



# StorageGRID 11.5 Documentation

## StorageGRID 11.5

NetApp  
November 07, 2024

# Table of Contents

- StorageGRID 11.5 Documentation . . . . . 1
- Release notes . . . . . 2
- Get started . . . . . 3
  - Grid primer . . . . . 3
  - Networking guidelines . . . . . 71
- Install and upgrade software . . . . . 100
  - Install Red Hat Enterprise Linux or CentOS . . . . . 100
  - Install Ubuntu or Debian . . . . . 169
  - Install VMware . . . . . 238
  - Upgrade software . . . . . 287
- Install and maintain hardware . . . . . 327
  - SG6000 storage appliances . . . . . 327
  - SG5700 storage appliances . . . . . 496
  - SG5600 storage appliances . . . . . 616
  - SG100 & SG1000 services appliances . . . . . 730
- Configure and manage . . . . . 836
  - Administer StorageGRID . . . . . 836
  - Manage objects with ILM . . . . . 1103
  - System hardening . . . . . 1266
  - Configure StorageGRID for FabricPool . . . . . 1274
- Use StorageGRID . . . . . 1293
  - Use a tenant account . . . . . 1293
  - Use S3 . . . . . 1394
  - Use Swift . . . . . 1515
- Monitor and troubleshoot . . . . . 1547
  - Monitor a StorageGRID system . . . . . 1547
  - Troubleshoot a StorageGRID system . . . . . 1829
  - Review audit logs . . . . . 1887
- Maintain . . . . . 1976
  - Expand your grid . . . . . 1976
  - Maintain & recover . . . . . 2029
- Other versions of NetApp StorageGRID documentation . . . . . 2263
- Legal notices . . . . . 2264
  - Copyright . . . . . 2264
  - Trademarks . . . . . 2264
  - Patents . . . . . 2264
  - Privacy policy . . . . . 2264
  - Open source . . . . . 2264

# StorageGRID 11.5 Documentation

# Release notes

Obtain release-specific information about new features, removed and deprecated features, fixed issues, and known issues.

Release Notes are available outside this doc site. You will be prompted to log in using your NetApp Support Site credentials.

- [HTML](#)
- [PDF](#)

# Get started

## Grid primer

Learn the basics of a NetApp StorageGRID system.

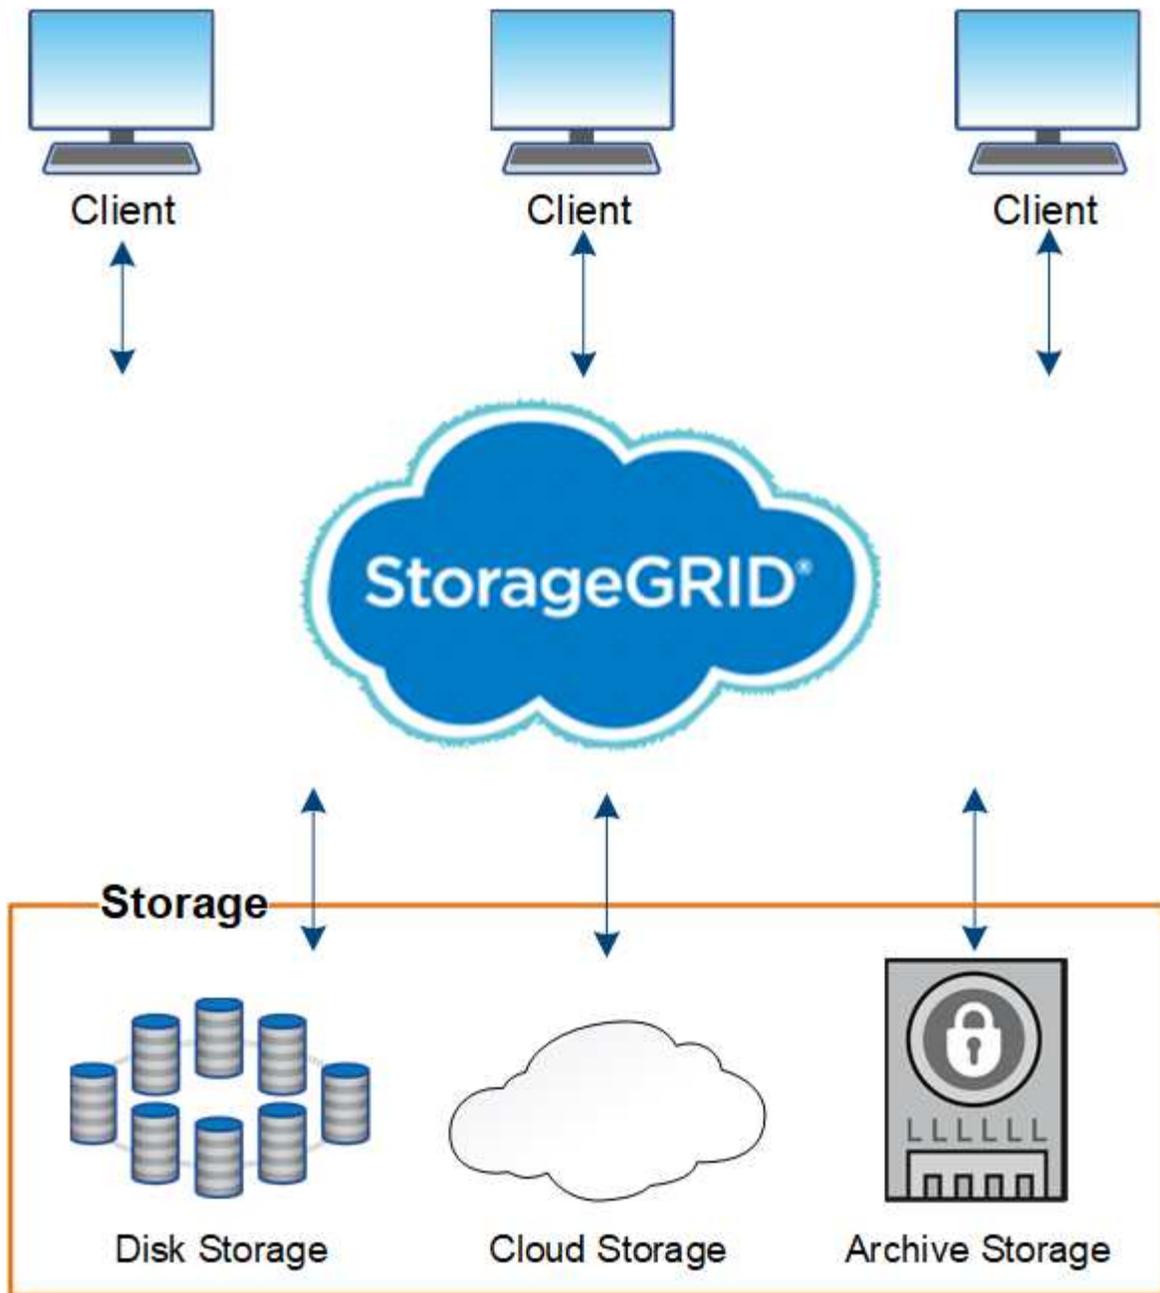
- [About StorageGRID](#)
- [StorageGRID architecture and network topology](#)
- [How StorageGRID manages data](#)
- [Exploring the Grid Manager](#)
- [Exploring the Tenant Manager](#)
- [Using StorageGRID](#)

### About StorageGRID

NetApp StorageGRID is a software-defined, object-based storage solution that supports industry-standard object APIs, including the Amazon Simple Storage Service (S3) API and the OpenStack Swift API.

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.



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
- 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 on 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, and customize data locality, durability, performance, cost, and retention

time. Tape is available as an integrated archive tier.

- 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.
  - Docker containers 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).

## **Related information**

[Hybrid clouds with StorageGRID](#)

[StorageGRID architecture and network topology](#)

[Controlling StorageGRID access](#)

[Managing tenants and client connections](#)

[Using information lifecycle management](#)

[Monitoring StorageGRID operations](#)

[Configuring network settings](#)

[Performing maintenance procedures](#)

## **Hybrid clouds with StorageGRID**

You can use StorageGRID in a hybrid cloud configuration by implementing policy-driven data management to store objects in Cloud Storage Pools, by leveraging StorageGRID platform services, and by moving data 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, 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.



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 (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 with StorageGRID

You can reduce the cost of ONTAP storage by tiering data to StorageGRID using FabricPool. FabricPool is a NetApp Data Fabric technology that 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

[Administer StorageGRID](#)

[Use a tenant account](#)

[Manage objects with ILM](#)

[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.

For additional information about StorageGRID network topology, requirements, and grid communications, see the networking guidelines.

## Related information


[Network guidelines](#)

## Deployment topologies

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

### Single site

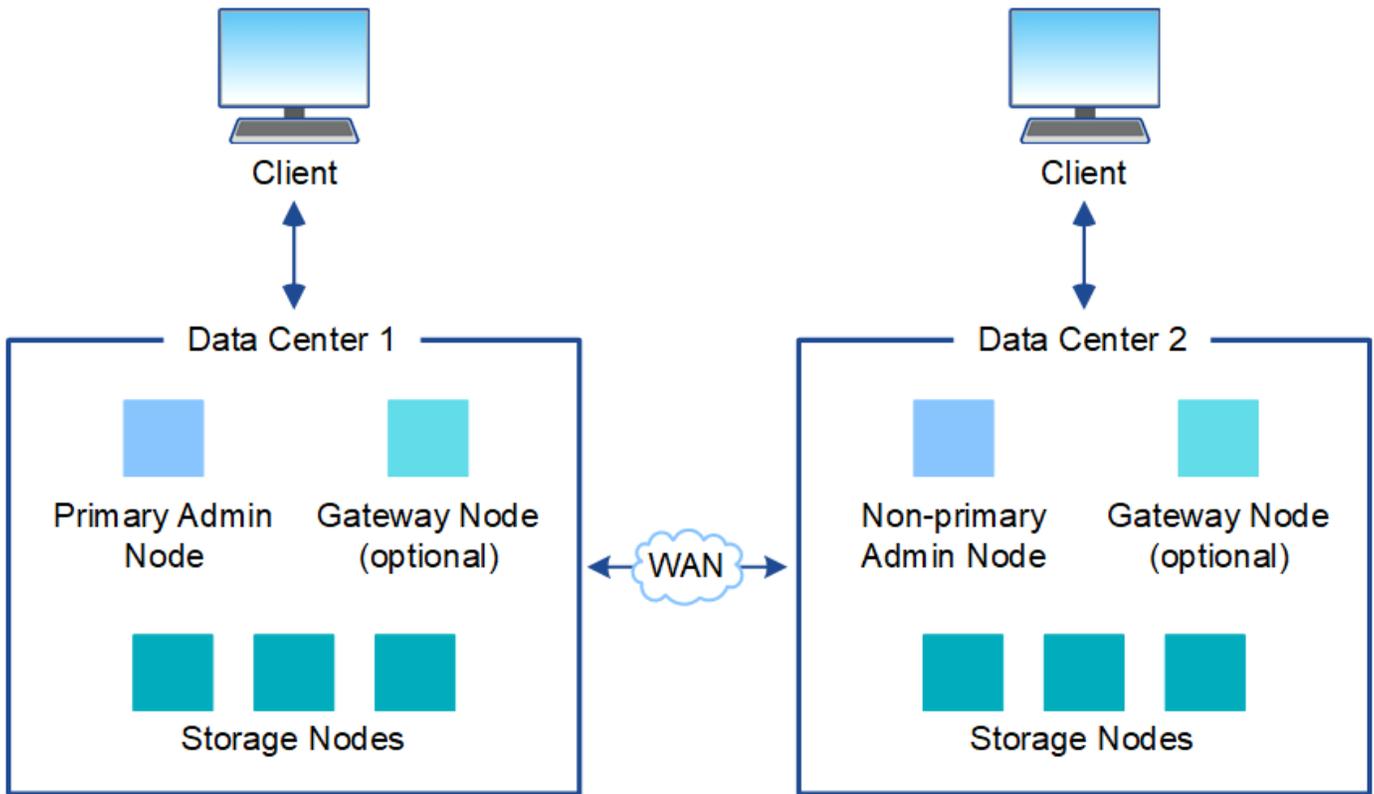
In a deployment with a single site, the infrastructure and operations of the StorageGRID system are centralized.



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

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

**Related information**

[Administer StorageGRID](#)

## 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.

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 and Swift client traffic.

- **Storage Nodes** manage and store object data and metadata. Each StorageGRID system must have at least three Storage Nodes. If you have multiple sites, each site within your StorageGRID system must also have three Storage Nodes.
- **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. You can use a combination of Gateway Nodes and Admin Nodes for load balancing, or you can implement a third-party HTTP load balancer.
- **Archive Nodes (optional)** provide an interface through which object data can be archived to tape.

### Software-based nodes

Software-based grid nodes can be deployed in the following ways:

- As virtual machines (VMs) in VMware vSphere Web Client
- Within Docker containers on Linux hosts. The following operating systems are supported:
  - Red Hat Enterprise Linux
  - CentOS
  - Ubuntu
  - Debian

Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

### 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.

Four types of StorageGRID appliances are available:

- The **SG100 and SG1000 services appliances** are 1-rack-unit (1U) servers that can each operate as the primary Admin Node, a non-primary Admin Node, or a Gateway Node. Both appliances can operate as Gateway Nodes and Admin Nodes (primary and non-primary) at the same time.
- The **SG6000 storage appliance** operates as a Storage Node and combines the 1U SG6000-CN compute controller with a 2U or 4U storage controller shelf. The SG6000 is available in two models:

- **SGF6024**: Combines the SG6000-CN compute controller with a 2U storage controller shelf that includes 24 solid state drives (SSDs) and redundant storage controllers.
- **SG6060**: Combines the SG6000-CN compute controller with a 4U enclosure that includes 58 NL-SAS drives, 2 SSDs, and redundant storage controllers. Each SG6060 appliance supports one or two 60-drive expansion shelves, providing up to 178 drives dedicated to object storage.
- The **SG5700 storage appliance** is an integrated storage and computing platform that operates as a Storage Node. The SG5700 is available in two models:
  - **SG5712**: a 2U enclosure that includes 12 NL-SAS drives and integrated storage and compute controllers.
  - **SG5760**: a 4U enclosure that includes 60 NL-SAS drives and integrated storage and compute controllers.
- The **SG5600 storage appliance** is an integrated storage and computing platform that operates as a Storage Node. The SG5600 is available in two models:
  - **SG5612**: a 2U enclosure that includes 12 NL-SAS drives and integrated storage and compute controllers.
  - **SG5660**: a 4U enclosure that includes 60 NL-SAS drives and integrated storage and compute controllers.

See the NetApp Hardware Universe for complete specifications.

#### 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.
Configuration Management Node (CMN)	Manages system-wide configuration. Primary Admin Node only.
Management Application Program Interface (mgmt-api)	Processes requests from the Grid Management API and the Tenant Management API.
High Availability	Manages high availability virtual IP addresses for groups of Admin Nodes and Gateway Nodes.  <b>Note:</b> This service is also found on Gateway Nodes.
Load Balancer	Provides load balancing of S3 and Swift traffic from clients to Storage Nodes.  <b>Note:</b> This service is also found on Gateway Nodes.
Network Management System (NMS)	Provides functionality for the Grid Manager.

Service	Key function
Prometheus	Collects and stores metrics.
Server Status Monitor (SSM)	Monitors the operating system and underlying hardware.

### 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.
Cassandra	Stores and protects object metadata.
Cassandra Reaper	Performs automatic repairs of object metadata.
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.
Identity (idnt)	Federates user identities from LDAP and Active Directory.
Local Distribution Router (LDR)	Processes object storage protocol requests and manages object data on disk.
Replicated State Machine (RSM)	Ensures that S3 platform service requests are sent to their respective endpoints.
Server Status Monitor (SSM)	Monitors the operating system and underlying hardware.

### 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
Connection Load Balancer (CLB)	Provides Layers 3 and 4 load balancing of S3 and Swift traffic from clients to Storage Nodes. Legacy load balancing mechanism.  <b>Note:</b> The CLB service is deprecated.
High Availability	Manages high availability virtual IP addresses for groups of Admin Nodes and Gateway Nodes.  <b>Note:</b> This service is also found on Admin Nodes.
Load Balancer	Provides Layer 7 load balancing of S3 and Swift traffic from clients to Storage Nodes. This is the recommended load balancing mechanism.  <b>Note:</b> This service is also found on Admin Nodes.
Server Status Monitor (SSM)	Monitors the operating system and underlying hardware.

#### Primary services for Archive Nodes

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

Service	Key function
Archive (ARC)	Communicates with a Tivoli Storage Manager (TSM) external tape storage system.
Server Status Monitor (SSM)	Monitors the operating system and underlying hardware.

#### StorageGRID services

The following is a complete list of StorageGRID services.

- **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. The Load Balancer service is present on Admin Nodes and Gateway Nodes.

- **ADC service (Administrative Domain Controller)**

Maintains topology information, provides authentication services, and responds to queries from the LDR and CMN services. The ADC service is present on each of the first three Storage Nodes installed at a site.

- **AMS service (Audit Management System)**

Monitors and logs all audited system events and transactions to a text log file. The AMS service is present on Admin Nodes.

- **ARC service (Archive)**

Provides the management interface with which you configure connections to external archival storage, such as the cloud through an S3 interface or tape through TSM middleware. The ARC service is present on Archive Nodes.

- **Cassandra Reaper service**

Performs automatic repairs of object metadata. The Cassandra Reaper service is present on all Storage Nodes.

- **Chunk service**

Manages erasure-coded data and parity fragments. The Chunk service is present on Storage Nodes.

- **CLB service (Connection Load Balancer)**

Deprecated service that provides a gateway into StorageGRID for client applications connecting through HTTP. The CLB service is present on Gateway Nodes. The CLB service is deprecated and will be removed in a future StorageGRID release.

- **CMN service (Configuration Management Node)**

Manages system-wide configurations and grid tasks. Each grid has one CMN service, which is present on the primary Admin Node.

- **DDS service (Distributed Data Store)**

Interfaces with the Cassandra database to manage object metadata. The DDS service is present on Storage Nodes.

- **DMV service (Data Mover)**

Moves data to cloud endpoints. The DMV service is present on Storage Nodes.

- **Dynamic IP service**

Monitors the grid for dynamic IP changes and updates local configurations. The Dynamic IP (dynip) service is present on all nodes.

- **Grafana service**

Used for metrics visualization in the Grid Manager. The Grafana service is present on Admin Nodes.

- **High Availability service**

Manages high availability Virtual IPs on nodes configured on the High Availability Groups page. The High Availability service is present on Admin Nodes and Gateway Nodes. This service is also known as the keepalived service.

- **Identity (idnt) service**

Federates user identities from LDAP and Active Directory. The Identity service (idnt) is present on three Storage Nodes at each site.

- **Load Balancer service**

Provides load balancing of S3 and Swift traffic from clients to Storage Nodes. The Load Balancer service

can be configured through the Load Balancer Endpoints configuration page. The Load Balancer service is present on Admin Nodes and Gateway Nodes. This service is also known as the nginx-gw service.

- **LDR service (Local Distribution Router)**

Manages the storage and transfer of content within the grid. The LDR service is present on Storage Nodes.

- **MISCd Information Service Control Daemon service**

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. The MISCd service is present on all nodes.

- **nginx service**

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. The nginx service is present on all nodes.

- **nginx-gw service**

Powers the Load Balancer service. The nginx-gw service is present on Admin Nodes and Gateway Nodes.

- **NMS service (Network Management System)**

Powers the monitoring, reporting, and configuration options that are displayed through the Grid Manager. The NMS service is present on Admin Nodes.

- **Persistence service**

Manages files on the root disk that need to persist across a reboot. The Persistence service is present on all nodes.

- **Prometheus service**

Collects time series metrics from services on all nodes. The Prometheus service is present on Admin Nodes.

- **RSM service (Replicated State Machine Service)**

Ensures platform service requests are sent to their respective endpoints. The RSM service is present on Storage Nodes that use the ADC service.

- **SSM service (Server Status Monitor)**

Monitors hardware conditions and reports to the NMS service. An instance of the SSM service is present on every grid node.

- **Trace collector service**

Performs trace collection to gather information for use by technical support. The trace collector service uses open source Jaeger software and is present on Admin Nodes.

## Related information

[NetApp Interoperability Matrix Tool](#)

[NetApp Hardware Universe](#)

[Install VMware](#)

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

[Install Ubuntu or Debian](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[Administer StorageGRID](#)

## How StorageGRID manages data


As you begin working with the StorageGRID system, it is helpful to understand how the StorageGRID system manages data.

- [What an object is](#)
- [How object data is protected](#)
- [The life of an object](#)

### What an object is

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.



### Object data

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

## 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 [Managing object metadata storage](#).


## How object data is 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, Archive 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](#)


[Using 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 or Swift 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 policy. 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, in a Cloud Storage Pool, or on Archive Node.
- **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, Cloud Storage Pool, or Archive Node.
- **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](#)


[Using 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 saves 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 policy and its 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.

### **Data flow**

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.

### Data flow

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. Once 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. Once broken into 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 moved to the external S3 bucket or Azure Blob storage container that was specified for the Cloud Storage Pool.

### Data flow

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.

### Content protection: archive

An archive operation consists of a defined data flow between the StorageGRID system and the client.

If the ILM policy requires that a copy of object data be archived, the ILM engine, which is a component of the LDR service, sends a request to the Archive Node, which in turn sends a copy of the object data to the targeted archival storage system.



1. The ILM engine sends a request to the ARC service to store a copy on archive media.

2. The ARC service queries the ADC service for the best source location and sends a request to the source LDR service.
3. The ARC service retrieves object data from the LDR service.
4. The ARC service sends the object data to the archive media destination.
5. The archive media notifies the ARC service that the object data has been stored.
6. The ARC service notifies the ILM engine that the object data has been stored.


### 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 or Archive Node.

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 or to an Archive Node.



If the only object copy is on AWS Glacier storage or the Azure Archive tier, the client application must issue an S3 POST Object restore 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 cannot be deleted by any method.
  - Before an object version's retain-until-date is reached, that version cannot 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.
- **Client delete request:** An S3 or Swift client can issue a delete object request. When a client deletes an object, all copies of the object are removed 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.




The Expiration action in an S3 bucket lifecycle always overrides ILM settings. As a result, an object might be retained on the grid even after any ILM instructions for placing the object have lapsed.

## 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.

## Exploring the Grid Manager

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

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.

### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84


You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## 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 includes summary information about system health, storage use, ILM processes, and S3 and Swift operations.




For an explanation of the information on each panel, click the help icon  for that panel.

## Related information

[Monitor & troubleshoot](#)

## 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:

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

#### **Related information**

[Monitoring and managing alerts](#)

[Monitor & troubleshoot](#)

#### **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, click the appropriate link on the left.

Dashboard

Alerts ▾

Nodes

Tenants

ILM ▾

Configuration ▾

Maintenance ▾

Support ▾

## StorageGRID Deployment

## Data Center 1

- ✓ DC1-ADM1
- ✓ DC1-ARC1
- ✓ DC1-G1
- ✓ DC1-S1
- ✓ DC1-S2
- ✓ DC1-S3

## Data Center 2

- ✓ DC2-ADM1
- ✓ DC2-S1
- ✓ DC2-S2
- ✓ DC2-S3

## Data Center 3

- ✓ DC3-S1
- ✓ DC3-S2
- ✓ DC3-S3

## StorageGRID Deployment

Network

Storage

Objects

ILM

Load Balancer


1 hour

1 day

1 week

1 month

Custom

**Related information**[Viewing the Nodes page](#)[Monitor & troubleshoot](#)**Tenant Accounts page**

The Tenant Accounts 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 Tenant Accounts 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.

### Tenant Accounts

View information for each tenant account.

**Note:** 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 and select **View Details**.

Display Name	Space Used	Quota Utilization	Quota	Object Count	Sign in
S3 tenant	0 bytes	0.00%	100.00 GB	0	
Swift tenant	0 bytes	0.00%	100.00 GB	0	

Show 20 rows per page

### Related information

[Managing tenants and client connections](#)

[Administer StorageGRID](#)

[Use a tenant account](#)

### 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.

Dashboard Alerts Nodes Tenants **ILM** Configuration Maintenance Support

Storage Pools

**Storage Pools**

A storage pool is a logical group of Storage Nodes or Archive Nodes that determine where object data is stored.

Pool Name	Archive Nodes	Storage Nodes	ILM Rule	Used in EC Profile
All Storage Nodes	0	5		
3 sites	0	9		

Displaying 2 pools.

### Related information

[Using information lifecycle management](#)

[Manage objects with ILM](#)

## Configuration menu

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

Configuration ▾	Maintenance ▾	Support ▾	
<b>Network Settings</b>	<b>System Settings</b>	<b>Monitoring</b>	<b>Access Control</b>
Domain Names	Display Options	Audit	Identity Federation
High Availability Groups	Grid Options	Events	Admin Groups
Link Cost	Key Management Server	SNMP Agent	Admin Users
Load Balancer Endpoints	S3 Object Lock		Single Sign-on
Proxy Settings	Storage Options		Client Certificates
Server Certificates			Grid Passwords
Traffic Classification			
Untrusted Client Network			

### Related information

[Configuring network settings](#)

[Managing tenants and client connections](#)

[Reviewing audit messages](#)

[Controlling StorageGRID access](#)


[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

[Review audit logs](#)

### Maintenance menu

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



## Decommission

Select **Decommission Nodes** to remove one or more nodes from a single site. Select **Decommission Site** to remove a site from the grid. Learn important details about removing grid nodes and sites in the "Decommission procedures" section.



## Maintenance Tasks

Maintenance tasks include:

- Decommission operations to remove unused grid nodes and sites.
- Expansion operations to add new grid nodes and sites.
- Recovery operations to replace a failed node and restore data.

## Network

Network tasks you can perform from the Maintenance menu include:

- Editing information about DNS servers.
- Configuring the subnets that are used on the Grid Network.
- Editing information about NTP servers.

## System

System tasks you can perform from the Maintenance menu include:

- Reviewing details for the current StorageGRID license or uploading a new license.
- Generating a Recovery Package.
- Performing StorageGRID software updates, including software upgrades, hotfixes, and updates to the SANtricity OS software on selected appliances.

## Related information

[Performing maintenance procedures](#)

[Downloading the Recovery Package](#)

[Expand your grid](#)

[Upgrade software](#)

[Maintain & recover](#)


[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

## Support menu

The Support menu provides options that help technical support analyze and troubleshoot your system. There are two parts to the Support menu: Tools and Alarms (legacy).



### Tools

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

- Enable AutoSupport.
- Perform a set of diagnostic checks on the current state of the grid.
- Access the Grid Topology tree to view detailed information about grid nodes, services, and attributes.
- Retrieve log files and system data.
- Review detailed metrics and charts.



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)

From the Alarms (legacy) section of the Support menu, you can review current, historical, and global alarms, and you can set up email notifications for legacy alarms and AutoSupport.

### Related information

[StorageGRID architecture and network topology](#)


[StorageGRID attributes](#)

[Using StorageGRID support options](#)

[Administer StorageGRID](#)

## Help menu

The Help option provides access to the StorageGRID Documentation Center for the current release and to the API documentation. You can also determine which version of StorageGRID is currently installed.



