will be left in an indeterminate state if you reset the procedure.

Info

Reset Recovery

Resetting the recovery procedure leaves the deployed grid node in an indeterminate state. To retry a recovery after resetting the procedure, you must restore the node to a pre-installed state:

- For VMware nodes, delete the deployed VM and then redeploy it.
- For StorageGRID appliance nodes, run "sgareinstall" on the node.
- For Linux nodes, run "storagegrid node force-recovery *node-name*" on the Linux host.

Do you want to reset recovery?

Cancel

OK

If you want to retry the recovery after resetting the procedure, you must restore the node to a pre-installed state, as follows:

- **VMware:** Delete the deployed virtual grid node. Then, when you are ready to restart the recovery, redeploy the node.
- **Linux:** Restart the node by running this command on the Linux host: `storagegrid node force-recovery node-name`
- **Appliance:** If you want to retry the recovery after resetting the procedure, you must restore the

appliance node to a pre-installed state by running `sgareinstall` on the node. See [Prepare appliance for reinstallation \(platform replacement only\)](#).

Recover from Archive Node failures

Recover from Archive Node failures

Support for Archive Nodes has been removed.

For information about recovering Archive Nodes, see [Recover from Archive Node failures \(StorageGRID 11.8 doc site\)](#).

Replace Linux node

Replace Linux node

If a failure requires that you deploy one or more new physical or virtual hosts or reinstall Linux on an existing host, deploy and configure the replacement host before you can recover the grid node. This procedure is one step of the grid node recovery process for all types of grid nodes.

"Linux" refers to a Red Hat® Enterprise Linux®, Ubuntu®, or Debian® deployment. For a list of supported versions, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

This procedure is only performed as one step in the process of recovering software-based Storage Nodes, primary or non-primary Admin Nodes, or Gateway Nodes. The steps are identical regardless of the type of grid node you are recovering.

If more than one grid node is hosted on a physical or virtual Linux host, you can recover the grid nodes in any order. However, recovering a primary Admin Node first, if present, prevents the recovery of other grid nodes from stalling as they try to contact the primary Admin Node to register for recovery.

Deploy new Linux hosts

With a few exceptions, you prepare the new hosts as you did during the initial installation process.

To deploy new or reinstalled physical or virtual Linux hosts, follow the procedure for preparing the hosts in the StorageGRID installation instructions for your Linux operating system:

- [Install Linux \(Red Hat Enterprise Linux\)](#)
- [Install Linux \(Ubuntu or Debian\)](#)

This procedure includes steps to accomplish the following tasks:

1. Install Linux.
2. Configure the host network.
3. Configure host storage.
4. Install the container engine.

5. Install the StorageGRID host service.



Stop after you complete the "Install StorageGRID host service" task in the installation instructions. Don't start the "Deploying grid nodes" task.

As you perform these steps, note the following important guidelines:

- Be sure to use the same host interface names you used on the original host.
- If you use shared storage to support your StorageGRID nodes, or you have moved some or all of the drives or SSDs from the failed to the replacement nodes, you must reestablish the same storage mappings that were present on the original host. For example, if you used WWIDs and aliases in `/etc/multipath.conf` as recommended in the installation instructions, be sure to use the same alias/WWID pairs in `/etc/multipath.conf` on the replacement host.
- If the StorageGRID node uses storage assigned from a NetApp ONTAP system, confirm that the volume does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

Restore grid nodes to the host

To restore a failed grid node to a new Linux host, you perform these steps to restore the node configuration file.

1. [Restore and validate the node](#) by restoring the node configuration file. For a new install, you create a node configuration file for each grid node to be installed on a host. When restoring a grid node to a replacement host, you restore or replace the node configuration file for any failed grid nodes.
2. [Start the StorageGRID host service](#).
3. As needed, [recover any nodes that fail to start](#).

If any block storage volumes were preserved from the previous host, you might have to perform additional recovery procedures. The commands in this section help you determine which additional procedures are required.

Restore and validate grid nodes

You must restore the grid configuration files for any failed grid nodes, and then validate the grid configuration files and resolve any errors.

About this task

You can import any grid node that should be present on the host, as long as its `/var/local` volume was not lost as a result of the failure of the previous host. For example, the `/var/local` volume might still exist if you used shared storage for StorageGRID system data volumes, as described in the StorageGRID installation instructions for your Linux operating system. Importing the node restores its node configuration file to the host.