### Related information

[Administer StorageGRID](#)

## Exploring the Tenant Manager

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

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

### Related information

[Exploring the Grid Manager](#)

[Use a tenant account](#)

## 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

## Total objects

8,418,886  
objects

## Tenant details

Name Human Resources  
ID 4955 9096 9804 4285 4354

View the instructions for Tenant Manager.

[Go to documentation](#) [↗](#)

## Storage menu (S3 tenants only)

The Storage menu is provided for S3 tenant accounts only. This menu allows S3 users to manage access keys, create and delete buckets, and manage platform service endpoints.



## 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 related to buckets:

- Create buckets
- Enable S3 Object Lock for a new bucket (assumes that S3 Object Lock is enabled for the StorageGRID system)
- Update consistency level settings
- Configure cross-origin resource sharing (CORS)
- Enable and disable last access time update settings for the buckets belonging to the tenant
- Delete empty buckets

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 AWS Simple Notification Service™ (SNS).
- 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.

## 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.



## Using StorageGRID

After you install grid nodes and StorageGRID networks, you can begin to configure and use StorageGRID. Some of the tasks you will perform include controlling user access to system administration functions, setting up tenant accounts, managing client connections, setting configuration options, managing object locations with ILM, monitoring the health and day-to-day activities of your StorageGRID system, and performing routine and non-routine maintenance activities.

- [Controlling StorageGRID access](#)
- [Managing tenants and client connections](#)
- [Configuring network settings](#)
- [Configuring system settings](#)
- [Using information lifecycle management](#)
- [Monitoring StorageGRID operations](#)
- [Performing maintenance procedures](#)
- [Using StorageGRID support options](#)

### Controlling 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.

#### Controlling 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 cannot make any changes or perform any operations in the Grid Manager or Grid Management API.

#### Enabling single sign-on

The StorageGRID system supports single sign-on (SSO) using the Security Assertion Markup Language 2.0 (SAML 2.0) standard. When SSO is enabled, 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 cannot sign in to StorageGRID.

When SSO is enabled and users sign in to StorageGRID, they are redirected to your organization's SSO page to validate their credentials. When users sign out of one Admin Node, they are automatically signed out of all Admin Nodes.

### Using client certificates

You can use client certificates to allow authorized external clients to access the StorageGRID Prometheus database. Client certificates provide a secure way to use external tools to monitor StorageGRID. You can provide your own client certificate or generate one using the Grid Manager.

### Changing grid passwords

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 this passphrase as required.

### Related information

[Administer StorageGRID](#)

[Use a tenant account](#)

## Managing tenants and client connections

As a grid administrator, you create and manage the tenant accounts that S3 and Swift clients use to store and retrieve objects, and manage the configuration options that control how clients connect to your StorageGRID system.

### Tenant accounts

A tenant account allows you to specify who can use your StorageGRID system to store and retrieve objects, and which functionality is available to them. Tenant accounts allow client applications that support the S3 REST API or the Swift REST API to store and retrieve objects on StorageGRID. Each tenant account uses either the S3 client protocol or the Swift client protocol.

You must create at least one tenant account for each client protocol that will be used to store objects on your StorageGRID system. Optionally, you can create additional tenant accounts if you want to segregate the objects stored on your system by different entities. Each tenant account has its own federated or local groups and users, and its own buckets (containers for Swift) and objects.

You can use the Grid Manager or the Grid Management API to create tenant accounts. When creating a tenant account, you specify the following information:

- Display name for the tenant (the tenant's account ID is assigned automatically and cannot be changed).
- Whether the tenant account will use the S3 or Swift.
- For S3 tenant accounts: Whether the tenant account is allowed to use platform services. If the use of platform services is allowed, the grid must be configured to support their use.
- Optionally, a storage quota for the tenant account—the maximum number of gigabytes, terabytes, or petabytes available for the tenant's objects. A tenant's storage quota represents a logical amount (object size), not a physical amount (size on disk).
- If identity federation is enabled for the StorageGRID system, which federated group has Root Access permission to configure the tenant account.

- If single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid's identity source, and the initial password for the tenant's local root user.

If S3 tenant accounts need to comply with regulatory requirements, grid administrators can enable the global S3 Object Lock setting for the StorageGRID system. When S3 Object Lock is enabled for the system, all S3 tenant accounts can create buckets with S3 Object Lock enabled and then specify retention and legal hold settings for the object versions in that bucket.

After a tenant account is created, tenant users can sign in to the Tenant Manager.

#### Client connections to StorageGRID nodes

Before tenant users can use S3 or Swift clients to store and retrieve data in StorageGRID, you must decide how these clients will connect to StorageGRID nodes.

Client applications can store or retrieve objects by connecting to any of the following:

- The Load Balancer service on Admin Nodes or Gateway Nodes. This is the recommended connection.
- The CLB service on Gateway Nodes.



The CLB service is deprecated.

- Storage Nodes, with or without an external load balancer.

When configuring StorageGRID so that clients can use the Load Balancer service, you perform the following steps:

1. Configure endpoints for the Load Balancer service. The Load Balancer service on Admin Nodes or Gateway Nodes distributes incoming network connections from client applications to Storage Nodes. When creating a load balancer endpoint, you specify a port number, whether the endpoint accepts HTTP or HTTPS connections, the type of client (S3 or Swift) that will use the endpoint, and the certificate to be used for HTTPS connections (if applicable).
2. Optionally specify that a node's Client Network is untrusted to ensure that all connections to the node's Client Network occur on the load balancer endpoints.
3. Optionally configure high availability (HA) groups. If you create an HA group, the interfaces of multiple Admin Nodes and Gateway Nodes are placed into an active-backup configuration. Client connections are made using the virtual IP address of the HA group.

#### Related information

[Administer StorageGRID](#)

[Use a tenant account](#)

[Use S3](#)

[Use Swift](#)

[Exploring the Tenant Manager](#)

[Configuring network settings](#)

## Configuring network settings

You can configure various network settings from the Grid Manager to fine tune the operation of your StorageGRID system.

### Domain names

If you plan to support S3 virtual hosted-style requests, you must configure the list of endpoint domain names that S3 clients connect to. Examples include s3.example.com, s3.example.co.uk, and s3-east.example.com.



The configured server certificates must match the endpoint domain names.

### High availability groups

High availability groups use virtual IP addresses (VIPs) to provide active-backup access to Gateway Node or Admin Node services. An HA group consists of one or more network interfaces on Admin Nodes and Gateway Nodes. When creating an HA group, you select network interfaces belonging to the Grid Network (eth0) or the Client Network (eth2).



The Admin Network does not support HA VIPs.

An HA group maintains one or more virtual IP addresses that are added to the active interface in the group. If the active interface becomes unavailable, the virtual IP addresses are moved to another interface. This 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.

You might want to use high availability (HA) groups for several reasons.

- 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 and Swift clients.
- An HA group that contains only one interface allows you to provide many VIP addresses and to explicitly set IPv6 addresses.

### Link costs

You can adjust link costs to reflect the latency between sites. When two or more data center sites exist, link costs prioritize which data center site should provide a requested service.

### Load balancer endpoints


You can use a load balancer to handle ingest and retrieval workloads from S3 and Swift clients. Load balancing maximizes speed and connection capacity by distributing the workloads and connections across multiple Storage Nodes.

If you want to use the StorageGRID load balancer service, which is included on Admin Nodes and Gateway Nodes, you must configure one or more load balancer endpoints. Each endpoint defines a Gateway Node or Admin Node port for S3 and Swift requests to Storage Nodes.

### Proxy settings

If you are using S3 platform services or Cloud Storage Pools, you can configure a non-transparent proxy server between Storage Nodes and the external S3 endpoints. If you send AutoSupport messages using

HTTPS or HTTP, you can configure a non-transparent proxy server between Admin Nodes and technical support.



### Server certificates

You can upload two types of server certificates:

- Management Interface Server Certificate, which is the certificate used for accessing the management interface.
- Object Storage API Service Endpoints Server Certificate, which secures the S3 and Swift endpoints for connections directly to Storage Nodes or when using the CLB service on a Gateway Node.



The CLB service is deprecated.

Load balancer certificates are configured on the Load Balancer Endpoints page. Key management server (KMS) certificates are configured on the Key Management Server page.

### Traffic classification policies

Traffic classification policies allow you to create rules for identifying and handling 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.

### Untrusted Client Networks

If you are using a Client Network, you can help secure StorageGRID from hostile attacks 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.

For example, you might want a Gateway Node to refuse all inbound traffic on the Client Network except for HTTPS S3 requests. Or, you might want to enable outbound S3 platform service traffic from a Storage Node, while preventing any inbound connections to that Storage Node on the Client Network.

### Related information

[Administer StorageGRID](#)

[Managing tenants and client connections](#)

### Configuring system settings

You can configure various system settings from the Grid Manager to fine tune the operation of your StorageGRID system.

### Display options

Display options allow you to specify the timeout period for user sessions and to suppress email notifications for legacy alarms and event-triggered AutoSupport messages.

## Grid options

You can use Grid Options to configure the settings for all of the objects stored in your StorageGRID system, including stored object compression, stored object encryption, and stored object hashing.

You can also use these options to specify global settings for S3 and Swift client operations.

## Key management servers


You can configure one or more external key management servers (KMS) to provide encryption keys to StorageGRID services and storage appliances. Each KMS or KMS cluster uses the Key Management Interoperability Protocol (KMIP) to provide an encryption key to the appliance nodes at the associated StorageGRID site. Using key management servers lets you protect StorageGRID data even if an appliance is removed from the data center. After the appliance volumes are encrypted, you cannot 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.

## 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). You can enable the global S3 Object Lock setting for a StorageGRID system to allow S3 tenant accounts to create buckets with S3 Object Lock enabled. The tenant can then use an S3 client application to optionally specify retention settings (retain until date, legal hold, or both) for the objects in those buckets.



## Storage options

Storage options allow you to control object segmentation and to define storage watermarks to manage a Storage Node's usable storage space.

## Using information lifecycle management

You use information lifecycle management (ILM) to control the placement, duration, and data protection 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.

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, in Cloud Storage Pools, or on Archive Nodes.
- 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. The copies are evenly distributed across all Storage Nodes.

### Example ILM rule

This example ILM rule applies to the objects belonging to Tenant A. It makes two replicated copies of those objects and stores each copy at a different site. The two copies are retained “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.

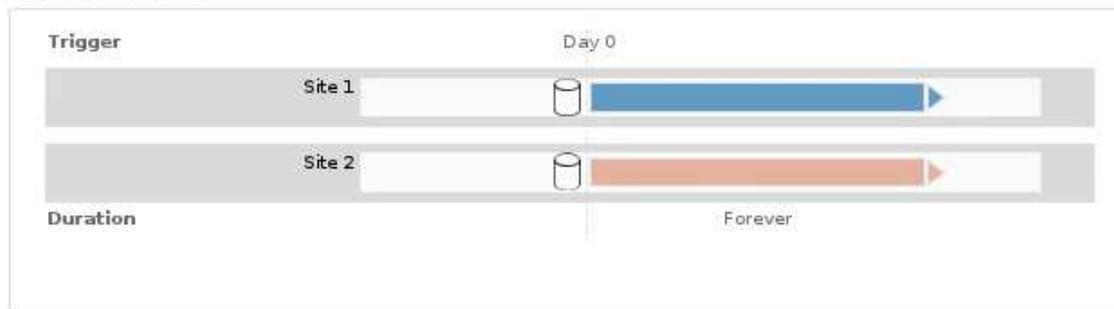
This rule uses the Balanced option for ingest behavior: the two-site placement instruction is applied as soon as Tenant A saves an object to StorageGRID, unless it is not possible to immediately make both required copies. For example, if Site 2 is unreachable when Tenant A saves an object, StorageGRID 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.

## Two copies at two sites for Tenant A

Description: Applies only to Tenant A  
Ingest Behavior: Balanced  
Tenant Accounts: Tenant A (34176783492629515782)  
Reference Time: Ingest Time  
Filtering Criteria:

Matches all objects.

### Retention Diagram:



### How an ILM policy evaluates objects

The active ILM policy for your StorageGRID system controls the placement, duration, and data protection 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 do not 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 cannot use any filters.

### Example ILM policy

This example ILM policy uses three ILM rules.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### 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.

	Default	Rule Name	Tenant Account	Actions
		Rule 1: 3 replicated copies for Tenant A	Tenant A (58889986524346589742)	
		Rule 2: Erasure coding for objects greater than 1 MB	—	
	<input checked="" type="checkbox"/>	Rule 3: 2 copies 2 data centers (default)	—	

In this example, Rule 1 matches all objects belonging to Tenant A. These objects are stored as three replicated copies at three sites. Objects belonging to other tenants are not matched by Rule 1, so they are evaluated against Rule 2.

Rule 2 matches all objects from other tenants but only if they are larger than 1 MB. These larger objects are stored using 6+3 erasure coding at three sites. Rule 2 does not match objects 1 MB or smaller, so these objects are evaluated against Rule 3.

Rule 3 is the last and default rule in the policy, and it does not use filters. Rule 3 makes 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](#)

#### Monitoring StorageGRID operations

The Grid Manager provides information for monitoring the daily activities of your StorageGRID system, including its health.

- [Viewing the Nodes page](#)
- [Monitoring and managing alerts](#)
- [Using SNMP monitoring](#)
- [Reviewing audit messages](#)

#### Viewing 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.

Dashboard

Alerts

Nodes

Tenants

ILM

Configuration

Maintenance

Support

## StorageGRID Deployment

## Data Center 1

- ✓ DC1-ADM1
- ✓ DC1-ARC1
- ✓ DC1-G1
- ✓ DC1-S1
- ✓ DC1-S2
- ✓ DC1-S3

## Data Center 2

- ✓ DC2-ADM1
- ✓ DC2-S1
- ✓ DC2-S2
- ✓ DC2-S3

## Data Center 3

- ✓ DC3-S1
- ✓ DC3-S2
- ✓ DC3-S3

## StorageGRID Deployment

Network

Storage

Objects

ILM

Load Balancer


1 hour

1 day

1 week

1 month

Custom



From the tree view on the left, you can see all the sites and all the nodes in your StorageGRID system. The icon for each node indicates if the node is connected or if there are any active alerts.

### Connection state icons

If a node is disconnected from the grid, the tree view shows a blue or gray connection state icon, not the icon for any underlying alerts.

- **Not connected - Unknown** : The node is not connected to the grid for an unknown reason. For example, the network connection between nodes has been lost or the power 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.

- **Not connected - 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.

### Alert icons

If a node is connected to the grid, the tree view shows one of the following icons, depending on if there are any current alerts for the node.

- **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 do not clear on their own to ensure they do not result in a more serious problem.
- **Normal** : No alerts are active, and the node is connected to the grid.

## Viewing details for a system, site, or node

To view the available information, click the appropriate links on the left, as follows:

- Select the grid name to see an aggregate summary of the statistics for your entire StorageGRID system. (The screenshot shows a system named StorageGRID Deployment.)
- 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.

## Related information

[Monitor & troubleshoot](#)

## Tabs for the Nodes page

The tabs at the top of the Nodes page are based on what you select from the tree at the left.

Tab name	Description	Included for
Overview	<ul style="list-style-type: none"><li>• Provides basic information about each node.</li><li>• Shows any current, unacknowledged alarms affecting the node.</li></ul>	All nodes
Hardware	<ul style="list-style-type: none"><li>• Displays CPU utilization and memory usage for each node</li><li>• For appliance nodes, provides additional hardware information.</li></ul>	All nodes
Network	Displays a graph showing the network traffic received and sent across the network interfaces.	All nodes, each site, and the entire grid
Storage	<ul style="list-style-type: none"><li>• Provides details for the disk devices and volumes on each node.</li><li>• For Storage Nodes, each site, and the entire grid, includes graphs showing object data storage and metadata storage used over time.</li></ul>	All nodes, each site, and the entire grid
Events	Displays a count of any system error or fault event, including errors such as network errors.	All nodes


Tab name	Description	Included for
Objects	<ul style="list-style-type: none"> <li>Provides information about S3 and Swift ingest and retrieve rates.</li> <li>For Storage Nodes, provides object counts and information about metadata store queries and background verification.</li> </ul>	Storage Nodes, each site, and the entire grid
ILM	<p>Provides information about Information Lifecycle Management (ILM) operations.</p> <ul style="list-style-type: none"> <li>For Storage Nodes, provides details about ILM evaluation and background verification for erasure coded objects.</li> <li>For each site and the entire grid, shows a graph of the ILM queue over time.</li> <li>For the entire grid, provides the estimated time to complete a full ILM scan of all objects.</li> </ul>	Storage Nodes, each site, and the entire grid
Load Balancer	<p>Includes performance and diagnostic graphs related to the Load Balancer service.</p> <ul style="list-style-type: none"> <li>For each site, provides an aggregate summary of the statistics for all nodes at that site.</li> <li>For the entire grid, provides an aggregate summary of the statistics for all sites.</li> </ul>	Admin Nodes and Gateway Nodes, each site, and the entire grid
Platform Services	Provides information about any S3 platform service operations at a site.	Each site
SANtricity System Manager	Provides access to SANtricity System Manager. From SANtricity System Manager, you can review hardware diagnostic and environmental information for the storage controller, as well as issues related to the drives.	Storage appliance nodes  <b>Note:</b> The SANtricity System Manager tab will not appear if the controller firmware on the storage appliance is less than 8.70.

### Prometheus metrics

The Prometheus service on Admin Nodes collects time series metrics from the services on all nodes.

The metrics collected by Prometheus are used in a number of 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 (**Help > 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 a number of 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.

### Related information

[Monitoring and managing alerts](#)

[Using StorageGRID support options](#)

[Monitor & troubleshoot](#)

### 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 a number of 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.)

- **Alarms:** When attributes reach defined threshold values, StorageGRID alarms (legacy system) are triggered at specific severity levels.
- **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.

## Related information

[Monitor & troubleshoot](#)

### Monitoring and managing alerts

The alert system provides an easy-to-use interface for detecting, evaluating, and resolving the issues that can occur during StorageGRID operation.

The alert system is designed to be your primary tool for monitoring any issues that might occur in your StorageGRID system.

- The alert system focuses on actionable problems in the system. Alerts are triggered for events that require your immediate attention, not for events that can safely be ignored.
- The Current Alerts and Resolved Alerts pages provide a user friendly interface for viewing current and historical problems. You can sort the listing by individual alerts and alert groups. For example, you might want to sort all alerts by node/site to see which alerts are affecting a specific node. Or, you might want to sort the alerts in a group by time triggered to find the most recent instance of a specific alert.
- Multiple alerts of the same type are grouped into one email to reduce the number of notifications. In addition, multiple alerts of the same type are shown as a group on the Current Alerts and Resolved Alerts pages. You can expand and collapse alert groups to show or hide the individual alerts. For example, if several nodes are reporting the **Unable to communicate with node** alert, only one email is sent and the alert is shown as a group on the Current Alerts page.

Current Alerts [Learn more](#)

View the current alerts affecting your StorageGRID system.

Name	Severity	Time triggered	Site / Node	Status	Current values
<b>Unable to communicate with node</b> One or more services are unresponsive or cannot be reached by the metrics collection job.	2 Major	9 minutes ago <i>(newest)</i> 19 minutes ago <i>(oldest)</i>		2 Active	
<b>Low root disk capacity</b> The space available on the root disk is low.	Minor	25 minutes ago	Data Center 1 / DC1-S1-99-51	Active	Disk space available: 2.00 GB Total disk space: 21.00 GB
<b>Expiration of server certificate for Storage API Endpoints</b> The server certificate used for the storage API endpoints is about to expire.	Major	31 minutes ago	Data Center 1 / DC1-ADM1-99-49	Active	Days remaining: 14
<b>Expiration of server certificate for Management Interface</b> The server certificate used for the management interface is about to expire.	Minor	31 minutes ago	Data Center 1 / DC1-ADM1-99-49	Active	Days remaining: 30
<b>Low installed node memory</b> The amount of installed memory on a node is low.	8 Critical	a day ago <i>(newest)</i> a day ago <i>(oldest)</i>		8 Active	

- Alerts use intuitive names and descriptions to help you understand more quickly what the problem is. Alert notifications include details about the node and site affected, the alert severity, the time when the alert rule was triggered, and the current value of metrics related to the alert.
- Alert email notifications and the alert listings on the Current Alerts and Resolved Alerts pages provide recommended actions for resolving an alert. These recommended actions often include direct links to StorageGRID documentation to make it easier to find and access more detailed troubleshooting procedures.

### Low installed node memory

The amount of installed memory on a node is low.

Status

Active ([silence this alert](#))

#### Recommended actions

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.

Site / Node

Data Center 2 / DC2-S1-99-56

Severity

✘ Critical

See the instructions for your platform:

- [VMware installation](#)
- [Red Hat Enterprise Linux or CentOS installation](#)
- [Ubuntu or Debian installation](#)

Total RAM size

8.38 GB

Condition

[View conditions](#) | [Edit rule](#)

#### Time triggered

2019-07-15 17:07:41 MDT (2019-07-15 23:07:41 UTC)

Close



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

## Managing alerts

All StorageGRID users can view alerts. If you have the Root Access or Manage Alerts permission, you can also manage alerts, as follows:

- If you need to temporarily suppress the notifications for an alert at one or more severity levels, you can easily silence a specific alert rule for a specified duration. You can silence an alert rule for the entire grid, a single site, or a single node.
- You can edit the default alert rules as required. You can disable an alert rule completely, or change its trigger conditions and duration.
- You can create custom alert rules to target the specific conditions that are relevant to your situation and to provide your own recommended actions. To define the conditions for a custom alert, you create expressions using the Prometheus metrics available from the Metrics section of the Grid Management API.