If it is not possible to import missing nodes, you must re-create their grid configuration files.

You must then validate the grid configuration file, and resolve any networking or storage issues that might

occur before going on to restart StorageGRID. When you re-create the configuration file for a node, you must use the same name for the replacement node that was used for the node you are recovering.

See the installation instructions for more information about the location of the `/var/local` volume for a node.

- [Install StorageGRID on Red Hat Enterprise Linux](#)
- [Install StorageGRID on Ubuntu or Debian](#)

Steps

1. At the command line of the recovered host, list all currently configured StorageGRID nodes:

```
sudo storagegrid node list
```

If no grid nodes are configured, there will be no output. If some grid nodes are configured, expect output in the following format:

Name	Metadata-Volume
=====	
dc1-adm1	/dev/mapper/sgws-adm1-var-local
dc1-gw1	/dev/mapper/sgws-gw1-var-local
dc1-sn1	/dev/mapper/sgws-sn1-var-local
dc1-arcl	/dev/mapper/sgws-arcl-var-local

If some or all of the grid nodes that should be configured on the host aren't listed, you need to restore the missing grid nodes.

2. To import grid nodes that have a `/var/local` volume:
 - a. Run the following command for each node you want to import:

```
sudo storagegrid node import node-var-local-volume-path
```

The `storagegrid node import` command succeeds only if the target node was shut down cleanly on the host on which it last ran. If that is not the case, you will observe an error similar to the following:

```
This node (node-name) appears to be owned by another host (UUID host-uuid).
```

Use the `--force` flag if you are sure import is safe.
 - b. If you see the error about the node being owned by another host, run the command again with the `--force` flag to complete the import:

```
sudo storagegrid --force node import node-var-local-volume-path
```



Any nodes imported with the `--force` flag will require additional recovery steps before they can rejoin the grid, as described in [What's next: Perform additional recovery steps, if required](#).

3. For grid nodes that don't have a `/var/local` volume, re-create the node's configuration file to restore it to the host. For instructions, see:
 - [Create node configuration files for Red Hat Enterprise Linux](#)
 - [Create node configuration files for Ubuntu or Debian](#)



When you re-create the configuration file for a node, you must use the same name for the replacement node that was used for the node you are recovering. For Linux deployments, ensure that the configuration file name contains the node name. You should use the same network interfaces, block device mappings, and IP addresses when possible. This practice minimizes the amount of data that needs to be copied to the node during recovery, which could make the recovery significantly faster (in some cases, minutes rather than weeks).



If you use any new block devices (devices that the StorageGRID node did not use previously) as values for any of the configuration variables that start with `BLOCK_DEVICE_` when you are re-creating the configuration file for a node, follow the guidelines in [Fix missing block device errors](#).

4. Run the following command on the recovered host to list all StorageGRID nodes.

```
sudo storagegrid node list
```

5. Validate the node configuration file for each grid node whose name was shown in the `storagegrid node list` output:

```
sudo storagegrid node validate node-name
```

You must address any errors or warnings before starting the StorageGRID host service. The following sections give more detail on errors that might have special significance during recovery.

Fix missing network interface errors

If the host network is not configured correctly or a name is misspelled, an error occurs when StorageGRID checks the mapping specified in the `/etc/storagegrid/nodes/node-name.conf` file.

You might see an error or warning matching this pattern:

```
Checking configuration file /etc/storagegrid/nodes/<node-name>.conf for
node <node-name>...
ERROR: <node-name>: GRID_NETWORK_TARGET = <host-interface-name>
       <node-name>: Interface <host-interface-name>' does not exist
```

The error could be reported for the Grid Network, the Admin Network, or the Client Network. This error means that the `/etc/storagegrid/nodes/node-name.conf` file maps the indicated StorageGRID network to the host interface named `host-interface-name`, but there is no interface with that name on the current host.

If you receive this error, verify that you completed the steps in [Deploy new Linux hosts](#). Use the same names for all host interfaces as were used on the original host.

If you are unable to name the host interfaces to match the node configuration file, you can edit the node configuration file and change the value of the `GRID_NETWORK_TARGET`, the `ADMIN_NETWORK_TARGET`, or the `CLIENT_NETWORK_TARGET` to match an existing host interface.

Make sure the host interface provides access to the appropriate physical network port or VLAN, and that the interface does not directly reference a bond or bridge device. You must either configure a VLAN (or other virtual interface) on top of the bond device on the host, or use a bridge and virtual Ethernet (veth) pair.

Fix missing block device errors

The system checks that each recovered node maps to a valid block device special file or a valid softlink to a block device special file. If StorageGRID finds invalid mapping in the `/etc/storagegrid/nodes/node-name.conf` file, a missing block device error displays.

If you observe an error matching this pattern:

```
Checking configuration file /etc/storagegrid/nodes/<node-name>.conf for
node <node-name>...
ERROR: <node-name>: BLOCK_DEVICE_PURPOSE = <path-name>
      <node-name>: <path-name> does not exist
```

It means that `/etc/storagegrid/nodes/node-name.conf` maps the block device used by *node-name* for `PURPOSE` to the given `path-name` in the Linux file system, but there is not a valid block device special file, or softlink to a block device special file, at that location.

Verify that you completed the steps in [Deploy new Linux hosts](#). Use the same persistent device names for all block devices as were used on the original host.

If you are unable to restore or re-create the missing block device special file, you can allocate a new block device of the appropriate size and storage category and edit the node configuration file to change the value of `BLOCK_DEVICE_PURPOSE` to point to the new block device special file.

Determine the appropriate size and storage category using the tables for your Linux operating system:

- [Storage and performance requirements for Red Hat Enterprise Linux](#)
- [Storage and performance requirements for Ubuntu or Debian](#)

Review the recommendations for configuring host storage before proceeding with the block device replacement:

- [Configure host storage for Red Hat Enterprise Linux](#)
- [Configure host storage for Ubuntu or Debian](#)



If you must provide a new block storage device for any of the configuration file variables starting with `BLOCK_DEVICE_` because the original block device was lost with the failed host, ensure the new block device is unformatted before attempting further recovery procedures. The new block device will be unformatted if you are using shared storage and have created a new volume. If you are unsure, run the following command against any new block storage device special files.



Run the following command only for new block storage devices. Don't run this command if you believe the block storage still contains valid data for the node being recovered, as any data on the device will be lost.

```
sudo dd if=/dev/zero of=/dev/mapper/my-block-device-name bs=1G count=1
```

Start StorageGRID host service

To start your StorageGRID nodes, and ensure they restart after a host reboot, you must enable and start the

StorageGRID host service.

Steps

1. Run the following commands on each host:

```
sudo systemctl enable storagegrid
sudo systemctl start storagegrid
```

2. Run the following command to ensure the deployment is proceeding:

```
sudo storagegrid node status node-name
```

3. If any node returns a status of "Not Running" or "Stopped," run the following command:

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

4. If you have previously enabled and started the StorageGRID host service (or if you are unsure if the service has been enabled and started), also run the following command:

```
sudo systemctl reload-or-restart storagegrid
```

Recover nodes that fail to start normally

If a StorageGRID node doesn't rejoin the grid normally and doesn't show up as recoverable, it might be corrupted. You can force the node into recovery mode.

Steps

1. Confirm that the node's network configuration is correct.

The node might have failed to rejoin the grid because of incorrect network interface mappings or an incorrect Grid Network IP address or gateway.

2. If the network configuration is correct, issue the `force-recovery` command:

```
sudo storagegrid node force-recovery node-name
```

3. Perform the additional recovery steps for the node. See [What's next: Perform additional recovery steps, if required](#).

What's next: Perform additional recovery steps, if required

Depending on the specific actions you took to get the StorageGRID nodes running on the replacement host, you might need to perform additional recovery steps for each node.

Node recovery is complete if you did not need to take any corrective actions while you replaced the Linux host or restored the failed grid node to the new host.

Corrective actions and next steps

During node replacement, you might have needed to take one of these corrective actions:

- You had to use the `--force` flag to import the node.
- For any `<PURPOSE>`, the value of the `BLOCK_DEVICE_<PURPOSE>` configuration file variable refers to a block device that does not contain the same data it did before the host failure.
- You issued `storagegrid node force-recovery node-name` for the node.
- You added a new block device.

If you took **any** of these corrective actions, you must perform additional recovery steps.