For example, 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 < 24000000000
```

## Related information

### [Monitor & troubleshoot](#)

#### Using SNMP monitoring

If you want to monitor StorageGRID using the Simple Network Management Protocol (SNMP), you can use the Grid Manager to configure the SNMP agent.

Each StorageGRID node runs an SNMP agent, or daemon, that provides a management information base (MIB). The StorageGRID MIB contains table and notification definitions for alerts and alarms. Each StorageGRID node also supports a subset of MIB-II objects.

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. The agent provides read-only MIB access for queries, and it can send two types of event-driven notifications to a management system:

- **Traps** are notifications sent by the SNMP agent that do not 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 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 whichever Admin Node is configured to be the preferred sender.
- Certain alarms (legacy system) are triggered at specified severity levels or higher.



SNMP notifications are not sent for every alarm or every alarm severity.

## Related information

### [Monitor & troubleshoot](#)

#### Reviewing 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 or Swift 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.

For easy access to audit logs, you can configure client access to the audit share for both NFS and CIFS (deprecated). You can also access audit log files directly from the command line of the Admin Node.

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 the instructions for audit messages. To learn how to configure audit client access, see the instructions for administering StorageGRID.

### **Related information**

[Review audit logs](#)

[Administer StorageGRID](#)

### **Performing maintenance procedures**

You perform various maintenance procedures to keep your StorageGRID system up-to-date and to ensure it is performing efficiently. The Grid Manager provides tools and options to facilitate the process of performing maintenance tasks.

#### **Software updates**

You can perform three types of software updates from the Software Update page in the Grid Manager:

- StorageGRID software upgrade
- StorageGRID hotfix
- SANtricity OS upgrade

#### **StorageGRID software upgrades**

When a new StorageGRID feature release is available, the Software Upgrade page guides you through the process of uploading the required file and upgrading your StorageGRID system. You must upgrade all grid nodes for all data center sites from the primary Admin Node.

During a StorageGRID software upgrade, client applications can continue to ingest and retrieve object data.

#### **Hotfixes**

If issues with the software are detected and resolved between feature releases, you might need to apply a hotfix to your StorageGRID system.

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.

The StorageGRID Hotfix page, shown below, allows you to upload a hotfix file.

## 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 

Browse

### Passphrase

Provisioning Passphrase 

Start

The hotfix is applied first to the primary Admin Node. Then, you must approve the application of the hotfix to other grid nodes until all nodes in your StorageGRID system 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.



While all 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.

### SANtricity OS upgrades

You might need to upgrade the SANtricity OS Software on the storage controllers of your storage appliances, if the controllers are not functioning optimally. You can upload the SANtricity OS file to the primary Admin Node in your StorageGRID system and apply the upgrade from the Grid Manager.

The SANtricity page, shown below, allows you to upload the SANtricity OS upgrade file.

## SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

### SANtricity OS Upgrade File

---

SANtricity OS Upgrade File



Browse

### Passphrase

---

Provisioning Passphrase



Start

After you upload the file, you can approve the upgrade on individual Storage Nodes or all nodes. The ability to selectively approve nodes makes it easier for you to schedule the upgrade. After you approve a node for upgrade, the system performs a health check and installs the upgrade if it is applicable to the node.

### Expansion procedures

You can expand a StorageGRID system by adding storage volumes to Storage Nodes, adding new grid nodes to an existing site, or adding a new data center site. If you have Storage Nodes that use the SG6060 storage appliance, you can add one or two expansion shelves to double or triple the storage capacity of the node.

You can perform expansions without interrupting the operation of your current system. When you add nodes or a site, you first deploy the new nodes and then perform the expansion procedure from the Grid Expansion page.

**i** A new Recovery Package has been generated as a result of the configuration change. Go to the [Recovery Package page](#) to download it.

**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
<b>Grid Node Status</b>						
Lists the installation and configuration status of each grid node included in the expansion.						
						Search <input type="text"/>
Name	Site	Grid Network IPv4 Address	Progress	Stage		
DC2-ADM1-184	Site A	172.17.3.184/21	<div style="width: 100%; height: 10px; background-color: #0070C0;"></div>	Waiting for NTP to synchronize		
DC2-S1-185	Site A	172.17.3.185/21	<div style="width: 100%; height: 10px; background-color: #0070C0;"></div>	Waiting for Dynamic IP Service peers		
DC2-S2-186	Site A	172.17.3.186/21	<div style="width: 100%; height: 10px; background-color: #0070C0;"></div>	Waiting for NTP to synchronize		
DC2-S3-187	Site A	172.17.3.187/21	<div style="width: 100%; height: 10px; background-color: #0070C0;"></div>	Waiting for NTP to synchronize		
DC2-S4-188	Site A	172.17.3.188/21	<div style="width: 100%; height: 10px; background-color: #0070C0;"></div>	Waiting for Dynamic IP Service peers		
DC2-ARC1-189	Site A	172.17.3.189/21	<div style="width: 100%; height: 10px; background-color: #0070C0;"></div>	Waiting for NTP to synchronize		
2. Initial Configuration						Pending
3. Distributing the new grid node's certificates to the StorageGRID system.						Pending
4. Starting services on the new grid nodes						Pending
5. Cleaning up unused Cassandra keys						Pending

**Node recovery procedures**

Grid nodes can fail if a hardware, virtualization, operating system, or software fault renders the node inoperable or unreliable.

The steps to recover a grid node depend on the platform where the grid node is hosted and on the type of grid node. Each type of grid node has a specific recovery procedure, which you must follow exactly. Generally, you try to preserve data from the failed grid node where possible, repair or replace the failed node, use the Recovery page to configure the replacement node, and restore the node's data.

For example, this flowchart shows the recovery procedure if an Admin Node has failed.



**Decommission procedures**

You might want to permanently remove grid nodes or an entire data center site from your StorageGRID system.

For example, you might want to decommission one or more grid nodes in these cases:

- You have added a larger Storage Node to the system and you want to remove one or more smaller Storage Nodes, while at the same time preserving objects.
- You require less total storage.
- You no longer require a Gateway Node or a non-primary Admin Node.
- Your grid includes a disconnected node that you cannot recover or bring back online.

You can use the Decommission Nodes page in the Grid Manager to remove the following types of grid nodes:

- Storage Nodes, unless not enough nodes would remain at the site to support certain requirements
- Gateway Nodes
- Non-primary Admin Nodes

### Decommission Nodes

Before decommissioning a grid node, review the health of all nodes. If possible, resolve any issues or alarms before proceeding.

Select the checkbox for each grid node you want to decommission. If decommission is not possible for a node, see the Recovery and Maintenance Guide to learn how to proceed.

#### Grid Nodes

	Name	Site	Type	Has ADC	Health	Decommission Possible
	DC1-ADM1	Data Center 1	Admin Node	-		No, primary Admin Node decommissioning is not supported.
<input type="checkbox"/>	DC1-ADM2	Data Center 1	Admin Node	-		
<input type="checkbox"/>	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.
<input type="checkbox"/>	DC1-S4	Data Center 1	Storage Node	No		
<input type="checkbox"/>	DC1-S5	Data Center 1	Storage Node	No		

#### Passphrase

Provisioning  
Passphrase

Start Decommission

You can use the Decommission Site page in the Grid Manager to remove a site. A connected site decommission removes an operational site and preserves data. A disconnected site decommission removes a failed site but does not preserve data. The Decommission Site wizard guides you through the process of selecting the site, viewing site details, revising the ILM policy, removing site references from ILM rules, and resolving any node conflicts.

## Decommission Site



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
<input checked="" type="radio"/>	Raleigh	3.93 MB	
<input checked="" type="radio"/>	Sunnyvale	3.97 MB	
<input type="radio"/>	Vancouver	3.90 MB	No. This site contains the primary Admin Node.

Next

### Network maintenance procedures

Some of the network maintenance procedures you might need to perform include the following:

- Updating the subnets on the Grid Network
- Using the Change IP tool to change the networking configuration that was initially set during grid deployment
- Adding, removing, or updating domain name system (DNS) servers
- Adding, removing, or updating network time protocol (NTP) servers to ensure that data is synchronized accurately between grid nodes
- Restoring network connectivity to nodes that might have become isolated from the rest of the grid

### Host-level and middleware procedures

Some maintenance procedures are specific to StorageGRID nodes that are deployed on Linux or VMware, or are specific to other components of the StorageGRID solution. For example, you might want to migrate a grid node to a different Linux host or perform maintenance on an Archive Node that is connected to Tivoli Storage Manager (TSM).

### Appliance node cloning

Appliance node cloning lets you easily replace an existing appliance node (source) in your grid with a compatible appliance (target) that is part of the same logical StorageGRID site. The process transfers all data to the new appliance, placing it in service to replace the old appliance node and leaving the old appliance in a pre-install state. Cloning provides a hardware-upgrade process that is easy to perform, and provides an alternate method for replacing appliances.

## Grid node procedures

You might need to perform certain procedures on a specific grid node. For example, you might need to reboot a grid node or manually stop and restart a specific grid node service. Some grid node procedures can be performed from the Grid Manager; others require you to log in to the grid node and use the node's command line.

### Related information

[Administer StorageGRID](#)

[Upgrade software](#)

[Expand your grid](#)

[Maintain & recover](#)

### Downloading the Recovery Package

The Recovery Package is a downloadable .zip file that contains deployment-specific files and software needed to install, expand, upgrade, and maintain a StorageGRID system.

The Recovery Package file also contains system-specific configuration and integration information, including server hostnames and IP addresses, and highly confidential passwords needed during system maintenance, upgrade, and expansion. The Recovery Package is required to recover from the failure of the primary Admin Node.

When installing a StorageGRID system, you are required to download the Recovery Package file and to confirm that you can successfully access the contents of this file. You should also download the file each time the grid topology of the StorageGRID system changes because of maintenance or upgrade procedures.

#### Recovery Package

Enter your provisioning passphrase and click Start Download to save a copy of the Recovery Package file. Download the file each time the grid topology of the StorageGRID system changes because of maintenance or upgrade procedures, so that you can restore the grid if a failure occurs.

When the download completes, copy the Recovery Package file to two safe, secure, and separate locations.

**Important:** 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.

Provisioning Passphrase

[Start Download](#)

After downloading the Recovery Package file and confirming you can extract the contents, copy the Recovery Package 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.

### Related information

[Upgrade software](#)

[Expand your grid](#)

## Using StorageGRID support options

The Grid Manager provides options to help you work with technical support if an issue arises with your StorageGRID system.

### Configuring AutoSupport

The AutoSupport feature enables your StorageGRID system to send health and status messages to technical support. Using AutoSupport can significantly speed 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 messages to be sent to one additional destination.

### Information included in AutoSupport messages

AutoSupport messages include information such as the following:

- StorageGRID software version
- Operating system version
- System-level and location-level attribute information
- Recent alerts and alarms (legacy system)
- Current status of all grid tasks, including historical data
- Events information as listed on the **Nodes** > *node* > **Events** page
- Admin Node database usage
- Number of lost or missing objects
- Grid configuration settings
- NMS entities
- Active ILM policy
- Provisioned grid specification file
- Diagnostic metrics

You can enable the AutoSupport feature and the individual AutoSupport options when you first install StorageGRID, or you can enable them later. If AutoSupport is not enabled, a message appears on the Grid ManagerDashboard. The message includes a link to the AutoSupport configuration page.

The AutoSupport feature is disabled. You should enable AutoSupport to allow StorageGRID to send health and status messages to technical support for proactive monitoring and troubleshooting.



You can select the “x” symbol  to close the message. The message will not appear again until your browser cache is cleared, even if AutoSupport remains disabled.

### Using 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.

You must enable AutoSupport if you want to use the Active IQ dashboards and functionality on the NetApp Support site.

[Active IQ Digital Advisor Documentation](#)

## Accessing AutoSupport settings

You configure AutoSupport using the Grid Manager (**Support > Tools > AutoSupport**). The **AutoSupport** page has two tabs: **Settings** and **Results**.

### AutoSupport

The AutoSupport feature enables your StorageGRID system to send periodic and event-driven health and status messages to technical support to allow proactive monitoring and troubleshooting. StorageGRID AutoSupport also enables the use of Active IQ for predictive recommendations.

Settings Results

**Protocol Details**

Protocol ?  HTTPS  HTTP  SMTP

NetApp Support Certificate Validation ? Use NetApp support certificate

**AutoSupport Details**

Enable Weekly AutoSupport ?

Enable Event-Triggered AutoSupport ?

Enable AutoSupport on Demand ?

**Additional AutoSupport Destination**

Enable Additional AutoSupport Destination ?

Save Send User-Triggered AutoSupport

## Protocols for sending AutoSupport messages

You can choose one of three protocols for sending AutoSupport messages:

- HTTPS
- HTTP
- SMTP

If you send AutoSupport messages using HTTPS or HTTP, you can configure a non-transparent proxy server between Admin Nodes and technical support.

If you use SMTP as the protocol for AutoSupport messages, you must configure an SMTP mail server.

## AutoSupport options

You can use any combination of the following options to send AutoSupport messages to technical support:

- **Weekly:** Automatically send AutoSupport messages once per week. Default setting: Enabled.
- **Event-triggered:** Automatically send AutoSupport messages every hour or when significant system events occur. Default setting: Enabled.
- **On Demand:** Allow technical support to request that your StorageGRID system send AutoSupport messages automatically, which is useful when they are actively working an issue (requires HTTPS AutoSupport transmission protocol). Default setting: Disabled.
- **User-triggered:** Manually send AutoSupport messages at any time.

## Related information

[Administer StorageGRID](#)

[Configuring network settings](#)

## Collecting StorageGRID logs

To help troubleshoot a problem, you might need to collect log files and forward them to technical support.

StorageGRID uses log files to capture events, diagnostic messages, and error conditions. The bycast.log file is maintained for every grid node and is the primary troubleshooting file. StorageGRID also creates log files for individual StorageGRID services, log files related to deployment and maintenance activities, and log files related to third-party applications.

Users who have the appropriate permissions and who know the provisioning passphrase for your StorageGRID system can use the Logs page in the Grid Manager to gather log files, system data, and configuration data. When you collect logs, you select a node or nodes and specify a time period. Data is collected and archived in a `.tar.gz` file, which you can download to a local computer. Inside this file, there is one log file archive for each grid node.

### Logs

Collect log files from selected grid nodes for the given time range. Download the archive package after all logs are ready.

▲ ▲  StorageGRID Webscale Deployment

- ▲ ▲  Data Center 1
  - DC1-ADM1
  - ▲  DC1-ARC1
  - DC1-G1
  - DC1-S1
  - DC1-S2
  - DC1-S3
- ▲  Data Center 2
  - DC2-ADM1
  - DC2-S1
  - DC2-S2
  - DC2-S3
- ▲  Data Center 3
  - DC3-S1
  - DC3-S2
  - DC3-S3

Log Start Time: 2018-04-18 [calendar icon] 01 : 38 PM MDT

Log End Time: 2018-04-18 [calendar icon] 05 : 38 PM MDT

Notes: [text area]

Provisioning Passphrase: [text field]

## Related information

[Monitor & troubleshoot](#)

[Administer StorageGRID](#)

### Using metrics and running diagnostics

When troubleshooting an issue, you can work with technical support to review detailed metrics and charts for your StorageGRID system. You can also run pre-constructed diagnostic queries to proactively assess key values for your StorageGRID system.

### Metrics page

The Metrics page provides access to 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.

## Metrics

Access charts and metrics to help troubleshoot issues.

**i** 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://\[redacted\] /metrics/graph](https://[redacted] /metrics/graph)

### 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.

<a href="#">ADE</a>	<a href="#">Node</a>
<a href="#">Account Service Overview</a>	<a href="#">Node (Internal Use)</a>
<a href="#">Alertmanager</a>	<a href="#">Platform Services Commits</a>
<a href="#">Audit Overview</a>	<a href="#">Platform Services Overview</a>
<a href="#">Cassandra Cluster Overview</a>	<a href="#">Platform Services Processing</a>
<a href="#">Cassandra Network Overview</a>	<a href="#">Replicated Read Path Overview</a>
<a href="#">Cassandra Node Overview</a>	<a href="#">S3 - Node</a>
<a href="#">Cloud Storage Pool Overview</a>	<a href="#">S3 Overview</a>
<a href="#">EC - ADE</a>	<a href="#">Site</a>
<a href="#">EC - Chunk Service</a>	<a href="#">Support</a>
<a href="#">Grid</a>	<a href="#">Traces</a>
<a href="#">ILM</a>	<a href="#">Traffic Classification Policy</a>
<a href="#">Identity Service Overview</a>	<a href="#">Usage Processing</a>
<a href="#">Ingests</a>	<a href="#">Virtual Memory (vmstat)</a>

The link in the Prometheus section of the Metrics page allows you to query the current values of StorageGRID metrics and to view graphs of the values over time.

Enable query history

Expression (press Shift+Enter for newlines)

Execute - insert metric at cursor - ▾

Graph Console

Element	Value
no data	

[Remove Graph](#)

Add Graph



Metrics that include *private* in their names are intended for internal use only and are subject to change between StorageGRID releases without notice.

The links in the Grafana section of the Metrics page allow you to access pre-constructed dashboards containing graphs of StorageGRID metrics over time.



## Diagnostics page

The Diagnostics page performs a set of pre-constructed diagnostic checks on the current state of the grid. In the 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 blocked task queue too large**

✓ **Cassandra commit log latency**

✓ **Cassandra commit log queue depth**

✓ **Cassandra compaction queue too large**

Clicking a specific diagnostic lets you see details about the diagnostic and its current results.

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 >= 75%  
✖ Caution >= 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%

## Related information

[Monitor & troubleshoot](#)

# Networking guidelines

Learn about StorageGRID architecture and networking topologies. Become familiar with the requirements for network configuration and provisioning.

- [StorageGRID networking overview](#)
- [Networking requirements and guidelines](#)
- [Deployment-specific networking considerations](#)
- [Network installation and provisioning](#)
- [Post-installation guidelines](#)
- [Network port reference](#)

## StorageGRID networking overview

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 the *Grid primer*.

Before you deploy and configure StorageGRID, you must configure the networking infrastructure. Communication needs to 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.

After you determine which of the three 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.

## Related information

[Grid primer](#)

[Administer StorageGRID](#)

[Release notes](#)

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

[Install Ubuntu or Debian](#)

[Install VMware](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

### 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)
Client traffic	The traffic that travels between external client applications and the grid, including all object storage requests from S3 and Swift clients.	Client 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.

## 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.

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 do not want to configure a gateway for a network, you should use the IP address as the gateway address.

High availability (HA) groups provide the ability to add virtual IP addresses to the Grid or Client Network interface. For more information, see the instructions for administering StorageGRID.

### 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:

- 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.

By using an Admin Network, administrative and maintenance traffic does not need to travel across the Grid Network. Typical uses of the Admin Network include access to the Grid Manager user interface; access to critical services such as NTP, DNS, external key management (KMS), and Lightweight Directory Access Protocol (LDAP); access to audit logs on Admin Nodes; and 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:

- 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 and Swift. 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 do not 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. For more information about the Untrusted Client Network feature and the Load Balancer service, see the instructions for administering StorageGRID.

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 and Swift client access to the grid.
- Storage Nodes, because these nodes provide access to the S3 and Swift protocols 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:

- 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.

### **Related information**

[Networking requirements and guidelines](#)

[Administer StorageGRID](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

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

[Install Ubuntu or Debian](#)

[Install VMware](#)

### **Network topology examples**

In addition to the required Grid Network, you can choose whether to configure Admin Network and Client Network interfaces when designing the network topology for a single or multi-site deployment.

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. For more information about internal and external ports, see the network port reference.

If you specify that a node's Client Network interface is untrusted, configure a load balancer endpoint to accept the inbound traffic. For information about configuring untrusted Client Networks and load balancer endpoints, see the instructions for administering StorageGRID.

### **Related information**

[Administer StorageGRID](#)

[Network port reference](#)

### **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 and Swift client requests and administrative and maintenance functions.



This topology is appropriate for single-site deployments that are not 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 and Swift 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,	172.16.0.0/16 → eth0	Static	GNSL
Storage 1,	10.10.1.0/24 → eth1	Link	Interface IP/mask
3, and 4	10.10.2.0/24 → 10.10.1.1	Static	AESL
	10.10.3.0/24 → 10.10.1.1	Static	AESL
Storage 2,	172.16.0.0/16 → eth0	Static	GNSL
Gateway	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 and Swift) 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 and Swift client requests and for administrative functions, while the Grid Network is dedicated to internal object management operations.

Topology example: Grid and Client Networks



*Provisioned*

**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

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 and Swift 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 `eithersupport.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 and Swift clients.
- Outbound requests from StorageGRID platform services such as Cloud Mirror replication or from Cloud Storage Pools.

If StorageGRID is unable to make contact with 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 are not 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 has been completed by using the Change IP tool as described in the recovery and maintenance instructions.

## 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. See the information about high availability groups in the instructions for administering StorageGRID.

## Using network address translation (NAT)

Do not 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.

### Related information

[Grid primer](#)

[Administer StorageGRID](#)

[Maintain & recover](#)

## 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 do not 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 cannot 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) cannot overlap the subnet of any other interface configured on the same node.
- The Grid Network subnet for each node must be included in the GNSL.
- The Admin Network subnet cannot overlap the Grid Network subnet, the Client Network subnet, or any subnet in the GNSL.
- The subnets in the AESL cannot overlap with any subnets in the GNSL.
- The Client Network subnet cannot 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.

## 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 are not 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 and Swift 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 the information about managing load balancing and managing untrusted Client Networks in the instructions for administering StorageGRID.

## Related information

[Administer StorageGRID](#)

## Deployment-specific networking considerations

Depending on the deployment platforms you use, you might have additional considerations for your StorageGRID network design.

Grid nodes can be deployed as:

- Software-based grid nodes deployed as virtual machines in VMware vSphere Web Client
- Software-based grid nodes deployed within Docker containers on Linux hosts
- Appliance-based nodes

For additional information about grid nodes, see the *Grid primer*.

## Related information

[Grid primer](#)

## Linux deployments

For efficiency, reliability, and security, the StorageGRID system runs on Linux as a collection of Docker containers. Docker-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.



Do not 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/CentOS or Ubuntu/Debian deployments.

## Related information

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

[Install Ubuntu or Debian](#)

## Host network configuration for Docker deployments

Before starting your StorageGRID deployment on a Docker container 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 cannot share IP addresses.

- Open the required ports to the host.

## Minimum bandwidth recommendations

The following table provides the minimum 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
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 information about configuring the host network in the installation instructions for your Linux platform.

### Related information

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

[Install 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. 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. In this example, the Untrusted Client Network feature can be used to restrict inbound connections.

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.

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.

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

For more information about untrusted Client Networks, see the instructions for [administering StorageGRID](#). For more information about platform services, see the instructions for [using tenant accounts](#). For more information about Cloud Storage Pools, see the instructions for [managing objects with information lifecycle management](#).

#### **Related information**

[Network port reference](#)

[Grid primer](#)

[Administer StorageGRID](#)

[Use a tenant account](#)

[Manage objects with ILM](#)

#### **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 in the installation and maintenance instructions for your appliance.

#### **Related information**

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

## 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 the information about network maintenance procedures in the recovery and maintenance instructions.

### Related information

[Maintain & recover](#)

## 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 cannot set up DHCP during configuration.



Nodes reboot when their IP addresses change, which can cause outages if a DHCP address 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 the information about configuring IP addresses in the recovery and maintenance instructions.
- 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 re-establish these connections.

## Related information

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

[Install Ubuntu or Debian](#)

[Install VMware](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[Maintain & recover](#)

## Network port reference

You must ensure the network infrastructure can provide internal and external communication between nodes within the grid and to external clients and services. You might need access across internal and external firewalls, switching systems, and routing systems.

Use the details provided for internal grid node communications and external communications to determine how to configure each required port.

- [Internal grid node communications](#)
- [External communications](#)

### Internal grid node communications

The StorageGRID internal firewall only allows incoming connections to specific ports on the Grid Network, with the exception of ports 22, 80, 123, and 443 (see the information about external communications). 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 cannot 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 (HA) 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 the installation instructions for your Linux platform.

### 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 the installation instructions for VMware.

### Guidelines for appliance Storage Nodes

If enterprise networking policies restrict access to any of these ports, you can remap ports using the StorageGRID Appliance Installer. For more information about port remapping for appliances, see the installation instructions for your storage 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.
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.
1509	TCP	All nodes	Archive Nodes	Archive Node internal traffic.
1511	TCP	All nodes	Storage Nodes	Metadata internal traffic.

5353	UDP	All nodes	All nodes	Optionally used for full-grid IP changes and for primary Admin Node discovery during installation, expansion, and recovery.
7001	TCP	Storage Nodes	Storage Nodes	Cassandra TLS inter-node cluster communication.
7443	TCP	All Nodes	Admin Nodes	Internal traffic for maintenance procedures and error reporting.
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 messages from E-Series SANtricity System Manager to the primary Admin Node.
11139	TCP	Archive/Storage Nodes	Archive/Storage Nodes	Internal traffic between Storage Nodes and Archive Nodes.
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.
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.
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](#)

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

[Install Ubuntu or Debian](#)

[Install VMware](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

## 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.

If enterprise networking policies restrict access to any of the ports, you can use load balancer endpoints to allow access on user-defined ports. The untrusted Client Networks feature can be used to allow access only on load balancer endpoint ports.



To use systems and protocols such as SMTP, DNS, SSH, or DHCP, you must remap ports when deploying nodes. However, you should not remap balancer endpoints. For information about port remapping, see the installation instructions for your platform.

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. For more information, see the instructions for configuring 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 domain name system.
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.

Port	TCP or UDP	Protocol	From	To	Details
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.
111	TCP/ UDP	RPCBind	NFS client	Admin Nodes	Used by NFS-based audit export (portmap).  <b>Note:</b> This port is required only if NFS-based audit export is enabled.
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.
137	UDP	NetBIOS	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support.  <b>Note:</b> This port is required only if SMB-based audit export is enabled.
138	UDP	NetBIOS	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support.  <b>Note:</b> This port is required only if SMB-based audit export is enabled.

Port	TCP or UDP	Protocol	From	To	Details
139	TCP	SMB	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support.  <b>Note:</b> This port is required only if SMB-based audit export is enabled.
161	TCP/ UDP	SNMP	SNMP client	All nodes	Used for SNMP polling. All nodes provide basic information; Admin Nodes also provide alert and alarm data. Defaults to UDP port 161 when configured.  <b>Note:</b> 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.  <b>Note:</b> 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.  <b>Note:</b> 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.  <b>Note:</b> 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.
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	Archive Nodes	Amazon S3	Used for accessing Amazon S3 from Archive Nodes.
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.
445	TCP	SMB	SMB client	Admin Nodes	Used by SMB-based audit export.  <b>Note:</b> This port is required only if SMB-based audit export is enabled.
903	TCP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export ( <code>rpc.mountd</code> ).  <b>Note:</b> This port is required only if NFS-based audit export is enabled.
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.

Port	TCP or UDP	Protocol	From	To	Details
2049	TCP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export (nfs).  <b>Note:</b> This port is required only if NFS-based audit export is enabled.
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.
8082	TCP	HTTPS	S3 clients	Gateway Nodes	S3-related external traffic to Gateway Nodes (HTTPS).
8083	TCP	HTTPS	Swift clients	Gateway Nodes	Swift-related external traffic to Gateway Nodes (HTTPS).
8084	TCP	HTTP	S3 clients	Gateway Nodes	S3-related external traffic to Gateway Nodes (HTTP).
8085	TCP	HTTP	Swift clients	Gateway Nodes	Swift-related external traffic to Gateway Nodes (HTTP).

Port	TCP or UDP	Protocol	From	To	Details
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.
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.  <b>Note:</b> This port is required only if certificate-based Prometheus access is enabled.
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-related external traffic to Storage Nodes (HTTPS).
18083	TCP	HTTPS	Swift clients	Storage Nodes	Swift-related external traffic to Storage Nodes (HTTPS).
18084	TCP	HTTP	S3 clients	Storage Nodes	S3-related external traffic to Storage Nodes (HTTP).
18085	TCP	HTTP	Swift clients	Storage Nodes	Swift-related external traffic to Storage Nodes (HTTP).

## Related information

[Internal grid node communications](#)

Install Red Hat Enterprise Linux or CentOS

Install Ubuntu or Debian

Install VMware

SG100 & SG1000 services appliances

SG6000 storage appliances

SG5700 storage appliances

SG5600 storage appliances

# Install and upgrade software

## Install Red Hat Enterprise Linux or CentOS

Learn how to install StorageGRID software in Red Hat Enterprise Linux or CentOS deployments.

- [Installation overview](#)
- [Planning and preparation](#)
- [Deploying virtual grid nodes](#)
- [Configuring the grid and completing installation](#)
- [Automating the installation](#)
- [Overview of the installation REST API](#)
- [Where to go next](#)
- [Troubleshooting installation issues](#)
- [Example /etc/sysconfig/network-scripts](#)

### Installation overview

Installing a StorageGRID system in a Red Hat Enterprise Linux (RHEL) or CentOS Linux environment includes three primary steps.

1. **Preparation:** During planning and preparation, you perform the following tasks:
  - Learn about the hardware and storage requirements for StorageGRID.
  - Learn about the specifics of StorageGRID networking so you can configure your network appropriately. For more information, see the StorageGRID networking guidelines.
  - Identify and prepare the physical or virtual servers you plan to use to host your StorageGRID grid nodes.
  - On the servers you have prepared:
    - Install Linux
    - Configure the host network
    - Configure host storage
    - Install Docker
    - Install the StorageGRID host services
2. **Deployment:** Deploy grid nodes using the appropriate user interface. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.
  - a. Use the Linux command line and node configuration files to deploy software-based grid nodes on the hosts you prepared in step 1.
  - b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

3. **Configuration:** When all nodes have been deployed, use the StorageGRID Grid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system. See also the information about the following alternative approaches:

- Use a standard orchestration framework such as Ansible, Puppet, or Chef to install RHEL or CentOS, configure networking and storage, install Docker and the StorageGRID host service, and deploy virtual grid nodes.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).
- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

#### Related information

[Planning and preparation](#)

[Deploying virtual grid nodes](#)

[Configuring the grid and completing installation](#)

[Automating the installation](#)

[Overview of the installation REST API](#)

[Network guidelines](#)

## Planning and preparation

Before deploying grid nodes and configuring the StorageGRID grid, you must be familiar with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operation of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum requirement of having at least three Storage Nodes.

Before starting a StorageGRID installation, you must:

- Understand StorageGRID's compute requirements, including the minimum CPU and RAM requirements for each node.
- Understand how StorageGRID supports multiple networks for traffic separation, security, and administrative convenience, and have a plan for which networks you intend to attach to each StorageGRID node.

See the StorageGRID networking guidelines.

- Understand the storage and performance requirements of each type of grid node.
- 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.
- Understand the requirements for node migration, if you want to perform scheduled maintenance on physical hosts without any service interruption.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the domain name system (DNS) and network time protocol (NTP) servers that will be used.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

- Decide which of the available deployment and configuration tools you want to use.

## Related information

[Network guidelines](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

## Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	You must have a valid, digitally signed NetApp license.  <b>Note:</b> 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	You must 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: <ul style="list-style-type: none"> <li>• Network port</li> <li>• SSH client (for example, PuTTY)</li> <li>• Supported web browser</li> </ul>

Item	Notes
StorageGRID documentation	<ul style="list-style-type: none"> <li>• Release Notes</li> <li>• Instructions for administering StorageGRID</li> </ul>

### Related information

[Downloading and extracting the StorageGRID installation files](#)

[Web browser requirements](#)

[Administer StorageGRID](#)

[Release notes](#)

### Downloading and extracting the StorageGRID installation files

You must download the StorageGRID installation archive and extract the required files.

#### Steps

1. Go to the NetApp Downloads page for StorageGRID.

[NetApp Downloads: 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 check box.

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 check box, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column, select the appropriate software.

Download the `.tgz` or `.zip` archive file for your platform.

The compressed files contain the RPM files and scripts for Red Hat Enterprise Linux or CentOS.



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

7. Save and extract the archive file.
8. Choose the files you need from the following list.

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 or CentOS hosts.
<code>./rpms/StorageGRID-Webscale-Service-version-SHA.rpm</code>	RPM package for installing the StorageGRID host service on your RHEL or CentOS 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>	A sample 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.
<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 or CentOS hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.

## Related information

[Maintain & recover](#)

## CPU and RAM requirements

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

For information about supported servers, see the Interoperability Matrix.

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node

- **RAM:** At least 24 GB per node, and 2 to 16 GB less than the total system RAM, depending on the total RAM available and the amount of non-StorageGRID software running on the system

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 are not 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 the information about storage requirements.

#### **Related information**

[NetApp Interoperability Matrix Tool](#)

[Storage and performance requirements](#)

[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

[Upgrade software](#)

#### **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 are not supported. You can use shared or local RAID storage for any of the storage categories; however, if you want to use StorageGRID's node migration capability, you must store both system data and object data on shared storage.

### 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 AFF storage

If the StorageGRID node uses storage assigned from a NetApp AFF 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, do not 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
Docker storage pool	Container pool	1	Total number of nodes × 100 GB

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
/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  <b>Note:</b> A software-based Storage Node can have 1 to 16 storage volumes; at least 3 storage volumes are recommended.	4,000 GB See <a href="#">Storage requirements for Storage Nodes</a> for more information.
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. As a general rule, 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 cannot be used to restore grid nodes. Instead, refer to the recovery and maintenance 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>
Archive 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	/var/local volume	1	90 GB
Storage Node	Object data	3	4,000 GB
Admin Node	/var/local volume	1	90 GB
Admin Node	Admin Node audit logs	1	200 GB
Admin Node	Admin Node tables	1	200 GB
Gateway Node	/var/local volume	1	90 GB
<b>Total</b>		<b>9</b>	<b>Container pool:</b> 300 GB <b>System data:</b> 670 GB <b>Object data:</b> 12,000 GB


### 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 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 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 are installing a new StorageGRID 11.5 system and each Storage Node has 128 GB or more of RAM, you should 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 the instructions for administering StorageGRID and search for “managing object metadata storage.”

[Administer StorageGRID](#)

#### Related information

[Node container migration requirements](#)

[Maintain & recover](#)

## 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 simply 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 StorageGRID's node migration feature. In this case, you can ignore these additional requirements.

Before performing migration or hypervisor maintenance, shut down the nodes gracefully. See the recovery and maintenance instructions for shutting down a grid node.

### VMware Live Migration not supported

OpenStack Live Migration and VMware live vMotion cause the virtual machine clock time to jump and are not 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 procedure for shutting down a grid node in the recovery and maintenance instructions.

### Consistent network interface names

In order 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

In order 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
```

### Related information

[Configuring the host network](#)

[Maintain & recover](#)

### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## Deployment tools

You might benefit from automating all or part of the StorageGRID installation.

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.

If you are interested in automating all or part of your StorageGRID deployment, review “Automating the installation” before beginning the installation process.

## Related information

[Overview of the installation REST API](#)

[Automating the installation](#)

## Preparing the hosts

You must complete the following steps to prepare your physical or virtual hosts for StorageGRID. Note that you can automate many or all of these steps using standard server configuration frameworks such as Ansible, Puppet, or Chef.

## Related information

## Installing Linux

You must install Red Hat Enterprise Linux or CentOS Linux on all grid hosts. Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

### 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, NetApp recommends selecting the “compute node” software configuration, if available, or “minimal install” base environment. Do not 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.

### Related information

[NetApp Interoperability Matrix Tool](#)

### Configuring the host network

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.

### What you'll need

- You have reviewed the StorageGRID networking guidelines.

[Network guidelines](#)

- You have reviewed the information about node container migration requirements.

[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.

[Considerations and recommendations for MAC address cloning](#)



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 cannot 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 simply 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 do not need to use either of these examples; you can use any approach that meets your needs.



Do not 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

[Network guidelines](#)

[Node container migration requirements](#)

[Creating node configuration files](#)

### Considerations and recommendations for MAC address cloning

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

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 show 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


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](#)

## Configuring host storage

You must allocate block storage volumes to each host.

### What you'll need

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 do not need to partition, format, or mount any of these volumes; you just need to ensure they are visible to the hosts.

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
    }
    ...
}

```

This will cause 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.

#### Related information

[Installing Docker](#)

## Configuring the Docker storage volume

Before installing Docker, you might need to format the Docker storage volume and mount it on `/var/lib/docker`.

### 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.

## Installing Docker

The StorageGRID system runs on Red Hat Enterprise Linux or CentOS as a collection of Docker containers. Before you can install StorageGRID, you must install Docker.

### 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 1.10.3 or later.

```
Client:
  Version: 1.10.3
  API version: 1.22
  Package version: docker-common-1.10.3-46.el7.14.x86_64
  Go version: go1.6.2
  Git commit: 5206701-unsupported
  Built: Mon Aug 29 14:00:01 2016
  OS/Arch: linux/amd64

Server:
  Version: 1.10.3
  API version: 1.22
  Package version: docker-common-1.10.3-46.el7.14.x86_64
  Go version: go1.6.2
  Git commit: 5206701-unsupported
  Built: Mon Aug 29 14:00:01 2016
  OS/Arch: linux/amd64
```

## Related information

[Configuring host storage](#)

## Installing 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 Yum repository metadata included in the installation archive to install the RPM packages remotely. See the Yum 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 yum --nogpgcheck localinstall /tmp/StorageGRID-Webscale-Images-
version-SHA.rpm
```

```
sudo yum --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.

## Deploying virtual grid nodes

To deploy virtual grid nodes on Red Hat Enterprise Linux or CentOS hosts, you create node configuration files for all nodes, validate the files, and start the StorageGRID host service, which starts the nodes. If you need to deploy any StorageGRID appliance Storage Nodes, see the installation and maintenance instructions for the appliance after you have deployed all virtual nodes.

- [Creating node configuration files](#)
- [Validating the StorageGRID configuration](#)
- [Starting the StorageGRID host service](#)

### Related information

[SG100 & SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

### Creating node configuration files

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 are not used for appliance nodes.

#### Where do I put the node configuration files?

You must 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.

#### What do I name the 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 do not 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.

#### What is in a node configuration file?

The configuration files contain key/value pairs, with one key and one value per line. For each key/value pair, you must 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.

Some keys are required for every node, while others are optional or only required for certain node types.

The table defines the acceptable values for all supported keys. In the middle column:

**R**: required

**BP**: best practice

**O**: optional

Key	R, BP, or O?	Value
ADMIN_IP	BP	<p>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.</p> <p>See “How grid nodes discover the primary Admin Node.”</p> <p><b>Note:</b> This value is ignored, and might be prohibited, on the primary Admin Node.</p>
ADMIN_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
ADMIN_NETWORK_ESL	O	<p>Comma-separated list of subnets in CIDR notation to which this node should communicate via the Admin Network gateway.</p> <p>Example: 172.16.0.0/21,172.17.0.0/21</p>
ADMIN_NETWORK_GATEWAY	O (R)	<p>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.</p> <p><b>Note:</b> This parameter is required if ADMIN_NETWORK_ESL is specified.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>

Key	R, BP, or O?	Value
ADMIN_NETWORK_IP	O	<p>IPv4 address of this node on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>
ADMIN_NETWORK_MAC	O	<p>The MAC address for the Admin Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:10</p>
ADMIN_NETWORK_MASK	O	<p>IPv4 netmask for this node, on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 255.255.255.0</li> <li>• 255.255.248.0</li> </ul>

Key	R, BP, or O?	Value
ADMIN_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Admin Network. Do not specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>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.</p> <p><b>IMPORTANT:</b> 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1500</li> <li>• 8192</li> </ul>

Key	R, BP, or O?	Value
ADMIN_NETWORK_TARGET	BP	<p>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.</p> <p><b>Note:</b> Do not 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.</p> <p><b>Best practice:</b> 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• bond0.1002</li> <li>• ens256</li> </ul>
ADMIN_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>

Key	R, BP, or O?	Value
ADMIN_NETWORK_TARGET_TY PE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container use the MAC address of the host host target interface on the Admin Network.</p> <p><b>Best practice:</b> In networks where promiscuous mode would be required, use the ADMIN_NETWORK_TARGET_TY PE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning, see the considerations and recommendations for MAC address cloning.</p> <p><a href="#">Considerations and recommendations for MAC address cloning</a></p>
ADMIN_ROLE	R	<p>Primary or Non-Primary</p> <p>This key is only required when NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p>
BLOCK_DEVICE_AUDIT_LOGS	R	<p>Path and name of the block device special file this node will use for persistent storage of audit logs. This key is only required for nodes with NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</li> <li>• /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</li> <li>• /dev/mapper/sgws-adm1-audit-logs</li> </ul>

Key	R, BP, or O?	Value
BLOCK_DEVICE_RANGEDB_00	R	<p>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>; do not specify it for other node types.</p> <p>Only <code>BLOCK_DEVICE_RANGEDB_00</code> is required; the rest are optional. The block device specified for <code>BLOCK_DEVICE_RANGEDB_00</code> must be at least 4 TB; the others can be smaller.</p> <p><b>Note:</b> Do not leave gaps. If you specify <code>BLOCK_DEVICE_RANGEDB_05</code>, you must also specify <code>BLOCK_DEVICE_RANGEDB_04</code>.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li><code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</code></li> <li><code>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</code></li> <li><code>/dev/mapper/sgws-snl-rangedb-0</code></li> </ul>
BLOCK_DEVICE_RANGEDB_01		
BLOCK_DEVICE_RANGEDB_02		
BLOCK_DEVICE_RANGEDB_03		
BLOCK_DEVICE_RANGEDB_04		
BLOCK_DEVICE_RANGEDB_05		
BLOCK_DEVICE_RANGEDB_06		
BLOCK_DEVICE_RANGEDB_07		
BLOCK_DEVICE_RANGEDB_08		
BLOCK_DEVICE_RANGEDB_09		
BLOCK_DEVICE_RANGEDB_10		
BLOCK_DEVICE_RANGEDB_11		
BLOCK_DEVICE_RANGEDB_12		
BLOCK_DEVICE_RANGEDB_13		
BLOCK_DEVICE_RANGEDB_14		
BLOCK_DEVICE_RANGEDB_15		

Key	R, BP, or O?	Value
BLOCK_DEVICE_TABLES	R	<p>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; do not specify it for other node types.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</li> <li>• /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</li> <li>• /dev/mapper/sgws-adml-tables</li> </ul>
BLOCK_DEVICE_VAR_LOCAL	R	<p>Path and name of the block device special file this node will use for its /var/local persistent storage.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</li> <li>• /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</li> <li>• /dev/mapper/sgws-sn1-var-local</li> </ul>
CLIENT_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED

Key	R, BP, or O?	Value
CLIENT_NETWORK_GATEWAY	O	<p>IPv4 address of the local Client Network gateway for this node, which must be on the subnet defined by CLIENT_NETWORK_IP and CLIENT_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>
CLIENT_NETWORK_IP	O	<p>IPv4 address of this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>
CLIENT_NETWORK_MAC	O	<p>The MAC address for the Client Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:20</p>
CLIENT_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 255.255.255.0</li> <li>• 255.255.248.0</li> </ul>

Key	R, BP, or O?	Value
CLIENT_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Client Network. Do not specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>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.</p> <p><b>IMPORTANT:</b> 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1500</li> <li>• 8192</li> </ul>

Key	R, BP, or O?	Value
CLIENT_NETWORK_TARGET	BP	<p>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.</p> <p><b>Note:</b> Do not 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.</p> <p><b>Best practice:</b> 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• bond0.1003</li> <li>• ens423</li> </ul>
CLIENT_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is only supported value.)</p>

Key	R, BP, or O?	Value
CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Client Network.</p> <p><b>Best practice:</b> In networks where promiscuous mode would be required, use the CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning, see the considerations and recommendations for MAC address cloning.</p> <p><a href="#">Considerations and recommendations for MAC address cloning</a></p>
GRID_NETWORK_CONFIG	BP	<p>STATIC or DHCP</p> <p>(Defaults to STATIC if not specified.)</p>
GRID_NETWORK_GATEWAY	R	<p>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.</p> <p>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.</p>

Key	R, BP, or O?	Value
GRID_NETWORK_IP	R	<p>IPv4 address of this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>
GRID_NETWORK_MAC	O	<p>The MAC address for the Grid Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:30</p>
GRID_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 255.255.255.0</li> <li>• 255.255.248.0</li> </ul>

Key	R, BP, or O?	Value
GRID_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Grid Network. Do not specify if GRID_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>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.</p> <p><b>IMPORTANT:</b> 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.</p> <p><b>IMPORTANT:</b> For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The <b>Grid Network MTU mismatch</b> alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values do not have to be the same for all network types.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1500</li> <li>• 8192</li> </ul>

Key	R, BP, or O?	Value
GRID_NETWORK_TARGET	R	<p>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.</p> <p><b>Note:</b> Do not 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• bond0.1001</li> <li>• ens192</li> </ul>
GRID_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>
GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>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.</p> <p><b>Best practice:</b> In networks where promiscuous mode would be required, use the GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning, see the considerations and recommendations for MAC address cloning.</p> <p><a href="#">Considerations and recommendations for MAC address cloning</a></p>

Key	R, BP, or O?	Value
MAXIMUM_RAM	O	<p>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.</p> <p><b>Note:</b> The RAM value affects a node's actual metadata reserved space. See the instructions for administering StorageGRID for a description of what Metadata Reserved Space is.</p> <p>The format for this field is &lt;number&gt;&lt;unit&gt;, where &lt;unit&gt; can be b, k, m, or g.</p> <p>Examples:</p> <p>24g</p> <p>38654705664b</p> <p><b>Note:</b> If you want to use this option, you must enable kernel support for memory cgroups.</p>
NODE_TYPE	R	<p>Type of node:</p> <ul style="list-style-type: none"> <li>• VM_Admin_Node</li> <li>• VM_Storage_Node</li> <li>• VM_Archive_Node</li> <li>• VM_API_Gateway</li> </ul>

Key	R, BP, or O?	Value
PORT_REMAP	O	<p>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.”</p> <p><b>IMPORTANT:</b> Do not remap the ports you are planning to use to configure load balancer endpoints.</p> <p><b>Note:</b> 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.</p> <p>The format used is: &lt;network type&gt;/&lt;protocol&gt;/&lt;default port used by grid node&gt;/&lt;new port&gt;, where &lt;network type&gt; is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <div style="border: 1px solid #ccc; border-radius: 5px; padding: 10px; margin-top: 10px;"> <pre>PORT_REMAP = client/tcp/18082/443</pre> </div>

Key	R, BP, or O?	Value
PORT_REMAP_INBOUND	O	<p>Remaps inbound communications to the specified port. If you specify PORT_REMAP_INBOUND but do not specify a value for PORT_REMAP, outbound communications for the port are unchanged.</p> <p><b>IMPORTANT:</b> Do not remap the ports you are planning to use to configure load balancer endpoints.</p> <p>The format used is: &lt;network type&gt;/&lt;protocol:&gt;/&lt;remapped port &gt;/&lt;default port used by grid node&gt;, where &lt;network type&gt; is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <pre>PORT_REMAP_INBOUND = grid/tcp/3022/22</pre>

### Related information

[How grid nodes discover the primary Admin Node](#)

[Network guidelines](#)

[Administer StorageGRID](#)

### 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 cannot 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/dc1-adm1.conf`

**Example file contents:**

```

NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
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 Archive Node

**Example file name:** /etc/storagegrid/nodes/dcl-ar1.conf

**Example file contents:**

```
NODE_TYPE = VM_Archive_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dcl-arcl-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.4
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
```

### Validating 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.

```
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

```

## Starting 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
```

For any node that returns a status of “Not-Running” or “Stopped”, run the following command:

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

3. 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
```

## Configuring the grid and completing installation

You complete installation by configuring the StorageGRID system from the Grid Manager on the primary Admin Node.

- [Navigating to the Grid Manager](#)
- [Specifying the StorageGRID license information](#)
- [Adding sites](#)
- [Specifying Grid Network subnets](#)
- [Approving pending grid nodes](#)
- [Specifying Network Time Protocol server information](#)
- [Specifying Domain Name System server information](#)
- [Specifying the StorageGRID system passwords](#)
- [Reviewing your configuration and completing installation](#)
- [Post-installation guidelines](#)

### Navigating to the Grid Manager

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

#### What you'll need

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

#### Steps

1. Open your web browser and navigate to one of the following addresses:

```
https://primary_admin_node_ip
```

```
client_network_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. Click **Install a StorageGRID system.**

The page used to configure a StorageGRID system appears.

NetApp® StorageGRID® Help ▾

Install

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

### Specifying 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 **Grid Name**.

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

2. Click **Browse**, locate the NetApp License File (`NLFunique_id.txt`), and click **Open**.

The license file is validated, and the serial number and licensed storage capacity are 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.

Install



## License

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

Grid Name	<input type="text" value="Grid1"/>
New License File	<input type="button" value="Browse"/>
License Serial Number	<input type="text" value="950719"/>
Storage Capacity (TB)	<input type="text" value="240"/>

3. Click **Next**.

## Adding 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.

Install



## 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	<input type="text" value="Raleigh"/>	✕
Site Name 2	<input type="text" value="Atlanta"/>	+ ✕

3. Click **Next**.

## Specifying 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 via 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.

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.

Install



### 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**.

### Approving pending grid nodes

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

#### What you'll need

All virtual and StorageGRID appliance grid nodes must have been deployed.


#### 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.

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 <input type="text"/>			
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address		
<input checked="" type="radio"/>	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 <input type="text"/>		
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address		
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21		
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21		
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21		
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21		
<input type="radio"/>	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:

## Storage Node Configuration

### General Settings

Site	<input type="text" value="Raleigh"/>
Name	<input type="text" value="NetApp-SGA"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

### Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.16.5.20/21"/>
Gateway	<input type="text" value="172.16.5.20"/>

### Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.224.5.20/21"/>
Gateway	<input type="text" value="10.224.0.1"/>
Subnets (CIDR)	<input type="text" value="10.0.0.0/8"/> <b>x</b>
	<input type="text" value="172.19.0.0/16"/> <b>x</b>
	<input type="text" value="172.21.0.0/16"/> <b>+ x</b>

### Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="47.47.5.20/21"/>
Gateway	<input type="text" value="47.47.0.1"/>

- **Site:** The name of the site with which this grid node will be associated.
- **Name:** The name that will be assigned to the node, and the name that will be displayed in the Grid Manager. The name defaults to the name you specified when you configured the node. During this step of the installation process, you can change the name as required.



After you complete the installation, you cannot change the name of the node.



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.

- **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 cannot 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 resulting 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 resulting 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 cannot 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, reset the node.
- Remove the node from the Pending Nodes table.
- 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.

For additional information, see the installation and maintenance 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 resulting 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 cannot 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, reset 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.

For additional information, see the installation and maintenance 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 Q

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
<i>No results found.</i>				

◀
▶

### 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
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
<input type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21
<input type="radio"/>	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**.

### Specifying 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, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID. See [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 interface. At the top, there is a blue header with "NetApp® StorageGRID®" and a "Help" dropdown. Below the header is a navigation bar with an "Install" button. A progress indicator shows 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 indicator, the "Network Time Protocol" section is displayed. 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". The values entered are: Server 1: 10.60.248.183, Server 2: 10.227.204.142, Server 3: 10.235.48.111, and Server 4: 0.0.0.0. A plus sign (+) is located to the right of the Server 4 input field.

3. Select **Next**.

## Specifying Domain Name System server information

You must specify Domain Name System (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. Specifying at least two DNS servers is recommended.




Provide two to six IPv4 addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see the information about modifying the DNS configuration in the recovery and maintenance instructions.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

### 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 best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

### Specifying 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 may 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.

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with "NetApp® StorageGRID®" and a "Help" dropdown. Below the header is a navigation bar with "Install" and a progress indicator. The progress indicator consists of eight numbered steps: 1. License, 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP, 6. DNS, 7. Passwords (highlighted in blue), and 8. Summary. Below the progress indicator, the "Passwords" section is displayed. It contains the following text: "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." There are four password input fields, each with a label and a masked password field (represented by dots): "Provisioning Passphrase", "Confirm Provisioning Passphrase", "Grid Management Root User Password", and "Confirm Root User Password". At the bottom of the form, there is a checked checkbox labeled "Create random command line passwords."

5. If you are installing a grid for proof of concept or demo purposes, optionally deselect the **Create random command line passwords** check box.

For production deployments, random passwords should always be used for security reasons. Deselect **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**.

## Reviewing your configuration and completing installation

You must carefully review the configuration information you have entered to ensure that the installation completes successfully.

### Steps

1. View the **Summary** page.

**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**

<b>Grid Name</b>	Grid1	<a href="#">Modify License</a>
<b>Passwords</b>	Auto-generated random command line passwords	<a href="#">Modify Passwords</a>

**Networking**

<b>NTP</b>	10.60.248.183 10.227.204.142 10.235.48.111	<a href="#">Modify NTP</a>
<b>DNS</b>	10.224.223.130 10.224.223.136	<a href="#">Modify DNS</a>
<b>Grid Network</b>	172.16.0.0/21	<a href="#">Modify Grid Network</a>

**Topology**

<b>Topology</b>	<b>Atlanta</b>	<a href="#">Modify Sites</a>	<a href="#">Modify Grid Nodes</a>
	<b>Raleigh</b>		
	<a href="#">dc1-adm1</a>	<a href="#">dc1-g1</a>	<a href="#">dc1-s1</a>
	<a href="#">dc1-s2</a>	<a href="#">dc1-s3</a>	<a href="#">NetApp-SGA</a>

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 cannot 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** check box, and click **Next**.

### Download Recovery Package

Before proceeding, you must download the Recovery Package file. This file is necessary to recover the StorageGRID system if a failure occurs.

When the download completes, open the .zip file and confirm it includes a "gpt-backup" directory and a second .zip file. Then, extract this inner .zip file and confirm you can open the passwords.txt file.

After you have verified the contents, copy the Recovery Package 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.

The Recovery Package is required for recovery procedures and must be stored in a secure location.

[Download Recovery Package](#)

I have successfully downloaded and verified the Recovery Package file.

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.

Name	Site	Grid Network IPv4 Address	Progress	Stage
dc1-adm1	Site1	172.16.4.215/21	<div style="width: 100%; background-color: #0070C0;"></div>	Starting services
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%; background-color: #0070C0;"></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 50%; background-color: #0070C0;"></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 20%; background-color: #0070C0;"></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 10%; background-color: #0070C0;"></div>	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 cannot set up DHCP during configuration.



Nodes reboot when their IP addresses change, which can cause outages if a DHCP address 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 the information about configuring IP addresses in the recovery and maintenance instructions.
- 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 re-establish these connections.

## Automating the installation

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.

If you are interested in automating all or part of your StorageGRID deployment, review “Automating the installation” before beginning the installation process.

### Automating 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 can be prepared ahead of time (or programmatically) to enable automated installation. If you already use a standard orchestration framework to install and configure RHEL or CentOS, adding StorageGRID to your playbooks or

recipes should be straightforward.

An 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 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.

## Automating the configuration of StorageGRID

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

### What you'll need

- 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>	Sample 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 sample 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

- 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, you need to extract 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

[Configuring the grid and completing installation](#)

[Overview of the installation REST API](#)

## Overview of the 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 (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs 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 in the event that 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.

## Where to go next

After completing an installation, you must perform a series of integration and configuration steps. Some steps are required; others are optional.

### Required tasks

- Create a tenant account for each client protocol (Swift or S3) 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 or Swift API client applications you will use to upload objects to your StorageGRID system.
- When you are ready, configure the information lifecycle management (ILM) rules and ILM policy you want to use to protect object data.



When you install StorageGRID, the default ILM policy, Baseline 2 Copies Policy, is active. This policy includes the stock ILM rule (Make 2 Copies), and it applies if no other policy has been activated.

- If your installation includes appliance Storage Nodes, use SANtricity software to complete the following

tasks:

- Connect to each StorageGRID appliance.
- Verify receipt of AutoSupport data.
- If your StorageGRID system includes any Archive Nodes, configure the Archive Node's connection to the target external archival storage system.



If any Archive Nodes will use Tivoli Storage Manager as the external archival storage system, you must also configure Tivoli Storage Manager.

- Review and follow the StorageGRID system hardening guidelines to eliminate security risks.
- Configure email notifications for system alerts.

### Optional tasks

- If you want to receive notifications from the (legacy) alarm system, configure mailing lists and email notifications for alarms.
- Update grid node IP addresses if they have changed since you planned your deployment and generated the Recovery Package. See information about changing IP addresses in the recovery and maintenance instructions.
- Configure storage encryption, if required.
- Configure storage compression to reduce the size of stored objects, if required.
- Configure audit client access. You can configure access to the system for auditing purposes through an NFS or a CIFS file share. See the instructions for administering StorageGRID.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

## Troubleshooting 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 the instructions for monitoring and troubleshooting StorageGRID. For help troubleshooting appliance installation issues, see the installation and maintenance instructions for your appliances. If you need additional help, contact technical support.

### Related information

[Monitor & troubleshoot](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[NetApp Support](#)

## 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 Ubuntu or Debian

Learn how to install StorageGRID software in Ubuntu or Debian deployments.

- [Installation overview](#)
- [Planning and preparation](#)
- [Deploying virtual grid nodes](#)
- [Configuring the grid and completing installation](#)
- [Automating the installation](#)
- [Overview of the installation REST API](#)
- [Where to go next](#)
- [Troubleshooting installation issues](#)
- [Example /etc/network/interfaces](#)

### Installation overview

Installing a StorageGRID system in an Ubuntu or Debian environment includes three primary steps.

1. **Preparation:** During planning and preparation, you perform the following tasks:
  - Learn about the hardware and storage requirements for StorageGRID.
  - Learn about the specifics of StorageGRID networking so you can configure your network appropriately. For more information, see the StorageGRID networking guidelines.
  - Identify and prepare the physical or virtual servers you plan to use to host your StorageGRID grid nodes.
  - On the servers you have prepared:
    - Install Ubuntu or Debian
    - Configure the host network
    - Configure host storage
    - Install Docker
    - Install the StorageGRID host services
2. **Deployment:** Deploy grid nodes using the appropriate user interface. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.
  - a. Use the Ubuntu or Debian command line and node configuration files to deploy virtual grid nodes on the hosts you prepared in step 1.
  - b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

3. **Configuration:** When all nodes have been deployed, use the Grid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system in an Ubuntu or Debian environment. See also the information about the following alternative approaches:

- Use a standard orchestration framework such as Ansible, Puppet, or Chef to install Ubuntu or Debian, configure networking and storage, install Docker and the StorageGRID host service, and deploy virtual grid nodes.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).
- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

## Related information

[Planning and preparation](#)

[Deploying virtual grid nodes](#)

[Configuring the grid and completing installation](#)

[Automating the installation and configuration of the StorageGRID host service](#)

## Planning and preparation

Before deploying grid nodes and configuring the StorageGRID grid, you must be familiar with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operation of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum requirement of having at least three Storage Nodes.

Before starting a StorageGRID installation, you must:

- Understand StorageGRID's compute requirements, including the minimum CPU and RAM requirements for each node.
- Understand how StorageGRID supports multiple networks for traffic separation, security, and administrative convenience, and have a plan for which networks you intend to attach to each StorageGRID node.

See the StorageGRID networking guidelines.

- Understand the storage and performance requirements of each type of grid node.
- 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.
- Understand the requirements for node migration, if you want to perform scheduled maintenance on physical hosts without any service interruption.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the domain name system (DNS) and network time protocol (NTP) servers that will be used.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

- Decide which of the available deployment and configuration tools you want to use.

### Related information

## Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	You must have a valid, digitally signed NetApp license.  <b>Note:</b> 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	You must 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: <ul style="list-style-type: none"><li>• Network port</li><li>• SSH client (for example, PuTTY)</li><li>• Supported web browser</li></ul>
StorageGRID documentation	<ul style="list-style-type: none"><li>• Release notes</li><li>• Instructions for administering StorageGRID</li></ul>

## Related information

[Downloading and extracting the StorageGRID installation files](#)

[Web browser requirements](#)

[Administer StorageGRID](#)

[Release notes](#)

## Downloading and extracting the StorageGRID installation files

You must download the StorageGRID installation archive and extract the required files.

### Steps

1. Go to the NetApp Downloads page for StorageGRID.

[NetApp Downloads: 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 check box.

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 check box, and then select **Accept & Continue**.

The downloads page for the version you selected appears. The page contains three columns:

6. In the **Install StorageGRID** column, select the appropriate software.

Select the `.tgz` or `.zip` archive file for your platform.

- `StorageGRID-Webscale-version-DEB-uniqueID.zip`
- `StorageGRID-Webscale-version-DEB-uniqueID.tgz`

The compressed files contain the DEB files and scripts for Ubuntu or Debian.



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

7. Save and extract the archive file.
8. Choose the files you need from the following list.

The set of files you need depends on your planned grid topology and how you will deploy your StorageGRID grid.



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.
<b>Deployment scripting tool</b>	<b>Description</b>

Path and file name	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.
<code>./debs/configure-storagegrid.sample.json</code>	A sample 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.

## Related information

[Maintain & recover](#)

## CPU and RAM requirements

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

For information about supported servers, see the Interoperability Matrix.

Each StorageGRID node requires the following minimum resources:

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

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 are not 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 the information about storage requirements.

### Related information

[NetApp Interoperability Matrix Tool](#)

[Storage and performance requirements](#)

[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

[Upgrade software](#)

### 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 are not supported. You can use shared or local RAID storage for any of the storage categories; however, if you want to use StorageGRID's node migration capability, you must store both system data and object data on shared storage.

### 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 AFF storage

If the StorageGRID node uses storage assigned from a NetApp AFF 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, do not 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
Docker storage pool	Container pool	1	Total number of nodes × 100 GB
<code>/var/local</code> volume	System data	1 for each node on this host	90 GB
Storage Node	Object data	3 for each Storage Node on this host <b>Note:</b> A software-based Storage Node can have 1 to 16 storage volumes; at least 3 storage volumes are recommended.	4,000 GB See storage requirements for Storage Nodes for more information.
Admin Node audit logs	System data	1 for each Admin Node on this host	200 GB

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
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. As a general rule, 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 cannot be used to restore grid nodes. Instead, refer to the recovery and maintenance 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>
Archive 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	4,000 GB
Admin Node	<code>/var/local</code> volume	1	90 GB
Admin Node	Admin Node audit logs	1	200 GB

Type of node	LUN purpose	Number of LUNs	LUN size
Admin Node	Admin Node tables	1	200 GB
Gateway Node	/var/local volume	1	90 GB
<b>Total</b>		<b>9</b>	<b>Container pool: 300 GB</b> <b>System data: 670 GB</b> <b>Object data: 12,000 GB</b>


### 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 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 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 are installing a new StorageGRID 11.5 system and each Storage Node has 128 GB or more of RAM, you should 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 the instructions for administering StorageGRID and search for “managing object metadata storage.”

## [Administer StorageGRID](#)

### **Related information**

[Node container migration requirements](#)

[Maintain & recover](#)

### **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 simply 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 StorageGRID’s node migration feature. In this case, you can ignore these additional requirements.

Before performing migration or hypervisor maintenance, shut down the nodes gracefully. See the recovery and maintenance instructions for shutting down a grid node.

### **VMware Live Migration not supported**

OpenStack Live Migration and VMware live vMotion cause the virtual machine clock time to jump and are not 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 procedure for shutting down a grid node in the recovery and maintenance instructions.

## Consistent network interface names

In order 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

In order 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`

#### Related information

[Configuring the host network](#)

[Maintain & recover](#)

#### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

#### Deployment tools

You might benefit from automating all or part of the StorageGRID installation.

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.

If you are interested in automating all or part of your StorageGRID deployment, review “Automating the installation” before beginning the installation process.

#### Related information

[Automating the installation](#)

#### Preparing the hosts

You must complete the following steps to prepare your physical or virtual hosts for StorageGRID. Note that you can automate many or all of these steps using standard server configuration frameworks such as Ansible, Puppet, or Chef.

#### Related information

[Automating the installation and configuration of the StorageGRID host service](#)

#### Installing Linux

You must install Ubuntu or Debian on all grid hosts. Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

#### Steps

1. Install Ubuntu or Debian on all physical or virtual grid hosts according to the distributor’s instructions or your standard procedure.



Do not 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 unselected.

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.

#### Related information

[NetApp Interoperability Matrix Tool](#)

#### Understanding 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](#).

Once you disable AppArmor, no entries indicating that AppArmor is enabled will appear in the StorageGRID log files.

### Configuring the host network

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.

#### What you'll need

- You have reviewed the StorageGRID networking guidelines.

[Network guidelines](#)

- You have reviewed the information about node container migration requirements.

[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.

[Considerations and recommendations for MAC address cloning](#)



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 cannot 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 simply 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 do not need to use either of these examples; you can use any approach that meets your needs.



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

To enable MAC cloning, see the instructions for creating node configuration files.

[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 show 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<sub>1</sub> through 10G<sub>3</sub> 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.

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

## Related information

[Example /etc/network/interfaces](#)

## Configuring host storage

You must allocate block storage volumes to each host.

## What you'll need

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 do not need to partition, format, or mount any of these volumes; you just need to ensure they are visible to the hosts.

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
    }
    ...
}