Type of recovery	Next step
Primary Admin Node	Configure replacement primary Admin Node
Non-primary Admin Node	Select Start Recovery to configure non-primary Admin Node
Gateway Node	Select Start Recovery to configure Gateway Node
Storage Node (software-based): • If you had to use the <code>--force</code> flag to import the node, or you issued <code>storagegrid node force-recovery node-name</code> • If you had to do a full node reinstall, or you needed to restore <code>/var/local</code>	Select Start Recovery to configure Storage Node
Storage Node (software-based): • If you added a new block device. • If, for any <code><PURPOSE></code>, the value of the <code>BLOCK_DEVICE_<PURPOSE></code> configuration file variable refers to a block device that does not contain the same data it did before the host failure.	Recover from storage volume failure where system drive is intact

Replace VMware node

When you recover a failed StorageGRID node that was hosted on VMware, you remove the failed node and deploy a recovery node.

Before you begin

You have determined that the virtual machine can't be restored and must be replaced.

About this task

You use the VMware vSphere Web Client to first remove the virtual machine associated with the failed grid

node. Then, you can deploy a new virtual machine.

This procedure is only one step in the grid node recovery process. The node removal and deployment procedure is the same for all VMware nodes, including Admin Nodes, Storage Nodes, and Gateway Nodes.

Steps

1. Log in to VMware vSphere Web Client.
2. Navigate to the failed grid node virtual machine.
3. Make a note of all of the information required to deploy the recovery node.
 - a. Right-click the virtual machine, select the **Edit Settings** tab, and note the settings in use.
 - b. Select the **vApp Options** tab to view and record the grid node network settings.
4. If the failed grid node is a Storage Node, determine if any of the virtual hard disks used for data storage are undamaged and preserve them for reattachment to the recovered grid node.
5. Power off the virtual machine.
6. Select **Actions > All vCenter Actions > Delete from Disk** to delete the virtual machine.
7. Deploy a new virtual machine to be the replacement node, and connect it to one or more StorageGRID networks. For instructions see [Deploying a StorageGRID node as a virtual machine](#).

When you deploy the node, you can optionally remap node ports or increase CPU or memory settings.



After deploying the new node, you can add new virtual disks according to your storage requirements, reattach any virtual hard disks preserved from the previously removed failed grid node, or both.

8. Complete the node recovery procedure, based on the type of node you are recovering.

Type of node	Go to
Primary Admin Node	Configure replacement primary Admin Node
Non-primary Admin Node	Select Start Recovery to configure non-primary Admin Node
Gateway Node	Select Start Recovery to configure Gateway Node
Storage Node	Select Start Recovery to configure Storage Node

Replace failed node with services appliance

Replace failed node with services appliance

You can use a services appliance to recover a failed Gateway Node, a failed non-primary Admin Node, or a failed primary Admin Node that was hosted on VMware, a Linux host, or a services appliance. This procedure is one step of the grid node recovery procedure.

Before you begin

- You have determined that one of the following situations is true:

- The virtual machine hosting the node can't be restored.
- The physical or virtual Linux host for the grid node has failed, and must be replaced.
- The services appliance hosting the grid node must be replaced.
- You have confirmed that the StorageGRID Appliance Installer version on the services appliance matches the software version of your StorageGRID system. See [Verify and upgrade StorageGRID Appliance Installer version](#).



Don't deploy both an SG110 and an SG1100 services appliance or an SG100 and an SG1000 services appliance in the same site. Unpredictable performance might result.

About this task

You can use a services appliance to recover a failed grid node in the following cases:

- The failed node was hosted on VMware or Linux ([platform change](#))
- The failed node was hosted on a services appliance ([platform replacement](#))

Install services appliance (platform change only)

When you are recovering a failed grid node that was hosted on VMware or a Linux host and you are using a services appliance for the replacement node, you must first install the new appliance hardware using the same node name (system name) as the failed node.

Before you begin

You have the following information about the failed node:

- **Node name:** You must install the services appliance using the same node name as the failed node. The node name is the hostname (system name).
- **IP addresses:** You can assign the services appliance the same IP addresses as the failed node, which is the preferred option, or you can select a new unused IP address on each network.

About this task

Perform this procedure only if you are recovering a failed node that was hosted on VMware or Linux and are replacing it with a node hosted on a services appliance.

Steps

1. Follow the instructions for installing a new services appliance. See [Quick start for hardware installation](#).
2. When prompted for a node name, use the node name of the failed node.

Prepare appliance for reinstallation (platform replacement only)

When recovering a grid node that was hosted on a services appliance, you must first prepare the appliance for reinstallation of StorageGRID software.