```

This will cause 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.

#### Related information

[Storage and performance requirements](#)

[Node container migration requirements](#)

## Configuring the Docker storage volume

Before installing Docker, you might need to format the Docker storage volume and mount it on `/var/lib/docker`.

### 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.

## Installing Docker

The StorageGRID system runs on Linux as a collection of Docker containers. Before you can install StorageGRID, you must install Docker.

### 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 1.10.3 or later.

```
Client:
  Version:      1.10.3
  API version:  1.22
  Go version:   go1.6.1
  Git commit:   20f81dd
  Built:        Wed, 20 Apr 2016 14:19:16 -0700
  OS/Arch:     linux/amd64

Server:
  Version:      1.10.3
  API version:  1.22
  Go version:   go1.6.1
  Git commit:   20f81dd
  Built:        Wed, 20 Apr 2016 14:19:16 -0700
  OS/Arch:     linux/amd64
```

## Related information

[Configuring host storage](#)

## Installing 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.

## Deploying virtual grid nodes

When you deploy grid nodes in an Ubuntu or Debian environment, you create node configuration files for all nodes, validate the files, and start the StorageGRID host service, which starts the nodes. If you need to deploy any StorageGRID appliance Storage Nodes, see the installation and maintenance instructions for the appliance after you have deployed all virtual nodes.

- [Creating node configuration files](#)
- [Validating the StorageGRID configuration](#)
- [Starting the StorageGRID host service](#)

### Related information

[SG100 & SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

### Creating node configuration files

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 are not used for appliance nodes.

#### Where do I put the node configuration files?

You must 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.

#### What do I name the 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 do not 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>-<node type>-<node number>.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.

#### What is in a node configuration file?

The configuration files contain key/value pairs, with one key and one value per line. For each key/value pair, you must 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.

Some keys are required for every node, while others are optional or only required for certain node types.

The table defines the acceptable values for all supported keys. In the middle column:

**R:** required

**BP:** best practice

**O:** optional

Key	R, BP, or O?	Value
ADMIN_IP	BP	<p>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.</p> <p>See “How grid nodes discover the primary Admin Node.”</p> <p><b>Note:</b> This value is ignored, and might be prohibited, on the primary Admin Node.</p>
ADMIN_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED
ADMIN_NETWORK_ESL	O	<p>Comma-separated list of subnets in CIDR notation to which this node should communicate via the Admin Network gateway.</p> <p>Example: 172.16.0.0/21,172.17.0.0/21</p>
ADMIN_NETWORK_GATEWAY	O (R)	<p>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.</p> <p><b>Note:</b> This parameter is required if ADMIN_NETWORK_ESL is specified.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>

Key	R, BP, or O?	Value
ADMIN_NETWORK_IP	O	<p>IPv4 address of this node on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>
ADMIN_NETWORK_MAC	O	<p>The MAC address for the Admin Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:10</p>
ADMIN_NETWORK_MASK	O	<p>IPv4 netmask for this node, on the Admin Network. This key is only required when ADMIN_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 255.255.255.0</li> <li>• 255.255.248.0</li> </ul>

Key	R, BP, or O?	Value
ADMIN_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Admin Network. Do not specify if ADMIN_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>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.</p> <p><b>IMPORTANT:</b> 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1500</li> <li>• 8192</li> </ul>

Key	R, BP, or O?	Value
ADMIN_NETWORK_TARGET	BP	<p>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.</p> <p><b>Note:</b> Do not 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.</p> <p><b>Best practice:</b> 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• bond0.1002</li> <li>• ens256</li> </ul>
ADMIN_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>

Key	R, BP, or O?	Value
ADMIN_NETWORK_TARGET_TY PE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container use the MAC address of the host host target interface on the Admin Network.</p> <p><b>Best practice:</b> In networks where promiscuous mode would be required, use the ADMIN_NETWORK_TARGET_TY PE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning, see the considerations and recommendations for MAC address cloning.</p> <p><a href="#">Considerations and recommendations for MAC address cloning</a></p>
ADMIN_ROLE	R	<p>Primary or Non-Primary</p> <p>This key is only required when NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p>
BLOCK_DEVICE_AUDIT_LOGS	R	<p>Path and name of the block device special file this node will use for persistent storage of audit logs. This key is only required for nodes with NODE_TYPE = VM_Admin_Node; do not specify it for other node types.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</li> <li>• /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</li> <li>• /dev/mapper/sgws-adm1-audit-logs</li> </ul>

Key	R, BP, or O?	Value
BLOCK_DEVICE_RANGEDB_00	R	<p>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>; do not specify it for other node types.</p> <p>Only <code>BLOCK_DEVICE_RANGEDB_00</code> is required; the rest are optional. The block device specified for <code>BLOCK_DEVICE_RANGEDB_00</code> must be at least 4 TB; the others can be smaller.</p> <p><b>Note:</b> Do not leave gaps. If you specify <code>BLOCK_DEVICE_RANGEDB_05</code>, you must also specify <code>BLOCK_DEVICE_RANGEDB_04</code>.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li><code>/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</code></li> <li><code>/dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</code></li> <li><code>/dev/mapper/sgws-snl-rangedb-0</code></li> </ul>
BLOCK_DEVICE_RANGEDB_01		
BLOCK_DEVICE_RANGEDB_02		
BLOCK_DEVICE_RANGEDB_03		
BLOCK_DEVICE_RANGEDB_04		
BLOCK_DEVICE_RANGEDB_05		
BLOCK_DEVICE_RANGEDB_06		
BLOCK_DEVICE_RANGEDB_07		
BLOCK_DEVICE_RANGEDB_08		
BLOCK_DEVICE_RANGEDB_09		
BLOCK_DEVICE_RANGEDB_10		
BLOCK_DEVICE_RANGEDB_11		
BLOCK_DEVICE_RANGEDB_12		
BLOCK_DEVICE_RANGEDB_13		
BLOCK_DEVICE_RANGEDB_14		
BLOCK_DEVICE_RANGEDB_15		

Key	R, BP, or O?	Value
BLOCK_DEVICE_TABLES	R	<p>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; do not specify it for other node types.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</li> <li>• /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</li> <li>• /dev/mapper/sgws-adml-tables</li> </ul>
BLOCK_DEVICE_VAR_LOCAL	R	<p>Path and name of the block device special file this node will use for its /var/local persistent storage.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0</li> <li>• /dev/disk/by-id/wwn-0x600a09800059d6df000060d757b475fd</li> <li>• /dev/mapper/sgws-sn1-var-local</li> </ul>
CLIENT_NETWORK_CONFIG	O	DHCP, STATIC, or DISABLED

Key	R, BP, or O?	Value
CLIENT_NETWORK_GATEWAY	O	<p>IPv4 address of the local Client Network gateway for this node, which must be on the subnet defined by CLIENT_NETWORK_IP and CLIENT_NETWORK_MASK. This value is ignored for DHCP-configured networks.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>
CLIENT_NETWORK_IP	O	<p>IPv4 address of this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>
CLIENT_NETWORK_MAC	O	<p>The MAC address for the Client Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:20</p>
CLIENT_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Client Network. This key is only required when CLIENT_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 255.255.255.0</li> <li>• 255.255.248.0</li> </ul>

Key	R, BP, or O?	Value
CLIENT_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Client Network. Do not specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>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.</p> <p><b>IMPORTANT:</b> 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1500</li> <li>• 8192</li> </ul>

Key	R, BP, or O?	Value
CLIENT_NETWORK_TARGET	BP	<p>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.</p> <p><b>Note:</b> Do not 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.</p> <p><b>Best practice:</b> 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• bond0.1003</li> <li>• ens423</li> </ul>
CLIENT_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is only supported value.)</p>

Key	R, BP, or O?	Value
CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>Set the key to "true" to cause the StorageGRID container to use the MAC address of the host target interface on the Client Network.</p> <p><b>Best practice:</b> In networks where promiscuous mode would be required, use the CLIENT_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning, see the considerations and recommendations for MAC address cloning.</p> <p><a href="#">Considerations and recommendations for MAC address cloning</a></p>
GRID_NETWORK_CONFIG	BP	<p>STATIC or DHCP</p> <p>(Defaults to STATIC if not specified.)</p>
GRID_NETWORK_GATEWAY	R	<p>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.</p> <p>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.</p>

Key	R, BP, or O?	Value
GRID_NETWORK_IP	R	<p>IPv4 address of this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1.1.1.1</li> <li>• 10.224.4.81</li> </ul>
GRID_NETWORK_MAC	O	<p>The MAC address for the Grid Network interface in the container.</p> <p>This field is optional. If omitted, a MAC address will be generated automatically.</p> <p>Must be 6 pairs of hexadecimal digits separated by colons.</p> <p>Example: b2:9c:02:c2:27:30</p>
GRID_NETWORK_MASK	O	<p>IPv4 netmask for this node on the Grid Network. This key is only required when GRID_NETWORK_CONFIG = STATIC; do not specify it for other values.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 255.255.255.0</li> <li>• 255.255.248.0</li> </ul>

Key	R, BP, or O?	Value
GRID_NETWORK_MTU	O	<p>The maximum transmission unit (MTU) for this node on the Grid Network. Do not specify if GRID_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1500 is used.</p> <p>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.</p> <p><b>IMPORTANT:</b> 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.</p> <p><b>IMPORTANT:</b> For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The <b>Grid Network MTU mismatch</b> alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values do not have to be the same for all network types.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• 1500</li> <li>• 8192</li> </ul>

Key	R, BP, or O?	Value
GRID_NETWORK_TARGET	R	<p>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.</p> <p><b>Note:</b> Do not 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.</p> <p>Examples:</p> <ul style="list-style-type: none"> <li>• bond0.1001</li> <li>• ens192</li> </ul>
GRID_NETWORK_TARGET_TYPE	O	<p>Interface</p> <p>(This is the only supported value.)</p>
GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC	BP	<p>True or False</p> <p>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.</p> <p><b>Best practice:</b> In networks where promiscuous mode would be required, use the GRID_NETWORK_TARGET_TYPE_INTERFACE_CLONE_MAC key instead.</p> <p>For more details on MAC cloning, see the considerations and recommendations for MAC address cloning.</p> <p><a href="#">Considerations and recommendations for MAC address cloning</a></p>

Key	R, BP, or O?	Value
MAXIMUM_RAM	O	<p>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.</p> <p><b>Note:</b> The RAM value affects a node's actual metadata reserved space. See the instructions for administering StorageGRID for a description of what Metadata Reserved Space is.</p> <p>The format for this field is &lt;number&gt;&lt;unit&gt;, where &lt;unit&gt; can be b, k, m, or g.</p> <p>Examples:</p> <p>24g</p> <p>38654705664b</p> <p><b>Note:</b> If you want to use this option, you must enable kernel support for memory cgroups.</p>
NODE_TYPE	R	<p>Type of node:</p> <ul style="list-style-type: none"> <li>• VM_Admin_Node</li> <li>• VM_Storage_Node</li> <li>• VM_Archive_Node</li> <li>• VM_API_Gateway</li> </ul>

Key	R, BP, or O?	Value
PORT_REMAP	O	<p>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.”</p> <p><b>IMPORTANT:</b> Do not remap the ports you are planning to use to configure load balancer endpoints.</p> <p><b>Note:</b> 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.</p> <p>The format used is: &lt;network type&gt;/&lt;protocol&gt;/&lt;default port used by grid node&gt;/&lt;new port&gt;, where network type is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <div style="border: 1px solid #ccc; border-radius: 10px; padding: 10px; background-color: #f9f9f9; margin-top: 10px;"> <pre>PORT_REMAP = client/tcp/18082/443</pre> </div>

Key	R, BP, or O?	Value
PORT_REMAP_INBOUND	O	<p>Remaps inbound communications to the specified port. If you specify PORT_REMAP_INBOUND but do not specify a value for PORT_REMAP, outbound communications for the port are unchanged.</p> <p><b>IMPORTANT:</b> Do not remap the ports you are planning to use to configure load balancer endpoints.</p> <p>The format used is: &lt;network type&gt;/&lt;protocol:&gt;/&lt;remapped port &gt;/&lt;default port used by grid node&gt;, where network type is grid, admin, or client, and protocol is tcp or udp.</p> <p>For example:</p> <pre>PORT_REMAP_INBOUND = grid/tcp/3022/22</pre>

### Related information

[How grid nodes discover the primary Admin Node](#)

[Network guidelines](#)

[Administer StorageGRID](#)

### 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 cannot 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/dc1-adm1.conf

**Example file contents:**

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-adm1-var-local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/dc1-adm1-audit-logs
BLOCK_DEVICE_TABLES = /dev/mapper/dc1-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/dc1-sn1.conf

**Example file contents:**

```
NODE_TYPE = VM_Storage_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-sn1-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/dc1-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/dc1-sn1-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/dc1-sn1-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/dc1-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 Archive Node

**Example file name:** /etc/storagegrid/nodes/dc1-ar1.conf

**Example file contents:**

```
NODE_TYPE = VM_Archive_Node
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-ar1-var-local
GRID_NETWORK_TARGET = bond0.1001
ADMIN_NETWORK_TARGET = bond0.1002
CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_IP = 10.1.0.4
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.1.0.1
```

### Example for Gateway Node

**Example file name:** /etc/storagegrid/nodes/dc1-gw1.conf

**Example file contents:**

```
NODE_TYPE = VM_API_Gateway
ADMIN_IP = 10.1.0.2
BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/dc1-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/dc1-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
```

### Validating 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.

```
Checking for misnamed node configuration files... PASSED
Checking configuration file for node dcl-adm1... PASSED
Checking configuration file for node dcl-gw1... PASSED
Checking configuration file for node dcl-sn1... PASSED
Checking configuration file for node dcl-sn2... PASSED
Checking configuration file for node dcl-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

```

## Starting 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
```

For any node that returns a status of “Not Running” or “Stopped”, run the following command:

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

3. 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
```

## Configuring the grid and completing installation

You complete installation by configuring the StorageGRID system from the Grid Manager on the primary Admin Node.

- [Navigating to the Grid Manager](#)
- [Specifying the StorageGRID license information](#)
- [Adding sites](#)
- [Specifying Grid Network subnets](#)
- [Approving pending grid nodes](#)
- [Specifying Network Time Protocol server information](#)
- [Specifying Domain Name System server information](#)
- [Specifying the StorageGRID system passwords](#)
- [Reviewing your configuration and completing installation](#)
- [Post-installation guidelines](#)

### Navigating to the Grid Manager

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

#### What you'll need

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

#### Steps

1. Open your web browser and navigate to one of the following addresses:

```
https://primary_admin_node_ip  
  
client_network_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.

### 1. Click **Install a StorageGRID system**.

The page used to configure a StorageGRID grid appears.

NetApp® StorageGRID® Help ▾

Install

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

### Specifying 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 **Grid Name**.

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

2. Click **Browse**, locate the NetApp License File (NLUnique\_id.txt), and click **Open**.

The license file is validated, and the serial number and licensed storage capacity are 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.



### License

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

Grid Name	<input type="text" value="Grid1"/>
New License File	<input type="button" value="Browse"/>
License Serial Number	<input type="text" value="950719"/>
Storage Capacity (TB)	<input type="text" value="240"/>

3. Click **Next**.

### Adding 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.

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.



### 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	<input type="text" value="Raleigh"/>	<input type="button" value="✕"/>
Site Name 2	<input type="text" value="Atlanta"/>	<input type="button" value="+ ✕"/>

3. Click **Next**.

## Specifying 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 via 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.

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.

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with the text "NetApp® StorageGRID®" and a "Help" dropdown menu. Below the header is a navigation bar with an "Install" button and a progress indicator consisting of eight numbered steps: 1. License, 2. Sites, 3. Grid Network (highlighted in blue), 4. Grid Nodes, 5. NTP, 6. DNS, 7. Passwords, and 8. Summary. Below the progress bar, the "Grid Network" section is displayed. It contains a paragraph of instructions: "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." Below this is a note: "Note: You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway." The main content area shows a "Subnet 1" label followed by a text input field containing "172.16.0.0/21" and a plus sign (+) to its right. Below the input field is a button labeled "Discover Grid Network subnets".

3. Click **Next**.

## Approving pending grid nodes

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

### What you'll need

All virtual and StorageGRID appliance grid nodes must have been deployed.


### 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.

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.

<input type="button" value="+ Approve"/> <input type="button" value="✕ Remove"/>		<input type="text" value="Search"/> <input type="button" value="Q"/>				
<input type="radio"/>	Grid Network MAC Address <i>↑↓</i>	Name <i>↑↓</i>	Type <i>↑↓</i>	Platform <i>↑↓</i>	Grid Network IPv4 Address <i>↓</i>	
<input checked="" type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Storage Node	StorageGRID Appliance	172.16.5.20/21	
<input type="button" value="◀"/> <input type="button" value="▶"/>						

#### 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.

<input type="button" value="✎ Edit"/> <input type="button" value="↺ Reset"/> <input type="button" value="✕ Remove"/>		<input type="text" value="Search"/> <input type="button" value="Q"/>				
<input type="radio"/>	Grid Network MAC Address <i>↑↓</i>	Name <i>↑↓</i>	Site <i>↑↓</i>	Type <i>↑↓</i>	Platform <i>↑↓</i>	Grid Network IPv4 Address <i>↓</i>
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
<input type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21
<input type="button" value="◀"/> <input type="button" value="▶"/>						

3. Click **Approve**.

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

## Storage Node Configuration

### General Settings

Site	<input type="text" value="Raleigh"/>
Name	<input type="text" value="NetApp-SGA"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

### Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.16.5.20/21"/>
Gateway	<input type="text" value="172.16.5.20"/>

### Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.224.5.20/21"/>
Gateway	<input type="text" value="10.224.0.1"/>
Subnets (CIDR)	<input type="text" value="10.0.0.0/8"/> <b>x</b>
	<input type="text" value="172.19.0.0/16"/> <b>x</b>
	<input type="text" value="172.21.0.0/16"/> <b>+ x</b>

### Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="47.47.5.20/21"/>
Gateway	<input type="text" value="47.47.0.1"/>

- **Site:** The name of the site with which this grid node will be associated.
- **Name:** The name that will be assigned to the node, and the name that will be displayed in the Grid Manager. The name defaults to the name you specified when you configured the node. During this step of the installation process, you can change the name as required.



After you complete the installation, you cannot change the name of the node.



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.

- **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 cannot 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 resulting 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 resulting 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 cannot 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, reset the node.
- Remove the node from the Pending Nodes table.
- 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.

For additional information, see the installation and maintenance 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 resulting 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 cannot 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, reset 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.

For additional information, see the installation and maintenance 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 Q

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
<i>No results found.</i>				

◀
▶

### 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
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
<input type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21
<input type="radio"/>	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**.

### Specifying 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, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

[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. The progress bar at the top indicates the current step is 5, NTP. Below the progress bar, the steps are: 1 License, 2 Sites, 3 Grid Network, 4 Grid Nodes, 5 NTP, 6 DNS, 7 Passwords, and 8 Summary. The NTP step is highlighted in blue. Below the progress bar, the text reads: "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." There are four input fields for Server 1, Server 2, Server 3, and Server 4. The IP addresses entered are 10.60.248.183, 10.227.204.142, 10.235.48.111, and 0.0.0.0. A plus sign is visible next to the Server 4 input field.

3. Select **Next**.

## Related information

[Network guidelines](#)

## Specifying Domain Name System server information

You must specify Domain Name System (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. Specifying at least two DNS servers is recommended.



Provide two to six IPv4 addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see the information about modifying the DNS configuration in the recovery and maintenance instructions.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

## 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. The progress bar at the top indicates the current step is 6, DNS. Below the progress bar, the steps are: 1 License, 2 Sites, 3 Grid Network, 4 Grid Nodes, 5 NTP, 6 DNS, 7 Passwords, and 8 Summary. The DNS step is highlighted in blue. Below the progress bar, the text 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." Below this text, there are two text boxes for DNS servers. The first is labeled "Server 1" and contains the IP address "10.224.223.130". The second is labeled "Server 2" and contains the IP address "10.224.223.136". To the right of the "Server 1" box is a minus sign icon, and to the right of the "Server 2" box is a plus sign icon.

The best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

## Specifying 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 may 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.

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there 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, 7. Passwords (highlighted in blue), and 8. Summary. Below the progress bar, the "Passwords" section is displayed. It contains the following text: "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." There are four password input fields, each with a label and a masked password field (represented by dots): "Provisioning Passphrase", "Confirm Provisioning Passphrase", "Grid Management Root User Password", and "Confirm Root User Password". At the bottom of the form, there is a checkbox labeled "Create random command line passwords." which is checked.

5. If you are installing a grid for proof of concept or demo purposes, optionally deselect the **Create random command line passwords** check box.

For production deployments, random passwords should always be used for security reasons. Deselect

**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**.

## Reviewing your configuration and completing 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	<a href="#">Modify License</a>
Passwords	Auto-generated random command line passwords	<a href="#">Modify Passwords</a>

**Networking**

NTP	10.60.248.183 10.227.204.142 10.235.48.111	<a href="#">Modify NTP</a>
DNS	10.224.223.130 10.224.223.136	<a href="#">Modify DNS</a>
Grid Network	172.16.0.0/21	<a href="#">Modify Grid Network</a>

**Topology**

Topology	Atlanta	<a href="#">Modify Sites</a>	<a href="#">Modify Grid Nodes</a>
	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 cannot 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** check box, and click **Next**.

### Download Recovery Package

Before proceeding, you must download the Recovery Package file. This file is necessary to recover the StorageGRID system if a failure occurs.

When the download completes, open the .zip file and confirm it includes a "gpt-backup" directory and a second .zip file. Then, extract this inner .zip file and confirm you can open the passwords.txt file.

After you have verified the contents, copy the Recovery Package 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.

The Recovery Package is required for recovery procedures and must be stored in a secure location.

[Download Recovery Package](#)

I have successfully downloaded and verified the Recovery Package file.

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.

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 cannot set up DHCP during configuration.



Nodes reboot when their IP addresses change, which can cause outages if a DHCP address 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 the information about configuring IP addresses in the recovery and maintenance instructions.
- 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 re-establish these connections.

## Automating the installation

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.

### Automating 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.

### Automating the configuration of StorageGRID

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

#### What you'll need

- 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>	Sample 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 sample 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, you need to extract 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

[Configuring the grid and completing installation](#)

[Overview of the installation REST API](#)

## Overview of the 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 (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs 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 in the event that 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.

### Related information

[Automating the installation](#)

### Where to go next

After completing an installation, you must perform a series of integration and configuration steps. Some steps are required; others are optional.

### Required tasks

- Create a tenant account for each client protocol (Swift or S3) 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 or Swift API client applications you will use to upload objects to your StorageGRID system.
- When you are ready, configure the information lifecycle management (ILM) rules and ILM policy you want to use to protect object data.



When you install StorageGRID, the default ILM policy, Baseline 2 Copies Policy, is active. This policy includes the stock ILM rule (Make 2 Copies), and it applies if no other policy has been activated.

- If your installation includes appliance Storage Nodes, use SANtricity software to complete the following tasks:
  - Connect to each StorageGRID appliance.
  - Verify receipt of AutoSupport data.
- If your StorageGRID system includes any Archive Nodes, configure the Archive Node's connection to the target external archival storage system.



If any Archive Nodes will use Tivoli Storage Manager as the external archival storage system, you must also configure Tivoli Storage Manager.

- Review and follow the StorageGRID system hardening guidelines to eliminate security risks.
- Configure email notifications for system alerts.

### Optional tasks

- If you want to receive notifications from the (legacy) alarm system, configure mailing lists and email notifications for alarms.
- Update grid node IP addresses if they have changed since you planned your deployment and generated the Recovery Package. See information about changing IP addresses in the recovery and maintenance instructions.
- Configure storage encryption, if required.
- Configure storage compression to reduce the size of stored objects, if required.
- Configure audit client access. You can configure access to the system for auditing purposes through an NFS or a CIFS file share. See the instructions for administering StorageGRID.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

### Troubleshooting 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 the instructions for monitoring and troubleshooting StorageGRID. For help troubleshooting appliance installation issues, see the installation and maintenance instructions for your appliances. If you need additional help, contact technical support.

### Related information

[Monitor & troubleshoot](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[NetApp Support](#)

## 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 VMware

Learn how to install StorageGRID in VMware deployments.

- [Installation overview](#)
- [Planning and preparation](#)
- [Deploying virtual machine grid nodes in VMware vSphere Web Client](#)
- [Configuring the grid and completing installation](#)
- [Automating the installation](#)
- [Overview of the installation REST API](#)
- [Where to go next](#)
- [Troubleshooting installation issues](#)

### Installation overview

Installing a StorageGRID system in a VMware environment includes three primary steps.

1. **Preparation:** During planning and preparation, you perform the following tasks:
  - Learn about the hardware, software, virtual machine, storage, and performance requirements for StorageGRID.
  - Learn about the specifics of StorageGRID networking so you can configure your network appropriately. For more information, see the StorageGRID networking guidelines.
  - Identify and prepare the physical servers you plan to use to host your StorageGRID grid nodes.
  - On the servers you have prepared:
    - Install VMware vSphere Hypervisor
    - Configure the ESX hosts
    - Install and configure VMware vSphere and vCenter

2. **Deployment:** Deploy grid nodes using the VMware vSphere Web Client. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.
  - a. 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.
  - b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

3. **Configuration:** When all nodes have been deployed, use the StorageGRIDGrid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system in a VMware environment. See also the information about the following alternative approaches:

- Use the `deploy-vsphere-ovftool.sh` Bash script (available from the installation archive) to deploy grid nodes in VMware vSphere.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).
- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

## Related information

[Planning and preparation](#)

[Deploying virtual machine grid nodes in VMware vSphere Web Client](#)

[Configuring the grid and completing installation](#)

[Automating the installation](#)

[Overview of the installation REST API](#)

[Network guidelines](#)

## Planning and preparation

Before deploying grid nodes and configuring the StorageGRID grid, you must be familiar with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operational functionality of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum requirement of having at least three Storage Nodes.

Before starting the node deployment and grid configuration procedure, you must:

- Plan the StorageGRID deployment.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

- Understand the available network options and how each network option should be implemented on grid nodes. See the StorageGRID networking guidelines.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the domain name system (DNS) and network time protocol (NTP) servers that will be used.
- Decide which of the available deployment and configuration tools you want to use.

### Related information

[Network guidelines](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

### Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	You must have a valid, digitally signed NetApp license.  <b>Note:</b> The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product.
StorageGRID installation archive for VMware	You must download the StorageGRID installation archive and extract the files.
VMware software and documentation	During installation, you deploy virtual grid nodes on virtual machines in VMware vSphere Web Client. For supported versions, see the Interoperability Matrix.

Item	Notes
Service laptop	<p>The StorageGRID system is installed through a service laptop. The service laptop must have:</p> <ul style="list-style-type: none"> <li>• Network port</li> <li>• SSH client (for example, PuTTY)</li> <li>• Supported web browser</li> </ul>
StorageGRID documentation	<ul style="list-style-type: none"> <li>• Release Notes</li> <li>• Instructions for administering StorageGRID</li> </ul>

### Related information

[NetApp Interoperability Matrix Tool](#)

[Downloading and extracting the StorageGRID installation files](#)

[Web browser requirements](#)

[Administer StorageGRID](#)

[Release notes](#)

### Downloading and extracting the StorageGRID installation files

You must download the StorageGRID installation archives and extract the files.

#### Steps

1. Go to the NetApp Downloads page for StorageGRID.

[NetApp Downloads: 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 check box.

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 check box, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column, select the appropriate software.

Download the `.tgz` or `.zip` archive file for your platform.

- `StorageGRID-Webscale-version-VMware-uniqueID.zip`
- `StorageGRID-Webscale-version-VMware-uniqueID.tgz`



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

7. Save and extract the archive file.
8. Choose the files you need from the following list.

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-archive.ovf</code> <code>./vsphere/vsphere-archive.mf</code>	The template file ( <code>.ovf</code> ) and manifest file ( <code>.mf</code> ) for deploying Archive 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>	A sample 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 is enabled.
<code>./vsphere/configure-storagegrid.sample.json</code>	A sample 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.

## Related information

[Maintain & recover](#)

## Software requirements

You can use a virtual machine to host any type of StorageGRID grid node. One virtual machine is required for each grid node installed on the VMware server.

### 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 grid nodes.

For supported versions of VMware vSphere Hypervisor and VMware vCenter Server software, see the [Interoperability Matrix](#).

For the steps required to install these VMware products, see the [VMware documentation](#).

## Related information

[NetApp Interoperability Matrix Tool](#)

## CPU and RAM requirements

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

For information about supported servers, see the Interoperability Matrix.

Each StorageGRID node requires the following minimum resources:

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

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 are not 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 the information about storage requirements.

## Related information

[NetApp Interoperability Matrix Tool](#)

[Storage and performance requirements](#)

[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

[Upgrade software](#)

## 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 AFF storage

If you are deploying a StorageGRID node as a virtual machine with storage assigned from a NetApp AFF 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.



In a production deployment, do not run more than one Storage Node on a single virtual machine server. Using a dedicated virtual machine 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 virtual machine host, or they can be deployed on their own dedicated virtual machine hosts as required. However, if you have multiple nodes of the same type (two Gateway Nodes, for example), do not install all instances on the same virtual machine host.

### Storage requirements by node type

In a production environment, the virtual machines for StorageGRID grid nodes must meet different requirements, depending on the types of nodes.



Disk snapshots cannot be used to restore grid nodes. Instead, refer to the recovery and maintenance 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  <b>Note:</b> 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.
Gateway Node	100 GB LUN for OS
Archive 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. As a general rule, 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 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 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 are installing a new StorageGRID 11.5 system and each Storage Node has 128 GB or more of RAM, you should 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 the instructions for administering StorageGRID and search for “managing object metadata storage.”

[Administer StorageGRID](#)

#### Related information

[Maintain & recover](#)

#### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## Deploying virtual machine grid nodes in VMware vSphere Web Client

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 networks. If you need to deploy any StorageGRID appliance Storage Nodes, see the installation and maintenance instructions for the appliance after you have deployed all virtual machine grid nodes.

- [Collecting information about your deployment environment](#)
- [How grid nodes discover the primary Admin Node](#)
- [Deploying a StorageGRID node as a virtual machine](#)

### Related information

[SG100 & SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

### Collecting information about your deployment environment

Before deploying grid nodes, you must collect information about your network configuration and VMware environment.

### 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 grid 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.
- If you are not using DHCP, the required networking details for each grid node (IP address, gateway, and network mask).
- If you are not 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).
- 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).

### **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
- Domain Name System (DNS) server IP addresses

### Related information

[How grid nodes discover the primary Admin Node](#)

[Storage and performance requirements](#)

### 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 cannot 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.

### Deploying 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. Optionally, you can remap node ports or increase CPU or memory settings for the node before powering it on.

### What you'll need

- You have reviewed the planning and preparation topics, and you understand the requirements for software, CPU and RAM, and storage and performance.

[Planning and preparation](#)

- 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 do not 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-<em>SG-version</em>-SHA.vmdk
```

- You have the `.ovf` and `.mf` files for each type of grid node you are deploying:

Filename	Description
<code>vsphere-primary-admin.ovf</code>	The template file and manifest file for the primary Admin Node.
<code>vsphere-primary-admin.mf</code>	
<code>vsphere-non-primary-admin.ovf</code>	The template file and manifest file for a non-primary Admin Node.
<code>vsphere-non-primary-admin.mf</code>	
<code>vsphere-archive.ovf</code>	The template file and manifest file for an Archive Node.
<code>vsphere-archive.mf</code>	
<code>vsphere-gateway.ovf</code>	The template file and manifest file for a Gateway Node.
<code>vsphere-gateway.mf</code>	
<code>vsphere-storage.ovf</code>	The template file and manifest file for a Storage Node.
<code>vsphere-storage.mf</code>	

- The `.vmdk`, `.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 virtual machine server.



In a production deployment, do not run more than one Storage Node on a single virtual machine server. Using a dedicated virtual machine host for each Storage Node provides an isolated failure domain.

- 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.
  - [Expand your grid](#)

- [Maintain & recover](#)

- If you are deploying a StorageGRID node as a virtual machine with storage assigned from a NetApp AFF 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.

### 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, Gateway Nodes, and Archive Nodes.

If you are installing a new StorageGRID system:

- You must deploy the primary Admin Node before you deploy any other grid node.
- 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 the primary Admin Node over the Grid Network.

If you need to remap any of the node's ports, do not 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  
sphere-primary-admin.ovf  
sphere-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 may provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).

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.
- If you use the Admin Network, select a different destination network in the vSphere environment. If you do not 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 do not use the Client Network, select the same destination you selected for the Grid Network.

7. Under **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. 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.

c. 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.

- d. In the **Admin Network (eth1)** section, select STATIC, DHCP, or DISABLED for the **Admin network IP configuration**.
  - If you do not 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.
- e. In the **Client Network (eth2)** section, select STATIC, DHCP, or DISABLED for the **Client network IP configuration**.
  - If you do not 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 may 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.

## Networking guidelines



Do not 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. Scroll back to the top of the table, and select **Edit**.
- ii. On the Type tab, select **User configurable**, and select **Save**.
- iii. Select **Set Value**.
- iv. 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
```

- v. 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 do not specify a value for PORT\_REMAP, outbound communications for the port are unchanged.

- i. Scroll back to the top of the table, and select **Edit**.
- ii. On the Type tab, select **User configurable**, and select **Save**.
- iii. Select **Set Value**.
- iv. 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
```

v. 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.

## Configuring the grid and completing installation

You complete installation by configuring the StorageGRID system from the Grid Manager on the primary Admin Node.

- [Navigating to the Grid Manager](#)
- [Specifying the StorageGRID license information](#)
- [Adding sites](#)
- [Specifying Grid Network subnets](#)
- [Approving pending grid nodes](#)
- [Specifying Network Time Protocol server information](#)
- [Specifying Domain Name System server information](#)
- [Specifying the StorageGRID system passwords](#)
- [Reviewing your configuration and completing installation](#)
- [Post-installation guidelines](#)

### Navigating to the Grid Manager

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

## StorageGRID system.

### What you'll need

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

### Steps

1. Open your web browser and navigate to one of the following addresses:

```
https://primary_admin_node_ip
```

```
client_network_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. Click **Install a StorageGRID system**.

The page used to configure a StorageGRID grid appears.

NetApp® StorageGRID® Help ▾

Install

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

### Specifying 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 **Grid Name**.

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

2. Click **Browse**, locate the NetApp License File (NLUnique\_id.txt) and click **Open**.

The license file is validated, and the serial number and licensed storage capacity are 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.

NetApp® StorageGRID® Help ▾

Install

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	<input type="text" value="Grid1"/>
New License File	<input type="button" value="Browse"/>
License Serial Number	<input type="text" value="950719"/>
Storage Capacity (TB)	<input type="text" value="240"/>

3. Click **Next**.

### Adding 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.

Install



## 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	<input type="text" value="Raleigh"/>	✕
Site Name 2	<input type="text" value="Atlanta"/>	+ ✕

3. Click **Next**.

## Specifying 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 via 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.

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.

Install



### 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**.

### Approving pending grid nodes

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

#### What you'll need

All virtual and StorageGRID appliance grid nodes must have been deployed.


#### 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.

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 <input type="text"/>		
Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address		
<a href="#">50:6b:4b:42:d7:00</a>	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 <input type="text"/>		
Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address			
<a href="#">00:50:56:87:42:ff</a>	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21			
<a href="#">00:50:56:87:c0:16</a>	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21			
<a href="#">00:50:56:87:79:ee</a>	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21			
<a href="#">00:50:56:87:db:9c</a>	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21			
<a href="#">00:50:56:87:62:38</a>	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:

## Storage Node Configuration

### General Settings

Site	<input type="text" value="Raleigh"/>
Name	<input type="text" value="NetApp-SGA"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

### Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.16.5.20/21"/>
Gateway	<input type="text" value="172.16.5.20"/>

### Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.224.5.20/21"/>
Gateway	<input type="text" value="10.224.0.1"/>
Subnets (CIDR)	<input type="text" value="10.0.0.0/8"/> <b>x</b>
	<input type="text" value="172.19.0.0/16"/> <b>x</b>
	<input type="text" value="172.21.0.0/16"/> <b>+ x</b>

### Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="47.47.5.20/21"/>
Gateway	<input type="text" value="47.47.0.1"/>

- **Site:** The name of the site with which this grid node will be associated.
- **Name:** The name that will be assigned to the node, and the name that will be displayed in the Grid Manager. The name defaults to the name you specified when you configured the node. During this step of the installation process, you can change the name as required.



After you complete the installation, you cannot change the name of the node.



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.

- **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 cannot 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 resulting 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 resulting 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 cannot 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, reset the node.
- Remove the node from the Pending Nodes table.
- 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.

For additional information, see the installation and maintenance 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 resulting 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 cannot 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, reset 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.

For additional information, see the installation and maintenance 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 Q

Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
<i>No results found.</i>				

◀
▶

### 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
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
<input type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21
<input type="radio"/>	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**.

### Specifying 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, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in high-accuracy environments, such as StorageGRID.

### [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.

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there 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 (highlighted in blue), 6. DNS, 7. Passwords, and 8. Summary. Below the progress bar, the "Network Time Protocol" section is visible. 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". A plus sign (+) is located to the right of the Server 4 input field.

3. Select **Next**.

### Specifying Domain Name System server information

You must specify Domain Name System (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. Specifying at least two DNS servers is recommended.



Provide two to six IPv4 addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see the information about modifying the DNS configuration in the recovery and maintenance instructions.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

## 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. The progress bar at the top indicates the current step is 6, DNS. Below the progress bar, the steps are: 1 License, 2 Sites, 3 Grid Network, 4 Grid Nodes, 5 NTP, 6 DNS, 7 Passwords, and 8 Summary. The DNS step is highlighted in blue. Below the progress bar, the text 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." Below this text, there are two text input fields. The first field is labeled "Server 1" and contains the IP address "10.224.223.130". To the right of this field is a red "x" icon. The second field is labeled "Server 2" and contains the IP address "10.224.223.136". To the right of this field are a red "+" icon and 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**.

## Related information

[Maintain & recover](#)

## Specifying 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 may 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.



### 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	<input type="password"/>
Confirm Provisioning Passphrase	<input type="password"/>
Grid Management Root User Password	<input type="password"/>
Confirm Root User Password	<input type="password"/>

Create random command line passwords.

- If you are installing a grid for proof of concept or demo purposes, optionally deselect the **Create random command line passwords** check box.

For production deployments, random passwords should always be used for security reasons. Deselect **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**.

## Reviewing your configuration and completing installation

You must carefully review the configuration information you have entered to ensure that the installation completes successfully.

### Steps

- View the **Summary** page.

The screenshot shows the NetApp StorageGRID Summary page. At the top, there is a blue header with the NetApp StorageGRID logo 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, 7. Passwords, and 8. Summary. Step 8 is highlighted in blue, indicating it is the current page.

Below the progress bar, the page title is "Summary". A paragraph of text reads: "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."

The page is divided into three sections: General Settings, Networking, and Topology.

**General Settings**

Grid Name	Grid1	<a href="#">Modify License</a>
Passwords	Auto-generated random command line passwords	<a href="#">Modify Passwords</a>

**Networking**

NTP	10.60.248.183 10.227.204.142 10.235.48.111	<a href="#">Modify NTP</a>
DNS	10.224.223.130 10.224.223.136	<a href="#">Modify DNS</a>
Grid Network	172.16.0.0/21	<a href="#">Modify Grid Network</a>

**Topology**

Topology	Atlanta	<a href="#">Modify Sites</a>	<a href="#">Modify Grid Nodes</a>
	Raleigh		
	<a href="#">dc1-adm1</a> <a href="#">dc1-g1</a> <a href="#">dc1-s1</a> <a href="#">dc1-s2</a> <a href="#">dc1-s3</a> <a href="#">NetApp-SGA</a>		

- 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 cannot 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** check box, and click **Next**.

### Download Recovery Package

Before proceeding, you must download the Recovery Package file. This file is necessary to recover the StorageGRID system if a failure occurs.

When the download completes, open the .zip file and confirm it includes a "gpt-backup" directory and a second .zip file. Then, extract this inner .zip file and confirm you can open the passwords.txt file.

After you have verified the contents, copy the Recovery Package 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.

The Recovery Package is required for recovery procedures and must be stored in a secure location.

[Download Recovery Package](#)

I have successfully downloaded and verified the Recovery Package file.

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.

Name	Site	Grid Network IPv4 Address	Progress	Stage
dc1-adm1	Site1	172.16.4.215/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Starting services
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%;"><div style="width: 100%;"></div></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 100%;"><div style="width: 75%;"></div></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 100%;"><div style="width: 25%;"></div></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 100%;"><div style="width: 25%;"></div></div>	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 cannot set up DHCP during configuration.



Nodes reboot when their IP addresses change, which can cause outages if a DHCP address 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 the information about configuring IP addresses in the recovery and maintenance instructions.
- 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 re-establish these connections.

## Automating the installation

You can automate the deployment of VMware virtual grid nodes, the configuration of grid nodes, and the configuration of StorageGRID appliances.

- [Automating grid node deployment in VMware vSphere](#)
- [Automating the configuration of StorageGRID](#)

## Automating grid node deployment in VMware vSphere

You can automate the deployment of StorageGRID grid nodes in VMware vSphere.

### What you'll need

- You have access to a Linux/Unix system with Bash 3.2 or later.
- You have VMware OVF Tool 4.1 installed and correctly configured.
- You know the username and password required to access VMware vSphere using the OVF Tool.

- 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 configuration file. See [Collecting information about your deployment environment](#) for information.
- You have access to the following files from the VMware installation archive for StorageGRID:

Filename	Description
<code>NetApp-SG-version-SHA.vmdk</code>	The virtual machine disk file that is used as a template for creating grid node virtual machines.  <b>Note:</b> This file must be in the same folder as the <code>.ovf</code> and <code>.mf</code> files.
<code>vsphere-primary-admin.ovf</code> <code>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-non-primary-admin.ovf</code> <code>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-archive.ovf</code> <code>vsphere-archive.mf</code>	The template file ( <code>.ovf</code> ) and manifest file ( <code>.mf</code> ) for deploying Archive Nodes.
<code>vsphere-gateway.ovf</code> <code>vsphere-gateway.mf</code>	The template file ( <code>.ovf</code> ) and manifest file ( <code>.mf</code> ) for deploying Gateway Nodes.
<code>vsphere-storage.ovf</code> <code>vsphere-storage.mf</code>	The template file ( <code>.ovf</code> ) and manifest file ( <code>.mf</code> ) for deploying virtual machine-based Storage Nodes.
<code>deploy-vsphere-ovftool.sh</code>	The Bash shell script used to automate the deployment of virtual grid nodes.
<code>deploy-vsphere-ovftool-sample.ini</code>	The sample configuration file for use with the <code>deploy-vsphere-ovftool.sh</code> script.

### Defining 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 a sample configuration file, so that you do not have to create the file from scratch.

### Steps

1. Make a copy of the sample 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=thin
--datastore='<em>datastore_name</em>'
```

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-<em>version</em>/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 = DHCP
```

- **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 = 8192
```

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 switch port 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 do not 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. 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. Do not 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. 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. Do not 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](#).



Do not 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.

- **PORT\_REMAP\_INBOUND:** Remaps inbound communications for the specified port. If you specify `PORT_REMAP_INBOUND` but do not specify a value for `PORT_REMAP`, outbound communications for the port are unchanged.



Do not 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.

## 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`, `VM_Archive_Node`, or `VM_API_Gateway_Node`
- **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

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 setting is 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
```

## Archive Node

The following additional setting is required for Archive Nodes:

- **NODE\_TYPE:** VM\_Archive\_Node

This example entry is for an Archive Node that is on the Grid and Admin Networks, but not on the Client Network.

```
[DC1-ARC1]
NODE_TYPE = VM_Archive_Node

GRID_NETWORK_IP = 10.1.0.4
ADMIN_NETWORK_IP = 10.3.0.4

ADMIN_IP = 10.1.0.2
```

## 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.

### Related information

[How grid nodes discover the primary Admin Node](#)

[Networking guidelines](#)

### Running 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 grid nodes in VMware vSphere.

### What you'll need

- 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

## Automating the configuration of StorageGRID

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

### What you'll need

- 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>	Sample 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 sample 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` 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, you need to extract 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

[Navigating to the Grid Manager](#)

[Overview of the installation REST API](#)

## Overview of the 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 (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs 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 in the event that 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.

## Where to go next

After completing an installation, you must perform a series of integration and configuration steps. Some steps are required; others are optional.

### 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 each client protocol (Swift or S3) 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 or Swift API client applications you will use to upload objects to your StorageGRID system.
- When you are ready, configure the information lifecycle management (ILM) rules and ILM policy you want to use to protect object data.



When you install StorageGRID, the default ILM policy, Baseline 2 Copies Policy, is active. This policy includes the stock ILM rule (Make 2 Copies), and it applies if no other policy has been activated.

- If your installation includes appliance Storage Nodes, use SANtricity software to complete the following tasks:
  - Connect to each StorageGRID appliance.
  - Verify receipt of AutoSupport data.
- If your StorageGRID system includes any Archive Nodes, configure the Archive Node's connection to the target external archival storage system.



If any Archive Nodes will use Tivoli Storage Manager as the external archival storage system, you must also configure Tivoli Storage Manager.

- Review and follow the StorageGRID system hardening guidelines to eliminate security risks.
- Configure email notifications for system alerts.

## Optional tasks

- If you want to receive notifications from the (legacy) alarm system, configure mailing lists and email notifications for alarms.
- Update grid node IP addresses if they have changed since you planned your deployment and generated the Recovery Package. See information about changing IP addresses in the recovery and maintenance instructions.
- Configure storage encryption, if required.
- Configure storage compression to reduce the size of stored objects, if required.
- Configure audit client access. You can configure access to the system for auditing purposes through an NFS or a CIFS file share. See the instructions for administering StorageGRID.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

## Troubleshooting 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)

To learn how to access the log files, see the instructions for monitoring and troubleshooting StorageGRID. For help troubleshooting appliance installation issues, see the installation and maintenance instructions for your appliances. If you need additional help, contact technical support.

### Related information

[Monitor & troubleshoot](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[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 are not 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.

## Upgrade software

Learn how to upgrade a StorageGRID system to a new release.

- [About StorageGRID 11.5](#)
- [Upgrade planning and preparation](#)
- [Performing the upgrade](#)
- [Troubleshooting upgrade issues](#)

### About StorageGRID 11.5

Before starting an upgrade, review this section to learn about the new features and enhancements in StorageGRID 11.5, determine whether any features have been deprecated or removed, and find out about changes to StorageGRID APIs.

- [What's new in StorageGRID 11.5](#)
- [Removed or deprecated features](#)
- [Changes to the Grid Management API](#)
- [Changes to the Tenant Management API](#)

### What's new in StorageGRID 11.5

StorageGRID 11.5 introduces S3 Object Lock, support for KMIP encryption of data, usability improvements to ILM, a redesigned Tenant Manager user interface, support for decommissioning a StorageGRID site, and an appliance node clone procedure.

#### S3 Object Lock for compliant data

The S3 Object Lock feature in StorageGRID 11.5 is an object-protection solution that is equivalent to S3 Object Lock in Amazon Simple Storage Service (Amazon S3). You can enable the global S3 Object Lock setting for a StorageGRID system to allow S3 tenant accounts to create buckets with S3 Object Lock enabled. The tenant can then use an S3 client application to optionally specify retention and legal hold settings for the objects in those buckets.

S3 Object Lock lets tenant users comply with regulations that require certain objects to be retained for a fixed amount of time or indefinitely.

#### Learn more

- [Manage objects with ILM](#)
- [Use S3](#)
- [Use a tenant account](#)

#### KMS encryption key management

You can now configure one or more external key management servers (KMS) in the Grid Manager to provide encryption keys to StorageGRID services and storage appliances. Each KMS or KMS cluster uses the Key Management Interoperability Protocol (KMIP) to provide an encryption key to the appliance nodes at the associated StorageGRID site. After the appliance volumes are encrypted, you cannot access any data on the appliance unless the node can communicate with the KMS.



If you want to use encryption key management, you must use the StorageGRID Appliance Installer to enable the **Node Encryption** setting for the appliance before you add the appliance to the grid.

#### Learn more

- [Administer StorageGRID](#)

#### Usability enhancements for information lifecycle management (ILM)

- You can now view the total capacity of a storage pool, including the amount of used and free space. You can also see which nodes are included in a storage pool and which ILM rules and Erasure Coding profiles use the storage pool.
- You can now design ILM rules that apply to more than one tenant account.
- When you create an ILM rule for erasure coding, you are now reminded to set the Object Size (MB) advanced filter to greater than 0.2 to ensure that very small objects are not erasure coded.
- The ILM policy interface now ensures that the default ILM rule will be always be used for any objects not matched by another rule. Starting in StorageGRID 11.5, the default rule cannot use any basic or advanced filters and is automatically placed as the last rule in the policy.