Perform this procedure only if you are replacing a failed node that was hosted on a services appliance. Don't follow these steps if the failed node was originally hosted on VMware or a Linux host.

Steps

1. Log in to the failed grid node:

- a. Enter the following command: `ssh admin@grid_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 `#`.

2. Prepare the appliance for the installation of StorageGRID software. Enter: `sgareinstall`

3. When prompted to continue, enter: `y`

The appliance reboots, and your SSH session ends. It usually takes about 5 minutes for the StorageGRID Appliance Installer to become available, although in some cases you might need to wait up to 30 minutes.

The services appliance is reset, and data on the grid node is no longer accessible. IP addresses configured during the original installation process should remain intact; however, it is recommended that you confirm this when the procedure completes.

After executing the `sgareinstall` command, all StorageGRID-provisioned accounts, passwords, and SSH keys are removed, and new host keys are generated.

Start software installation on services appliance

To install a Gateway Node or Admin Node on a services appliance, you use the StorageGRID Appliance Installer, which is included on the appliance.

Before you begin

- The appliance is installed in a rack, connected to your networks, and powered on.
- Network links and IP addresses are configured for the appliance using the StorageGRID Appliance Installer.
- If you are installing a Gateway Node or non-primary Admin Node, you know the IP address of the primary Admin Node for the StorageGRID grid.
- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer are defined in the Grid Network Subnet List on the primary Admin Node.

See [Quick start for hardware installation](#).

- You are using a [supported web browser](#).
- You have one of the IP addresses assigned to the appliance. You can use the IP address for the Admin Network, the Grid Network, or the Client Network.
- If you are installing a primary Admin Node, you have the Ubuntu or Debian install files for this version of StorageGRID available.



A recent version of StorageGRID software is preloaded onto the services appliance during manufacturing. If the preloaded version of software matches the version being used in your StorageGRID deployment, you don't need the installation files.

About this task

To install StorageGRID software on a services appliance:

- For a primary Admin Node, you specify the name of the node and then upload the appropriate software packages (if required).
- For a non-primary Admin Node or a Gateway Node, you specify or confirm the IP address of the primary Admin Node and the name of the node.
- You start the installation and wait as volumes are configured and the software is installed.
- Partway through the process, the installation pauses. To resume the installation, you must sign into the Grid Manager and configure the pending node as a replacement for the failed node.
- After you have configured the node, the appliance installation process completes, and the appliance is rebooted.

Steps

1. Open a browser and enter one of the IP addresses for the services appliance.

```
https://Controller_IP:8443
```

The StorageGRID Appliance Installer Home page appears.

NetApp® StorageGRID® Appliance Installer
Help

Home
Configure Networking
Configure Hardware
Monitor Installation
Advanced

Home

This Node

Node type
Gateway

Node name
NetApp-SGA

Cancel
Save

Primary Admin Node connection

Enable Admin Node discovery
☒
Uncheck to manually enter the Primary Admin Node IP

Connection state
Admin Node discovery is in progress

Cancel
Save

Installation

Current state
Unable to start installation. The Admin Node connection is not ready.

Start Installation

2. To install a Primary Admin Node:
 - a. In the This Node section, for **Node Type**, select **Primary Admin**.
 - b. In the **Node Name** field, enter the same name that was used for the node you are recovering, and click **Save**.
 - c. In the Installation section, check the software version listed under Current state

If the version of software that is ready to install is correct, skip ahead to the [Installation step](#).
 - d. If you need to upload a different version of software, under the **Advanced** menu, select **Upload StorageGRID Software**.

The Upload StorageGRID Software page appears.

Upload StorageGRID Software

If this node is the primary Admin Node of a new deployment, you must use this page to upload the StorageGRID software installation package, unless the version of the software you want to install has already been uploaded. If you are adding this node to an existing deployment, you can avoid network traffic by uploading the installation package that matches the software version running on the existing grid. If you do not upload the correct package, the node obtains the software from the grid's primary Admin Node during installation.

Current StorageGRID Installation Software

Version	None
Package Name	None

Upload StorageGRID Installation Software

Software Package	Browse
Checksum File	Browse

e. Click **Browse** to upload the **Software Package** and **Checksum File** for StorageGRID software.

The files are automatically uploaded after you select them.

f. Click **Home** to return to the StorageGRID Appliance Installer Home page.