If your current ILM policy does not conform to the new requirements, you can continue to use it after you upgrade to StorageGRID 11.5. However, if you attempt to clone a non-conforming policy after you upgrade, you are prompted to select a default rule that does not include filters and you are required to place the default rule at the end of the policy.

- The stock All Storage Nodes storage pool is no longer selected by default when you create a new ILM rule or a new Erasure Coding profile. In addition, you can now remove the All Storage Nodes storage pool as long as it not used in any rule.



Using the All Storage Nodes storage pool is not recommended because this storage pool contains all sites. Multiple copies of an object might be placed on the same site if you use this storage pool with a StorageGRID system that includes more than one site.

- You can now remove the stock Make 2 Copies rule (which uses the All Storage Nodes storage pool) as long as it is not used in an active or proposed policy.

- Objects stored in a Cloud Storage Pool can now be deleted immediately (synchronous deletion).

### Learn more

- [Manage objects with ILM](#)

### Enhancements to the Grid Manager

- The redesigned Tenant Accounts page makes it easier to view tenant account usage. The tenant summary table now includes columns for Space Used, Quota Utilization, Quota, and Object Count. A new **View Details** button accesses an overview of each tenant as well as details about the account's S3 buckets or Swift containers. In addition, you can now export two `.csv` files for tenant usage: one containing usage values for all tenants and one containing details about a tenant's buckets or containers.

Related to this change, three new Prometheus metrics were added to track tenant account usage:

- `storagegrid_tenant_usage_data_bytes`
- `storagegrid_tenant_usage_object_count`
- `storagegrid_tenant_usage_quota_bytes`

- The new **Access Mode** field on the Admin Groups page (**Configuration > Access Control**) allows you to specify whether the management permissions for the group are read-write (default) or read-only. Users who belong to a group with read-write access mode can change settings and perform operations in the Grid Manager and the Grid Management API. Users who belong to a group with read-only access mode can only view the settings and features that are selected for the group.



When you upgrade to StorageGRID 11.5, the read-write access mode option is selected for all existing admin groups.

- The AutoSupport user interface was redesigned. You can now configure event-triggered, user-triggered, and weekly AutoSupport messages from a single page in the Grid Manager. You can also configure an additional destination for AutoSupport messages.



If AutoSupport has not been enabled, a reminder message now appears on the Grid ManagerDashboard.

- When viewing the **Storage Used - Object Data** chart on the Nodes page, you can now see estimates for the amount of replicated object data and the amount of erasure-coded data on the grid, site, or Storage Node (**Nodes > grid/site/Storage Node > Storage**).
- Grid Manager menu options were reorganized to make options easier to find. For example, a new **Network Settings** submenu was added to the **Configuration** menu and options in the **Maintenance** and **Support** menus are now listed in alphabetic order.

### Learn more

- [Administer StorageGRID](#)

### Enhancements to the Tenant Manager

- The appearance and organization of the Tenant Manager user interface has been completely redesigned to improve the user experience.
- The new Tenant Manager dashboard provides a high-level summary of each account: it provides bucket details and shows the number of buckets or containers, groups, users, and platform services endpoints (if configured).

## Learn more

- [Use a tenant account](#)

## Client certificates for Prometheus metrics export

You can now upload or generate client certificates (**Configuration > Access Control > Client Certificates**), which can be used to provide secure, authenticated access to the StorageGRID Prometheus database. For example, you can use client certificates if you need to monitor StorageGRID externally using Grafana.

## Learn more

- [Administer StorageGRID](#)

## Load balancer enhancements

- When handling routing requests at a site, the Load Balancer service now performs load aware routing: it considers the CPU availability of the Storage Nodes at the same site. In some cases, information about CPU availability is limited to the site where the Load Balancer service is located.



CPU awareness will be not enabled until at least two-thirds of the Storage Nodes at a site have been upgraded to StorageGRID 11.5 and are reporting CPU statistics.

- For added security, you can now specify a binding mode for each load balancer endpoint. Endpoint pinning lets you restrict the accessibility of each endpoint to specific high availability groups or node interfaces.

## Learn more

- [Administer StorageGRID](#)

## Object metadata changes

- **New Actual reserved space metric:** To help you understand and monitor object metadata space usage on each Storage Node, a new Prometheus metric is shown on the Storage Used - Object Metadata graph for a Storage Node (**Nodes > Storage Node > Storage**).

```
storagegrid_storage_utilization_metadata_reserved
```

The **Actual reserved space** metric indicates how much space StorageGRID has reserved for object metadata on a specific Storage Node.

- **Metadata space increased for installations with larger Storage Nodes:** The system-wide Metadata Reserved Space setting has been increased for StorageGRID systems containing Storage Nodes with 128 GB or more of RAM, as follows:
  - **8 TB for new installations:** If you are installing a new StorageGRID 11.5 system and each Storage Node in the grid has 128 GB or more of RAM, the system-wide Metadata Reserved Space setting is now set to 8 TB instead of 3 TB.
  - **4 TB for upgrades:** If you are upgrading to StorageGRID 11.5 and each Storage Node at any one site has 128 GB or more of RAM, the system-wide Metadata Reserved Space setting is now set to 4 TB instead of 3 TB.

The new values for the Metadata Reserved Space setting increase the allowed metadata space for these larger Storage Nodes, up to 2.64 TB, and ensure that adequate metadata space is reserved for future hardware and software versions.



If your Storage Nodes have enough RAM and sufficient space on volume 0, you can manually increase the Metadata Reserved Space setting up to 8 TB after you upgrade. Reserving additional metadata space after the StorageGRID 11.5 upgrade will simplify future hardware and software upgrades.

#### [Increasing the Metadata Reserved Space setting](#)



If your StorageGRID system stores (or is expected to store) more than 2.64 TB of metadata on any Storage Node, the allowed metadata space can be increased in some cases. If your Storage Nodes each have available free space on storage volume 0 and more than 128 GB of RAM, contact your NetApp account representative. NetApp will review your requirements and increase the allowed metadata space for each Storage Node, if possible.

- **Automatic cleanup of deleted metadata:** When 20% or more of the metadata stored on a Storage Node is ready to be removed (because the corresponding objects were deleted), StorageGRID can now perform an automatic compaction on that Storage Node. This background process only runs if the load on the system is low—that is, when there is available CPU, disk space, and memory. The new compaction process removes metadata for deleted objects sooner than in previous releases and helps to free up space for new objects to be stored.

#### Learn more

- [Administer StorageGRID](#)

#### Changes to S3 REST API support

- You can now use the S3 REST API to specify [S3 Object Lock](#) settings:
  - To create a bucket with S3 Object Lock enabled, use a PUT Bucket request with the `x-amz-bucket-object-lock-enabled` header.
  - To determine if S3 Object Lock is enabled for a bucket, use a GET Object Lock Configuration request.
  - When adding an object version to a bucket with S3 Object Lock enabled, use the following request headers to specify the retention and legal hold settings: `x-amz-object-lock-mode`, `x-amz-object-lock-retain-until-date`, and `x-amz-object-lock-legal-hold`.
- You can now use DELETE Multiple Objects on a versioned bucket.
- You can now use PUT, GET, and DELETE Bucket encryption requests to manage encryption for an existing S3 bucket.
- A minor change was made to a field name for the `Expiration` parameter. This parameter is included in the response to a PUT Object, HEAD Object, or GET Object request if an expiration rule in the lifecycle configuration applies to a specific object. The field that indicates which expiration rule was matched was previously named `rule_id`. This field was renamed to `rule-id` to match the AWS implementation.
- By default, the S3 GET Storage Usage request now attempts to retrieve the storage used by a tenant account and its buckets using strong-global consistency. If strong-global consistency cannot be achieved, StorageGRID attempts to retrieve the usage information using strong-site consistency.
- The `Content-MD5` request header is now correctly supported.

#### Learn more

- [Use S3](#)

## Maximum size for CloudMirror objects increased to 5 TB

The maximum size for objects that can be replicated to a destination bucket by the CloudMirror replication service was increased to 5 TB, which is the maximum object size supported by StorageGRID.

### Learn more

- [Use S3](#)
- [Use Swift](#)

### New alerts added

The following new alerts were added for StorageGRID 11.5:

- Appliance BMC communication error
- Appliance Fibre Channel fault detected
- Appliance Fibre Channel HBA port failure
- Appliance LACP port missing
- Cassandra auto-compactor error
- Cassandra auto-compactor metrics out of date
- Cassandra compactions overloaded
- Disk I/O is very slow
- KMS CA certificate expiration
- KMS client certificate expiration
- KMS configuration failed to load
- KMS connectivity error
- KMS encryption key name not found
- KMS encryption key rotation failed
- KMS is not configured
- KMS key failed to decrypt an appliance volume
- KMS server certificate expiration
- Low free space for storage pool
- Node network reception frame error
- Services appliance storage connectivity degraded
- Storage appliance storage connectivity degraded (previously named Appliance storage connectivity degraded)
- Tenant quota usage high
- Unexpected node reboot

### Learn more

- [Monitor & troubleshoot](#)

### TCP support for SNMP traps

You can now select Transmission Control Protocol (TCP) as the protocol for SNMP trap destinations.

Previously, only the User Datagram Protocol (UDP) protocol was supported.

#### Learn more

- [Monitor & troubleshoot](#)

#### Installation and networking enhancements

- **MAC address cloning:** You can now use MAC address cloning to enhance the security of certain environments. MAC address cloning enables you to use a dedicated virtual NIC for the Grid Network, Admin Network, and Client Network. Having the Docker container use the MAC address of the dedicated NIC on the host allows you to avoid using promiscuous mode network configurations. Three new MAC address cloning keys were added to the node configuration file for Linux-based (bare metal) nodes.
- **Automatic discovery of DNS and NTP host routes:** Previously, there were restrictions on which network your NTP and DNS servers had to connect to, such as the requirement that you could not have all of your NTP and DNS servers on the Client Network. Now, those restrictions are removed.

#### Learn more

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

#### Support for rebalancing erasure-coded (EC) data after Storage Node expansion

The EC rebalance procedure is a new command-line script that might be required after you add new Storage Nodes. When you perform the procedure, StorageGRID redistributes erasure-coded fragments among the existing and the newly expanded Storage Nodes at a site.



You should only perform the EC rebalance procedure in limited cases. For example, if you cannot add the recommended number of Storage Nodes in an expansion, you can use the EC rebalance procedure to allow additional erasure-coded objects to be stored.

#### Learn more

- [Expand your grid](#)

#### New and revised maintenance procedures

- **Site decommission:** You can now remove an operational site from your StorageGRID system. The connected site decommission procedure removes an operational site and preserves data. The new Decommission Site wizard guides you through the process (**Maintenance > Decommission > Decommission Site**).
- **Appliance node cloning:** You can now clone an existing appliance node to upgrade the node to a new appliance model. For example, you can clone a smaller-capacity appliance node to a larger-capacity appliance. You can also clone an appliance node to implement new functionality, such as the new **Node Encryption** setting that is required for the KMS encryption.
- **Ability to change the provisioning passphrase:** You can now change the provisioning passphrase (**Configuration > Access Control > Grid Passwords**). The passphrase is required for recovery, expansion, and maintenance procedures.
- **Enhanced SSH password behavior:** To enhance the security of StorageGRID appliances, the SSH password is no longer changed when you place an appliance into maintenance mode. In addition, new SSH host certificates and host keys are generated when you upgrade a node to StorageGRID 11.5.



If you use SSH to log in to a node after upgrading to StorageGRID 11.5, you will receive a warning that the host key has changed. This behavior is expected and you can safely approve the new key.

#### Learn more

- [Maintain & recover](#)

#### Changes to StorageGRID appliances

- **Direct access to SANtricity System Manager for storage appliances:** You can now access the E-Series SANtricity System Manager user interface from the StorageGRID Appliance Installer and from the Grid Manager. Using these new methods enables access to SANtricity System Manager without using the management port on the appliance. Users who need to access SANtricity System Manager from the Grid Manager must have the new Storage Appliance Administrator permission.
- **Node encryption:** As part of the new KMS encryption feature, a new **Node Encryption** setting was added to the StorageGRID Appliance Installer. If you want to use encryption key management to protect appliance data, you must enable this setting during the hardware configuration stage of appliance installation.
- **UDP port connectivity:** You can now test the network connectivity of a StorageGRID appliance to UDP ports, such as those used for an external NFS or DNS server. From the StorageGRID Appliance Installer, select **Configure Networking > Port Connectivity Test (nmap)**.
- **Automating installation and configuration:** A new JSON configuration upload page was added to the StorageGRID Appliance Installer (**Advanced > Update Appliance Configuration**). This page enables you to use one file to configure multiple appliances in large grids. Additionally, the `configure-sga.py` Python script has been updated to match the capabilities of the StorageGRID Appliance Installer.

#### Learn more

- [SG100 & SG1000 services appliances](#)
- [SG6000 storage appliances](#)
- [SG5700 storage appliances](#)
- [SG5600 storage appliances](#)

#### Changes to audit messages

- **Automatic cleanup of overwritten objects:** Previously, objects that were overwritten were not removed from disk in specific cases, which resulted in additional space consumption. These overwritten objects, which are inaccessible to users, are now automatically removed to save storage space. Refer to the LKCU audit message for more information.
- **New audit codes for S3 Object Lock:** Four new audit codes were added to the SPUT audit message to include [S3 Object Lock](#) request headers:
  - LKEN: Object Lock Enabled
  - LKLH: Object Lock Legal Hold
  - LKMD: Object Lock Retention Mode
  - LKRU: Object Lock Retain Until Date
- **New fields for Last Modified Time and Previous Object Size:** You can now track when an object was overwritten as well as the original object size.
  - The MTME (Last Modified Time) field was added to the following audit messages:

- SDEL (S3 DELETE)
  - SPUT (S3 PUT)
  - WDEL (Swift DELETE)
  - WPUT (Swift PUT)
- The CSIZ (Previous Object Size) field was added to the OVWR (Object Overwrite) audit message.

#### Learn more

- [Review audit logs](#)

#### New nms.requestlog file

A new log file, `/var/local/log/nms.requestlog`, is maintained on all Admin Nodes. This file contains information about outgoing connections from the Management API to internal StorageGRID services.

#### Learn more

- [Monitor & troubleshoot](#)

#### StorageGRID documentation changes

- To make networking information and requirements easier to find and to clarify that the information also applies to StorageGRID appliance nodes, the networking documentation was moved from the software-based installation guides (RedHat Enterprise Linux/CentOS, Ubuntu/Debian, and VMware) to a new networking guide.

##### [Network guidelines](#)

- To make ILM-related instructions and examples easier to find, the documentation for managing objects with information lifecycle management was moved from the *Administrator Guide* to a new ILM guide.

##### [Manage objects with ILM](#)

- A new FabricPool guide provides an overview of configuring StorageGRID as a NetApp FabricPool cloud tier and describes the best practices for configuring ILM and other StorageGRID options for a FabricPool workload.

##### [Configure StorageGRID for FabricPool](#)

- You can now access several instructional videos from the Grid Manager. The current videos provide instructions for managing alerts, custom alerts, ILM rules, and ILM policies.

#### Removed or deprecated features

Some features were removed or deprecated in StorageGRID 11.5. You must review these items to understand whether you need to update client applications or modify your configuration before you upgrade.

##### Weak consistency control removed

The Weak consistency control was removed for StorageGRID 11.5. After you upgrade, the following behaviors will apply:

- Requests to set Weak consistency for an S3 bucket or Swift container will succeed, but the consistency

level will actually be set to Available.

- Existing buckets and containers that use Weak consistency will be silently updated to use Available consistency.
- Requests that have a Weak consistency-control header will actually use Available consistency, if applicable.

The Available consistency control behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations. The Available consistency control offers higher availability for HEAD operations than “read-after-new-write” if Storage Nodes are unavailable.

### Alarm for grid health deprecated

The `/grid/health/topology` API, which checks for active *alarms* on nodes, is deprecated. In its place, a new `/grid/node-health` endpoint was added. This API returns the current status of each node by checking for active *alerts* on nodes.

### Compliance feature deprecated

The S3 Object Lock feature in StorageGRID 11.5 replaces the Compliance feature that was available in previous StorageGRID versions. Because the new S3 Object Lock feature conforms to Amazon S3 requirements, it deprecates the proprietary StorageGRID Compliance feature, which is now referred to as “legacy Compliance.”

If you previously enabled the global Compliance setting, the new global S3 Object Lock setting is enabled automatically when you upgrade to StorageGRID 11.5. Tenant users will no longer be able to create new buckets with Compliance enabled in StorageGRID; however, as required, tenant users can continue to use and manage any existing legacy Compliant buckets.

In the Tenant Manager, a shield icon  indicates a legacy Compliant bucket. Legacy Compliant buckets might also have a hold badge  to indicate that the bucket is under a legal hold.

[KB: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

[Manage objects with ILM](#)

### “S3 multipart part too small” alert removed

The **S3 multipart part too small** alert was removed. Previous, this alert was triggered if an S3 client attempted to complete a multipart upload with parts that did not meet Amazon S3 size limits. After the upgrade to StorageGRID 11.5, any multipart upload requests that do not meet the following size limits will fail:

- 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. Since 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.

### "Appliance link down on Grid Network" alerts removed

The following alerts were removed. If the Grid Network is down, the metrics that would trigger these alerts are not accessible:

- Services appliance link down on Grid Network
- Storage appliance link down on Grid Network

### Support for fully qualified domain name removed from SNMP configuration

When configuring an SNMP server in the baseboard management controller (BMC) for the SG6000, SG100, or SG1000, you must now specify an IP address instead of a fully qualified domain name. If a fully qualified domain name was previously configured, change it to an IP address before upgrading to StorageGRID 11.5.

### Legacy attributes removed

The following legacy attributes were removed. As applicable, equivalent information is provided by Prometheus metrics:

Legacy attribute	Equivalent Prometheus metric
BREC	storagegrid_service_network_received_bytes
BTRA	storagegrid_service_network_transmitted_bytes
CQST	storagegrid_metadata_queries_average_latency_milliseconds
HAIS	storagegrid_http_sessions_incoming_attempted
HCCS	storagegrid_http_sessions_incoming_currently_established
HEIS	storagegrid_http_sessions_incoming_failed
HISC	storagegrid_http_sessions_incoming_successful
LHAC	<i>none</i>
NREC	<i>none</i>
NTSO (Chosen Time Source Offset)	storagegrid_ntp_chosen_time_source_offset_milliseconds
NTRA	<i>none</i>
SLOD	storagegrid_service_load
SMEM	storagegrid_service_memory_usage_bytes
SUTM	storagegrid_service_cpu_seconds
SVUT	storagegrid_service_uptime_seconds

Legacy attribute	Equivalent Prometheus metric
TRBS (Total bits per second received)	<i>none</i>
TRXB	storagegrid_network_received_bytes
TTBS (Total bits per second transmitted)	<i>none</i>
TTXB	storagegrid_network_transmitted_bytes

The following related changes were also made:

- The `network_received_bytes` and `network_transmitted_bytes` Prometheus metrics were changed from gauges to counters because the values of these metrics only increase. If you are currently using these metrics in Prometheus queries, you should start using the `increase()` function in the query.
- The Network Resources table was removed from the Resources tab for StorageGRID services. (Select **Support > Tools > Grid Topology**. Then, select **node > service > Resources**.)
- The HTTP Sessions page was removed for Storage Nodes. Previously, you could access this page by selecting **Support > Tools > Grid Topology** and then selecting **Storage Node > LDR > HTTP**.
- The HCCS (Currently Established Incoming Sessions) alarm was removed.
- The NTSO (Chosen Time Source Offset) alarm was removed.

### Changes to the Grid Management API

StorageGRID 11.5 uses version 3 of the Grid Management API. Version 3 deprecates version 2; however, version 1 and version 2 are still supported.



You can continue to use version 1 and version 2 of the management API with StorageGRID 11.5; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.5, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management` API.

### New client-certificates section

The new section, `/grid/client-certificates`, allows you to configure client certificates to provide secure, authenticated access to the StorageGRID Prometheus database. For example, you can monitor StorageGRID externally using Grafana.

### Legacy compliance endpoints moved to new s3-object-lock section

With the introduction of StorageGRID S3 Object Lock, the APIs used to manage the legacy compliance settings for the grid were moved to a new section of the Swagger user interface. The **s3-object-lock** section includes the two `/grid/compliance-global` API endpoints, which now control the global S3 Object Lock setting. The endpoint URIs remain unchanged for compatibility with existing applications.

## Swift-admin-password Accounts endpoint removed

The following Accounts API endpoint, which was deprecated in StorageGRID 10.4, has now been removed:

```
https://<IP-Address>/api/v1/grid/accounts/<AccountID>/swift-admin-password
```

## New grid-passwords section

The **grid-passwords** section enables operations for grid password management. The section includes two `/grid/change-provisioning-passphrase` API endpoints. The endpoints allow users to change the StorageGRID provisioning passphrase and retrieve the status of the passphrase change.

## storageAdmin permission added to Groups API

The `/grid/groups` API now includes the `storageAdmin` permission.

## New parameter for Storage Usage API

The `GET /grid/accounts/{id}/usage` API now has a `strictConsistency` parameter. To enforce a strong-global consistency when retrieving storage usage information across Storage Nodes, set this parameter to `true`. When this parameter is set to `false` (default), StorageGRID attempts to retrieve usage information using strong-global consistency, but falls back to strong-site consistency if strong-global consistency cannot be met.

## New Node Health API

A new `/grid/node-health` endpoint was added. This API returns the current status of each node by checking for active *alerts* on the nodes. The `/grid/health/topology` API, which checks for active *alarms* on nodes, is deprecated.

## Change to "ApplianceStorageShelvesPowerSupplyDegraded" alert rule ID

The alert rule ID "ApplianceStorageShelvesPowerSupplyDegraded" has been renamed to "ApplianceStorageShelvesDegraded" to better reflect the alert's actual behavior.

## Related information

[Administer StorageGRID](#)

## Changes to the Tenant Management API

StorageGRID 11.5 uses version 3 of the Tenant Management API. Version 3 deprecates version 2; however, version 1 and version 2 are still supported.



You can continue to use version 1 and version 2 of the management API with StorageGRID 11.5; however, support for these versions of the API will be removed in a future release of StorageGRID. After upgrading to StorageGRID 11.5, the deprecated v1 and v2 APIs can be deactivated using the `PUT /grid/config/management` API.

## New parameter for tenant Storage Usage API

The `GET /org/usage` API now has a `strictConsistency` parameter. To enforce a strong-global consistency when retrieving storage usage information across Storage Nodes, set this parameter to `true`.

When this parameter is set to `false` (default), StorageGRID attempts to retrieve usage information using strong-global consistency, but falls back to strong-site consistency if strong-global consistency cannot be met.

### Related information

[Use S3](#)

[Use a tenant account](#)

## Upgrade planning and preparation

You must plan the upgrade of your StorageGRID system to ensure that the system is ready for the upgrade, and that the upgrade can be completed with minimal disruption.

### Steps

1. [Estimating the time to complete an upgrade](#)
2. [How your system is affected during the upgrade](#)
3. [Impact of an upgrade on groups and user accounts](#)
4. [Verifying the installed version of StorageGRID](#)
5. [Obtaining the required materials for a software upgrade](#)
6. [Downloading the StorageGRID upgrade files](#)
7. [Downloading the Recovery Package](#)
8. [Checking the system's condition before upgrading software](#)

### Estimating the time to complete an upgrade

When planning an upgrade to StorageGRID 11.5, you must consider when to upgrade, based on how long the upgrade might take. You must also be aware of which operations you can and cannot 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.



During the upgrade from StorageGRID 11.4 to 11.5, Cassandra database tables on Storage Nodes will be upgraded. The **Upgrade Database** task occurs in the background, but might require an extensive amount of time to complete. While the database is being upgraded, you can safely use new features, apply hotfixes, and perform node recovery operations. However, you might be prevented from performing other maintenance procedures.



If an expansion is urgently required, perform the expansion before upgrading to 11.5.

Upgrade task	Description	Approximate time required	During this task
Start Upgrade Service	Upgrade prechecks are run, the software file is distributed, and the upgrade service is started.	3 minutes per grid node, unless validation errors are reported	As required, you can run the upgrade prechecks manually before the scheduled upgrade maintenance window.
Upgrade Grid Nodes (primary Admin Node)	The primary Admin Node is stopped, upgraded, and restarted.	Up to 30 minutes	You cannot access the primary Admin Node. Connection errors are reported, which you can ignore.
Upgrade Grid Nodes (all other 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 for several minutes each.	15 to 45 minutes per node, with appliance Storage Nodes requiring the most time  <b>Note:</b> For appliance nodes, the StorageGRID Appliance Installer is automatically updated to the latest release.	<ul style="list-style-type: none"> <li>• Do not change the grid configuration.</li> <li>• Do not change the audit level configuration.</li> <li>• Do not update the ILM configuration.</li> <li>• Do not perform another maintenance procedure, such as hotfix, decommission, or expansion.</li> </ul> <p><b>Note:</b> If you need to perform a recovery procedure, contact technical support.</p>
Enable Features	The new features for the new version are enabled.	Less than 5 minutes	<ul style="list-style-type: none"> <li>• Do not change the grid configuration.</li> <li>• Do not change the audit level configuration.</li> <li>• Do not update the ILM configuration.</li> <li>• Do not perform another maintenance procedure.</li> </ul>

Upgrade task	Description	Approximate time required	During this task
Upgrade Database	Cassandra database tables, which exist on all Storage Nodes, are upgraded.	Hours or days, based on the amount of metadata in your system	<p>During the <b>Upgrade Database</b> task, the upgraded grid will operate normally; however, the upgrade will still be in progress. During this task, you can:</p> <ul style="list-style-type: none"> <li>• Use the new features in the new StorageGRID version.</li> <li>• Change the audit level configuration.</li> <li>• Update the ILM configuration.</li> <li>• Apply a hotfix.</li> <li>• Recover a node.</li> </ul> <p><b>Note:</b> You cannot perform a decommission or expansion procedure until the <b>Final Upgrade Steps</b> complete.</p>
Final Upgrade Steps	Temporary files are removed and the upgrade to the new release completes.	5 minutes	When the <b>Final Upgrade Steps</b> task completes, you can perform all maintenance procedures.

## Steps

1. Estimate the time required to upgrade all grid nodes (consider all upgrade tasks except for **Upgrade Database**).
  - a. Multiply the number of nodes in your StorageGRID system by 30 minutes/node (average).
  - 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. Estimate the time required to upgrade the database.
  - a. From the Grid Manager, select **Nodes**.
  - b. Select the first entry in the tree (entire grid), and select the **Storage** tab.
  - c. Hover your cursor over the **Storage Used - Object Metadata** chart, and locate the **Used** value, which indicates how many bytes of object metadata are on your grid.
  - d. Divide the **Used** value by 1.5 TB/day to determine how many days will be needed to upgrade the database.