3. To install a Gateway Node or non-Primary Admin Node:

- In the This Node section, for **Node Type**, select **Gateway** or **Non-Primary Admin**, depending on the type of node you are restoring.
- In the **Node Name** field, enter the same name that was used for the node you are recovering, and click **Save**.
- In the Primary Admin Node connection section, determine whether you need to specify the IP address for the primary Admin Node.

The StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

d. If this IP address is not shown or you need to change it, specify the address:

Option	Description
Manual IP entry	- Clear the Enable Admin Node discovery checkbox. - Enter the IP address manually. - Click Save. - Wait while the connection state for the new IP address becomes "ready."

Option	Description
Automatic discovery of all connected primary Admin Nodes	- Select the Enable Admin Node discovery checkbox. - From the list of discovered IP addresses, select the primary Admin Node for the grid where this services appliance will be deployed. - Click Save. - Wait while the connection state for the new IP address becomes "ready."

- In the Installation section, confirm that the current state is Ready to start installation of node name and that the **Start Installation** button is enabled.

If the **Start Installation** button is not enabled, you might need to change the network configuration or port settings. For instructions, see the maintenance instructions for your appliance.

- From the StorageGRID Appliance Installer home page, click **Start Installation**.

The Current state changes to "Installation is in progress," and the Monitor Installation page is displayed.



If you need to access the Monitor Installation page manually, click **Monitor Installation** from the menu bar.

Monitor services appliance installation

The StorageGRID Appliance Installer provides status until installation is complete. When the software installation is complete, the appliance is rebooted.

Steps

- To monitor the installation progress, click **Monitor Installation** from the menu bar.

The Monitor Installation page shows the installation progress.

Monitor Installation

1. Configure storage		Complete
2. Install OS		Running
Step	Progress	Status
Obtain installer binaries		Complete
Configure installer		Complete
Install OS		Installer VM running
3. Install StorageGRID		Pending
4. Finalize installation		Pending

The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.



The installer ensures that tasks completed in a previous install aren't re-run. If you are re-running an installation, any tasks that don't need to be re-run are shown with a green status bar and a status of "Skipped."

2. Review the progress of first two installation stages.

◦ 1. Configure storage

During this stage, the installer clears any existing configuration from the drives, and configures host settings.

◦ 2. Install OS

During this stage, the installer copies the base operating system image for StorageGRID from the primary Admin Node to the appliance or installs the base operating system from the installation package for the primary Admin Node.

3. Continue monitoring the installation progress until one of the following occurs:

- For appliance Gateway Nodes or non-primary appliance Admin Nodes, the **Install StorageGRID** stage pauses and a message appears on the embedded console, prompting you to approve this node on the Admin Node using the Grid Manager.

Home

Configure Networking ▾

Configure Hardware ▾

Monitor Installation

Advanced ▾

Monitor Installation

1. Configure storage	Complete
2. Install OS	Complete
3. Install StorageGRID	Running
4. Finalize installation	Pending

Connected (unencrypted) to: QEMU

```

/platform.type: Device or resource busy
[2017-07-31T22:09:12.362566] INFO -- [INSG] NOTICE: seeding /var/local with c
ontainer data
[2017-07-31T22:09:12.366205] INFO -- [INSG] Fixing permissions
[2017-07-31T22:09:12.369633] INFO -- [INSG] Enabling syslog
[2017-07-31T22:09:12.511533] INFO -- [INSG] Stopping system logging: syslog-n
g.
[2017-07-31T22:09:12.570096] INFO -- [INSG] Starting system logging: syslog-n
g.
[2017-07-31T22:09:12.576360] INFO -- [INSG] Beginning negotiation for downloa
d of node configuration
[2017-07-31T22:09:12.581363] INFO -- [INSG]
[2017-07-31T22:09:12.585066] INFO -- [INSG]
[2017-07-31T22:09:12.588314] INFO -- [INSG]
[2017-07-31T22:09:12.591851] INFO -- [INSG]
[2017-07-31T22:09:12.594886] INFO -- [INSG]
[2017-07-31T22:09:12.598360] INFO -- [INSG]
[2017-07-31T22:09:12.601324] INFO -- [INSG]
[2017-07-31T22:09:12.604759] INFO -- [INSG]
[2017-07-31T22:09:12.607800] INFO -- [INSG]
[2017-07-31T22:09:12.610985] INFO -- [INSG]
[2017-07-31T22:09:12.614597] INFO -- [INSG]
[2017-07-31T22:09:12.618282] INFO -- [INSG] Please approve this node on the A
dmin Node GMI to proceed...

```

- For appliance primary Admin Nodes, a fifth phase (Load StorageGRID Installer) appears. If the fifth phase is in progress for more than 10 minutes, refresh the page manually.

NetApp® StorageGRID® Appliance Installer
Help

Home
Configure Networking
Configure Hardware
Monitor Installation
Advanced

Monitor Installation

1. Configure storage	Complete
2. Install OS	Complete
3. Install StorageGRID	Complete
4. Finalize installation	Complete
5. Load StorageGRID Installer	Running

Step	Progress	Status
Starting StorageGRID Installer		Do not refresh. You will be redirected when the installer is ready

4. Go to the next step of the recovery process for the type of appliance grid node that you are recovering.

Type of recovery	Reference
Gateway Node	Select Start Recovery to configure Gateway Node
Non-primary Admin Node	Select Start Recovery to configure non-primary Admin Node
Primary Admin Node	Configure replacement primary Admin Node

How technical support recovers a site

If an entire StorageGRID site fails or if multiple Storage Nodes fail, you must contact technical support. Technical support will assess your situation, develop a recovery plan, and then recover the failed nodes or site in a way that meets your business objectives, optimizes recovery time, and prevents unnecessary data loss.



Site recovery can only be performed by technical support.

StorageGRID systems are resilient to a wide variety of failures, and you can successfully perform many recovery and maintenance procedures yourself. However, it is difficult to create a simple, generalized site recovery procedure because the detailed steps depend on factors that are specific to your situation. For example:

- Your business objectives:** After the complete loss of a StorageGRID site, you should evaluate how best to meet your business objectives. For example, do you want to rebuild the lost site in-place? Do you want to replace the lost StorageGRID site in a new location? Every customer’s situation is different, and your recovery plan must be designed to address your priorities.
- Exact nature of the failure:** Before beginning a site recovery, establish if any nodes at the failed site are intact or if any Storage Nodes contain recoverable objects. If you rebuild nodes or storage volumes that contain valid data, unnecessary data loss could occur.
- Active ILM policies:** The number, type, and location of object copies in your grid is controlled by your active ILM policies. The specifics of your ILM policies can affect the amount of recoverable data, as well as

the specific techniques required for recovery.



If a site contains the only copy of an object and the site is lost, the object is lost.

- **Bucket (or container) consistency:** The consistency applied to a bucket (or container) affects whether StorageGRID fully replicates object metadata to all nodes and sites before telling a client that object ingest was successful. If the consistency value allows for eventual consistency, some object metadata might have been lost in the site failure. This can affect the amount of recoverable data and potentially the details of the recovery procedure.
- **History of recent changes:** The details of your recovery procedure can be affected by whether any maintenance procedures were in progress at the time of the failure or whether any recent changes were made to your ILM policies. Technical support must assess the recent history of your grid as well as its current situation before beginning a site recovery.



Site recovery can only be performed by technical support.

This is a general overview of the process that technical support uses to recover a failed site:

1. Technical support:
 - a. Makes a detailed assessment of the failure.
 - b. Works with you to review your business objectives.
 - c. Develops a recovery plan tailored for your situation.
2. If the primary Admin Node if it has failed, technical support recovers it.
3. Technical support recovers all Storage Nodes, following this outline:
 - a. Replace Storage Node hardware or virtual machines as required.
 - b. Restore object metadata to the failed site.
 - c. Restore object data to the recovered Storage Nodes.



Data loss will occur if the recovery procedures for a single failed Storage Node are used.



When an entire site has failed, technical support uses specialized commands to successfully restore objects and object metadata.

4. Technical support recovers other failed nodes.

After object metadata and data have been recovered, technical support uses standard procedures to recover failed Gateway Nodes or non-primary Admin Nodes.

Related information

[Site decommission](#)

How to enable StorageGRID in your environment

Go to [How to enable StorageGRID](#) to learn how to test and enable applications in your StorageGRID environment.

How to manage StorageGRID using BlueXP

Go to [StorageGRID management using BlueXP](#) to learn how to manage your StorageGRID systems from BlueXP using Grid Manager and use BlueXP's data services for backups, data tiering, and more.

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<https://www.netapp.com/company/legal/copyright/>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

https://library.netapp.com/ecm/ecm_download_file/ECMLP3330669

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.