4. Calculate the total estimated time for the upgrade by adding the results of steps 1, 2, and 3.

#### **Example: Estimating the time to upgrade from StorageGRID 11.4 to 11.5**

Suppose your system has 14 grid nodes, of which 8 are Linux nodes. Also, assume that the **Used** value for object metadata is 6 TB.

1. Multiply 14 by 30 minutes/node and add 1 hour. The estimated time to upgrade all nodes is 8 hours.
2. Multiple 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.
3. Divide 6 by 1.5 TB/day. The estimated number of days for the **Upgrade Database** task is 4 days.



While the **Upgrade Database** task is running, you can safely use new features, apply hotfixes, and perform node recovery operations.

4. Add the values together. You should allow 5 days to complete the upgrade of your system to StorageGRID 11.5.0.

#### **How your system is affected during the upgrade**

You must understand how your StorageGRID system will be affected during the upgrade.

##### **StorageGRID upgrades are non-disruptive**

The StorageGRID system can ingest and retrieve data from client applications throughout the upgrade process. Grid nodes are brought down one at a time during the upgrade, so there is not a time when all grid nodes are unavailable.

To allow for continued availability, you must ensure that objects are stored redundantly using the appropriate ILM policies. You must also ensure that all external S3 or Swift clients are configured to send requests to one of the following:

- A StorageGRID endpoint configured as a high availability (HA) group
- A high availability third-party load balancer
- Multiple Gateway Nodes for each client
- Multiple Storage Nodes for each client

##### **Appliance firmware is upgraded**

During the StorageGRID 11.5 upgrade:

- All StorageGRID appliance nodes are automatically upgraded to StorageGRID Appliance Installer firmware version 3.5.
- SG6060 and SGF6024 appliances are automatically upgraded to BIOS firmware version 3B03.EX and BMC firmware version BMC 3.90.07.
- SG100 and SG1000 appliances are automatically upgraded to BIOS firmware version 3B08.EC and BMC firmware version 4.64.07.

##### **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). 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 11.5 but other nodes are still running StorageGRID 11.4.

In general, these alerts will clear when the upgrade completes.

After the upgrade completes, you can review any upgrade-related alerts by selecting **Recently resolved alerts** from the Grid Manager Dashboard.



During the upgrade to StorageGRID 11.5, the **ILM placement unachievable** alert might be triggered when Storage Nodes are stopped. This alert might persist for 1 day after the upgrade is completed successfully.

### 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, unselect the **Enable SNMP Agent Notifications** check box (**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

Until the **Enable New Feature** task completes:

- Do not make any grid configuration changes.
- Do not change the audit level configuration.
- Do not enable or disable any new features.
- Do not update the ILM configuration. Otherwise, you might experience inconsistent and unexpected ILM behavior.
- Do not apply a hotfix or recover a grid node.

Until the **Final Upgrade Steps** task completes:

- Do not perform an expansion procedure.
- Do not perform a decommission procedure.

### Impact of an upgrade on groups and user accounts

You must understand the impact of the StorageGRID upgrade, so that you can update groups and user accounts appropriately after the upgrade is complete.

### Changes to group permissions and options

After upgrading to StorageGRID 11.5, optionally select the following new permissions and options (**Configuration > Access Control > Admin Groups**).

Permission or option	Description
Storage Appliance Administrator	Required to access the SANtricity System Manager user interface from Grid Manager.

Permission or option	Description
Access Mode	When managing groups, you can select <b>Read-only</b> for this new option to prevent users from changing the settings and features that are selected for the group. Users in groups with read-only access mode can view settings, but they cannot change them.

### Related information

[Administer StorageGRID](#)

### Verifying the installed version of StorageGRID

Before starting the upgrade, you must verify that the previous version of StorageGRID is currently installed with the latest available hotfix applied.

#### Steps

1. Sign in to the Grid Manager using a supported browser.
2. Select **Help > About**.
3. Verify that the **Version** is 11.4.x.y.

In the StorageGRID 11.4.x.y version number:

- The major release has an x value of 0 (11.4.0).
- A minor release, if available, has an x value other than 0 (for example, 11.4.1).
- A hotfix, if available, has a y value (for example, 11.4.0.1).



If you have an earlier version of StorageGRID, you must upgrade to any 11.4 version before upgrading to StorageGRID 11.5. You do not need to be at the highest 11.4 minor-version release to upgrade to StorageGRID 11.5.

4. If you are not at a StorageGRID 11.4 version you must upgrade to version 11.4, one release at a time, using the instructions for each release.

You must also apply the latest hotfix for each StorageGRID version before upgrading to the next level.

One possible upgrade path is shown in the example.

5. Once you are at StorageGRID 11.4, go to the NetApp Downloads page for StorageGRID and see if any hotfixes are available for your StorageGRID 11.4.x version.

[NetApp Downloads: StorageGRID](#)

6. Verify that your StorageGRID 11.4.x version has the latest hotfix applied.
7. If necessary, download and apply the latest StorageGRID 11.4.x.y hotfix for your StorageGRID 11.4.x version.

See the recovery and maintenance instructions for information about applying hotfixes.

### Example: Preparing to upgrade to StorageGRID 11.5 from version 11.3.0.8

The following example shows the upgrade steps to prepare for an upgrade from StorageGRID version 11.3.0.8 to version 11.5. Before you can upgrade to StorageGRID 11.5, your system must have a StorageGRID 11.4 version installed with the latest hotfix.

Download and install software in the following sequence to prepare your system for upgrade:

1. Apply the latest StorageGRID 11.3.0.y hotfix.
2. Upgrade to the StorageGRID 11.4.0 major release. (You do not need to install any 11.4.x minor releases.)
3. Apply the latest StorageGRID 11.4.0.y hotfix.

#### Related information

[Administer StorageGRID](#)

[Maintain & recover](#)

#### Obtaining the required materials for a software upgrade

Before you begin the software upgrade, you must obtain all required materials so you can complete the upgrade successfully.

Item	Notes
StorageGRID upgrade files	You must download the required files to your service laptop: <ul style="list-style-type: none"><li>• <b>All platforms:</b> .upgrade file</li><li>• <b>Any node on Red Hat Enterprise Linux or CentOS:</b> .upgrade file and RPM file (.zip or .tgz)</li><li>• <b>Any node on Ubuntu or Debian:</b> .upgrade file and DEB file (.zip or .tgz)</li></ul>
Service laptop	The service laptop must have: <ul style="list-style-type: none"><li>• Network port</li><li>• SSH client (for example, PuTTY)</li></ul>
Supported web browser	You must confirm that the web browser on the service laptop is supported for use with StorageGRID 11.5.  <a href="#">Web browser requirements</a>  <b>Note:</b> Browser support has changed for StorageGRID 11.5. Confirm you are using a supported version.

Item	Notes
Recovery Package (.zip) file	<p>Before upgrading, you should download the most recent Recovery Package file in case any problems occur during the upgrade.</p> <p>After you upgrade the primary Admin Node, you must 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.</p> <p><a href="#">Downloading the Recovery Package</a></p>
Passwords.txt file	This file is included in the SAID package, which is part of the Recovery Package .zip file. You must obtain the latest version of the Recovery Package.
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	<ul style="list-style-type: none"> <li>• Release notes for StorageGRID 11.5. Be sure to read these carefully before starting the upgrade.</li> <li>• Instructions for administering StorageGRID</li> <li>• If you are upgrading a Linux deployment, the StorageGRID installation instructions for your Linux platform.</li> <li>• Other StorageGRID documentation, as required.</li> </ul>

### Related information

[Web browser requirements](#)

[Administer StorageGRID](#)

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

[Install Ubuntu or Debian](#)

[Install VMware](#)

[Downloading the StorageGRID upgrade files](#)

[Downloading the Recovery Package](#)

[Release notes](#)

### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87

Web browser	Minimum supported version
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

### Downloading the StorageGRID upgrade files

You must download the required files to a service laptop before you upgrade your StorageGRID system.

#### What you'll need

You must have installed all required hotfixes for the StorageGRID software version you are upgrading. See the hotfix procedure in the recovery and maintenance instructions.

#### About this task

You must download the `.upgrade` archive for any platform. If any nodes are deployed on Linux hosts, you must also download an RPM or DEB archive, which you will install before you start the upgrade.

#### Steps

1. Go to the NetApp Downloads page for StorageGRID.

[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 statement appears, read it and select the check box.

This statement appears if there is a required hotfix for the release.

5. Read the End User License Agreement, select the check box, and then select **Accept & Continue**.

The downloads page for the version you selected appears. The page contains three columns:

- Install StorageGRID
- Upgrade StorageGRID
- Support files for StorageGRID Appliances

6. In the **Upgrade StorageGRID** column, select and download the `.upgrade` archive.

Every platform requires the `.upgrade` archive.

7. If any nodes are deployed on Linux hosts, also download the RPM or DEB archive in either `.tgz` or `.zip` format.

You must install the RPM or DEB archive on all Linux nodes before you start the upgrade.



No additional files are required for the SG100 or SG1000.



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

- Red Hat Enterprise Linux or CentOS

`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`

## Related information

[Linux: Installing the RPM or DEB package on all hosts](#)

[Maintain & recover](#)

## Downloading the Recovery Package

The Recovery Package file allows you to restore the StorageGRID system if a failure occurs.

### What you'll need

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

### About this task

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:
  - a. Open the `.zip` file.

- b. Confirm it includes a `gpt-backup` directory and an inner `.zip` file.
  - c. Extract the inner `.zip` file.
  - d. Confirm you can open 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.

### Related information

[Administer StorageGRID](#)

### Checking the system's condition before upgrading software

Before upgrading a StorageGRID system, you must verify the system is ready to accommodate the upgrade. You must 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 browser.
2. Check for and resolve any active alerts.

For information about specific alerts, see the monitoring and troubleshooting instructions.

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 the lists of internal and external ports in the 11.5 version of the networking guidelines, and ensure that all required ports are opened before you upgrade.



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.

### Related information

[Monitor & troubleshoot](#)

[Administer StorageGRID](#)

[Maintain & recover](#)

[Network guidelines](#)

## Performing the upgrade

The Software Upgrade page guides you through the process of uploading the required file and upgrading all of the grid nodes in your StorageGRID system.

### What you'll need

You are aware of the following:

- You must upgrade all grid nodes for all data center sites from the primary Admin Node, using the Grid Manager.
- To detect and resolve issues, you can manually run the upgrade prechecks before starting the actual upgrade. The same prechecks are performed when you start the upgrade. Precheck failures will stop the upgrade process and might require technical support involvement to resolve.
- When you start the upgrade, the primary Admin Node is upgraded automatically.
- After the primary Admin Node has been upgraded, you can select which grid nodes to upgrade next.
- You must upgrade all grid nodes in your StorageGRID system to complete the upgrade, but you can upgrade individual grid nodes in any order. You can select individual grid nodes, groups of grid nodes, or all grid nodes. You can repeat the process of selecting grid nodes as many times as necessary, until all grid nodes at all sites are upgraded.
- When the upgrade starts on a grid node, the services on that node are stopped. Later, the grid node is rebooted. Do not approve the upgrade for a grid node unless you are sure that node is ready to be stopped and rebooted.
- When all grid nodes have been upgraded, new features are enabled and you can resume operations; however, you must wait to perform a decommission or expansion procedure until the background **Upgrade Database** task and the **Final Upgrade Steps** task have completed.
- You must complete the upgrade on the same hypervisor platform you started with.

### Steps

1. [Linux: Installing the RPM or DEB package on all hosts](#)
2. [Starting the upgrade](#)
3. [Upgrading grid nodes and completing the upgrade](#)
4. [Increasing the Metadata Reserved Space setting](#)

### Related information

[Administer StorageGRID](#)

[Estimating the time to complete an upgrade](#)

### Linux: Installing the RPM or DEB package on all hosts

If any StorageGRID nodes are deployed on Linux hosts, you must install an additional RPM or DEB package on each of these hosts before you start the upgrade.

### What you'll need

You must have downloaded one of the following `.tgz` or `.zip` files from the NetApp Downloads page for StorageGRID.



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

Linux platform	Additional file (choose one)
Red Hat Enterprise Linux or CentOS	<ul style="list-style-type: none"><li>• <code>StorageGRID-Webscale-version-RPM-uniqueID.zip</code></li><li>• <code>StorageGRID-Webscale-version-RPM-uniqueID.tgz</code></li></ul>
Ubuntu or Debian	<ul style="list-style-type: none"><li>• <code>StorageGRID-Webscale-version-DEB-uniqueID.zip</code></li><li>• <code>StorageGRID-Webscale-version-DEB-uniqueID.tgz</code></li></ul>

### 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 for your Linux platform.

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

[Install Ubuntu or Debian](#)

The new packages are installed as additional packages. Do not remove the existing packages.

### Starting the upgrade

When you are ready to perform the upgrade, you select the downloaded file and enter the provisioning passphrase. As an option, you can run the upgrade prechecks before performing the actual upgrade.

### What you'll need

You have reviewed all of the considerations and completed all steps in [Upgrade planning and preparation](#).

### Steps

1. Sign in to the Grid Manager using a supported browser.
2. Select **Maintenance > System > Software Update**.

The Software Update page appears.

3. Select **StorageGRID Upgrade**.

The StorageGRID Upgrade page appears and shows the date and time of the most recently completed upgrade, unless the primary Admin Node has been rebooted or the management API restarted since that upgrade was performed.

4. Select the .upgrade file you downloaded.
  - a. Select **Browse**.
  - b. Locate and select the file: `NetApp_StorageGRID_version_Software_uniqueID.upgrade`

c. Select **Open**.

The file is uploaded and validated. When the validation process is done, a green checkmark appears next to the upgrade file name.

5. Enter the provisioning passphrase in the text box.

The **Run Prechecks** and **Start Upgrade** buttons become enabled.

#### StorageGRID Upgrade

Before starting the upgrade process, you must confirm that there are no active alerts and that all grid nodes are online and available.

After uploading the upgrade file, click the Run Prechecks button to detect problems that will prevent the upgrade from starting. These prechecks also run when you start the upgrade.

#### Upgrade file

Upgrade file

Browse

✔ NetApp\_StorageGRID\_11.5.0\_Software\_20210407.2135.8e126f1

Upgrade Version

StorageGRID® 11.5.0

#### Passphrase

Provisioning Passphrase

.....|

Run Prechecks

Start Upgrade

6. If you want to validate the condition of your system before you start the actual upgrade, select **Run Prechecks**. Then, resolve any precheck errors that are reported.



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.



The same prechecks are performed when you select **Start Upgrade**. Selecting **Run Prechecks** allows you to detect and resolve issues before starting the upgrade.

7. When you are ready to perform the upgrade, select **Start Upgrade**.

A warning appears to remind you that your browser's connection will be lost when the primary Admin Node is rebooted. When the primary Admin Node is available again, you need to clear your web browser's cache and reload the Software Upgrade page.

## Connection Will be Temporarily Lost

During the upgrade, your browser's connection to StorageGRID will be lost temporarily when the primary Admin Node is rebooted.

**Attention:** You must clear your cache and reload the page before starting to use the new version. Otherwise, StorageGRID might not respond as expected.

Are you sure you want to start the upgrade process?

Cancel

OK

8. Select **OK** to acknowledge the warning and start the upgrade process.

When the upgrade starts:

a. The upgrade prechecks are run.



If any precheck errors are reported, resolve them and select **Start Upgrade** again.

b. The primary Admin Node is upgraded, which includes stopping services, upgrading the software, and restarting services. You will not be able to 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.



While the primary Admin Node is being upgraded, multiple copies of the following error messages appear, which you can ignore.

## Error

Problem connecting to the server

Unable to communicate with the server. Please reload the page and try again. Contact technical support if the problem persists.

*2 additional copies of this message are not shown.*

OK

## ! Error

503: Service Unavailable

Service Unavailable

The StorageGRID API service is not responding. Please try again later. If the problem persists, contact Technical Support.

*4 additional copies of this message are not shown.*

OK

## ! Error

400: Bad Request

Clear your web browser's cache and reload the page to continue the upgrade.

*2 additional copies of this message are not shown.*

OK

9. After the primary Admin Node has been upgraded, clear your web browser's cache, sign back in, and reload the Software Upgrade page.

For instructions, see the documentation for your web browser.



You must clear the web browser's cache to remove outdated resources used by the previous version of the software.

### Related information

[Upgrade planning and preparation](#)

### Upgrading grid nodes and completing the upgrade

After the primary Admin Node has been upgraded, you must upgrade all other grid nodes in your StorageGRID system. You can customize the upgrade sequence by selecting to upgrade individual grid nodes, groups of grid nodes, or all grid nodes.

### Steps

1. Review the Upgrade Progress section on the Software Upgrade page, which provides information about each major upgrade task.
  - a. **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.

- b. When the **Start Upgrade Service** task is complete, the **Upgrade Grid Nodes** task starts.
  - c. While the **Upgrade Grid Nodes** task is in progress, the Grid Node Status table appears and shows the upgrade stage for each grid node in your system.
2. After the grid nodes appear in the Grid Node Status table, but before approving any grid nodes, download a new copy of the Recovery Package.



You must download a new copy of the Recovery Package file after you upgrade the software version on the primary Admin Node. The Recovery Package file allows you to restore the system if a failure occurs.

3. Review the information in the Grid Node Status table. Grid nodes are arranged in sections by type: Admin Nodes, API Gateway Nodes, Storage Nodes, and Archive Nodes.

## Upgrade Progress

Start Upgrade Service	Completed
Upgrade Grid Nodes	In Progress

### Grid Node Status

You must approve all grid nodes to complete an upgrade, but you can update grid nodes in any order.

During the upgrade of a node, the services on that node are stopped. Later, the node is rebooted. Do not click **Approve** for a node unless you are sure the node is ready to be stopped and rebooted.

When you are ready to add grid nodes to the upgrade queue, click one or more **Approve** buttons to add individual nodes to the queue, click the **Approve All** button at the top of the nodes table to add all nodes of the same type, or click the top-level **Approve All** button to add all nodes in the grid.

If necessary, you can remove nodes from the upgrade queue before node services are stopped by clicking **Remove** or **Remove All**.

**Approve All**

**Remove All**

Admin Nodes

Search

Site	Name	Progress	Stage	Error	Action
Data Center 1	DC1-ADM1	<div style="width: 100%; height: 10px; background-color: green;"></div>	Done		

◀ ▶

Storage Nodes

Approve All Remove All

Search

Site	Name	Progress	Stage	Error	Action
Data Center 1	DC1-S1	<div style="width: 25%; height: 10px; background-color: #00a0e3;"></div>	Waiting for you to approve		<b>Approve</b>
Data Center 1	DC1-S2	<div style="width: 25%; height: 10px; background-color: #00a0e3;"></div>	Waiting for you to approve		<b>Approve</b>
Data Center 1	DC1-S3	<div style="width: 25%; height: 10px; background-color: #00a0e3;"></div>	Waiting for you to approve		<b>Approve</b>

◀ ▶

A grid node can be in one of these stages when this page first appears:

- Done (primary Admin Node only)
- Preparing upgrade

- Software download queued
- Downloading
- Waiting for you to approve

4. Approve the grid nodes you are ready to add to the upgrade queue. Approved nodes of the same type are upgraded one at a time.

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. These operations might cause service interruptions for clients that are communicating with the node. Do not approve the upgrade for a node unless you are sure that node is ready to be stopped and rebooted.

- Select one or more **Approve** buttons to add one or more individual nodes to the upgrade queue.
- Select the **Approve All** button within each section to add all nodes of the same type to the upgrade queue.
- Select the top-level **Approve All** button to add all nodes in the grid to the upgrade queue.

5. If you need to remove a node or all nodes from the upgrade queue, select **Remove** or **Remove All**.

As shown in the example, when the Stage reaches **Stopping services**, the **Remove** button is hidden and you can no longer remove the node.

Storage Nodes		Approve All		Remove All	
Site	Name	Progress	Stage	Error	Action
Data Center 1	DC1-S1	<div style="width: 50%; background-color: #007bff;"></div>	Stopping services		
Data Center 1	DC1-S2	<div style="width: 25%; background-color: #007bff;"></div>	Queued		Remove
Data Center 1	DC1-S3	<div style="width: 25%; background-color: #007bff;"></div>	Queued		Remove

6. Wait for each node to proceed through the upgrade stages, which include Queued, Stopping services, Stopping container, Cleaning up Docker images, Upgrading base OS packages, Rebooting, and Starting services.



When an appliance node reaches the Upgrading base OS packages 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.

When all grid nodes have been upgraded, the **Upgrade Grid Nodes** task is shown as Completed. The remaining upgrade tasks are performed automatically and in the background.

7. As soon as the **Enable Features** task is complete (which occurs quickly), you can start using the new features in the upgraded StorageGRID version.

For example, if you are upgrading to StorageGRID 11.5, you can now enable S3 Object Lock, configure a key management server, or increase the Metadata Reserved Space setting.

### Increasing the Metadata Reserved Space setting

8. Periodically monitor the progress of the **Upgrade Database** task.


During this task, the Cassandra database is upgraded on each Storage Node.



The **Upgrade Database** task might take days to complete. As this background task runs, you can apply hotfixes or recover nodes. However, you must wait for the **Final Upgrade Steps** task to complete before performing an expansion or decommission procedure.

You can review the graph to monitor the progress for each Storage Node.

### Upgrade Progress



9. When the **Upgrade Database** task has completed, wait a few minutes for the **Final Upgrade Steps** task to

complete.

### StorageGRID Upgrade

The new features are enabled and can now be used. While the upgrade background tasks are in progress (which might take an extended time), you can apply hotfixes or recover nodes. You must wait for the upgrade to complete before performing an expansion or decommission.

Status	In Progress
Upgrade Version	11.5.0
Start Time	2021-04-08 09:01:48 MDT

#### Upgrade Progress

Start Upgrade Service	Completed
Upgrade Grid Nodes	Completed
Enable Features	Completed
Upgrade Database	Completed
Final Upgrade Steps	In Progress 

When the Final Upgrade Steps task has completed, the upgrade is done.

10. Confirm that the upgrade completed successfully.
  - a. Sign in to the Grid Manager using a supported browser.
  - b. Select **Help > About**.
  - c. Confirm that the displayed version is what you would expect.
  - d. Select **Maintenance > System > Software Update**. Then, select **StorageGRID Upgrade**.
  - e. Confirm that the green banner shows that the software upgrade was completed on the date and time you expected.

## StorageGRID Upgrade

Before starting the upgrade process, you must confirm that there are no active alerts and that all grid nodes are online and available.

After uploading the upgrade file, click the Run Prechecks button to detect problems that will prevent the upgrade from starting. These prechecks also run when you start the upgrade.

Software upgrade completed at 2021-04-08 12:14:40 MDT.

---

**Upgrade file**

Upgrade file

Upgrade Version No software upgrade file selected

**Passphrase**

Provisioning Passphrase

11. 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.
12. Check the NetApp Downloads page for StorageGRID to see if any hotfixes are available for the StorageGRID version that you just installed.

### [NetApp Downloads: StorageGRID](#)

In the StorageGRID 11.5.x.y version number:

- The major release has an x value of 0 (11.5.0).
  - A minor release, if available, has an x value other than 0 (for example, 11.5.1).
  - A hotfix, if available, has a y value (for example, 11.5.0.1).
13. If available, download and apply the latest hotfix for your StorageGRID version.

See the recovery and maintenance instructions for information about applying hotfixes.

### **Related information**

[Downloading the Recovery Package](#)

[Maintain & recover](#)

### **Increasing the Metadata Reserved Space setting**

After you upgrade to StorageGRID 11.5, 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 must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission or the Grid Topology Page Configuration and Other Grid Configuration permissions.
- You have started the StorageGRID 11.5 upgrade and the **Enable New Features** upgrade task has completed.

### About this task

You might be able to manually increase the system-wide Metadata Reserved Space setting up to 8 TB after upgrading to StorageGRID 11.5. Reserving additional metadata space after the 11.5 upgrade will simplify future hardware and software upgrades.

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, go to the instructions for administering StorageGRID and search for “managing object metadata storage.”

### Administer StorageGRID

#### Steps

1. Sign in to the Grid Manager using a supported browser.
2. Determine the current Metadata Reserved Space setting.
  - a. Select **Configuration > System Settings > Storage Options**.
  - b. In the Storage Watermarks section, note the value of **Metadata Reserved Space**.
3. 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.
4. 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.
5. Update the Metadata Reserved Space setting.
  - a. Select **Configuration > System Settings > 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 **800000000000** (8, followed by 12 zeros)

**Storage Options**

- Overview
- Configuration**

**Configure Storage Options**  
Updated: 2021-02-17 19:40:49 MST

**Object Segmentation**

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1000000000

**Storage Watermarks**

Description	Settings
Storage Volume Read-Write Watermark	30000000000
Storage Volume Soft Read-Only Watermark	10000000000
Storage Volume Hard Read-Only Watermark	5000000000
<b>Metadata Reserved Space</b>	<b>800000000000</b>

Apply Changes

- e. Select **Apply Changes**.

## Troubleshooting upgrade issues

If the upgrade does not complete successfully, you might be able to resolve the issue yourself. If you cannot resolve an issue, you should gather the required information before contacting technical support.

The following sections describe how to recover from situations where the upgrade has partially failed. Contact technical support if you cannot resolve an upgrade issue.

### **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. If you need help, contact technical support.

### **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 are not 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](#)

### **Troubleshooting user interface issues**

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

#### **Web interface does not respond as expected**

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

If you experience issues with the web interface:

- Make sure you are using a supported browser.



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

- Clear your web browser cache.

Clearing the cache removes outdated resources used by the previous version of StorageGRID software, and permits the user interface to operate correctly again. For instructions, see the documentation for your web browser.

## Related information

[Web browser requirements](#)

### “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.
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.  <b>Note:</b> 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.  For more information, see the installation instructions for your Linux platform.
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.

## Related information

Install Red Hat Enterprise Linux or CentOS

Install Ubuntu or Debian

# Install and maintain hardware

## SG6000 storage appliances

Learn how to install and maintain the StorageGRID SG6060 and SGF6024 appliances.

- [SG6000 appliances overview](#)
- [Installation and deployment overview](#)
- [Preparing for installation](#)
- [Installing the hardware](#)
- [Configuring the hardware](#)
- [Deploying an appliance Storage Node](#)
- [Monitoring the storage appliance installation](#)
- [Automating appliance installation and configuration](#)
- [Overview of installation REST APIs](#)
- [Troubleshooting the hardware installation](#)
- [Maintaining the SG6000 appliance](#)

### SG6000 appliances overview

The StorageGRIDSG6000 appliances are integrated storage and computing platforms that operate as Storage Nodes in a StorageGRID system. These appliances can be used in a hybrid grid environment that combines appliance Storage Nodes and virtual (software-based) Storage Nodes.

The SG6000 appliances provide the following features:

- Available in two models:
  - SG6060, which includes 60 drives and supports expansion shelves.
  - SGF6024, which offers 24 solid state drives (SSDs).
- Integrate the storage and computing elements for a StorageGRID Storage Node.
- Include the StorageGRID Appliance Installer to simplify Storage Node deployment and configuration.
- Include SANtricity System Manager for managing and monitoring the storage controllers and drives.
- Include a baseboard management controller (BMC) for monitoring and diagnosing the hardware in the compute controller.
- Support up to four 10-GbE or 25-GbE connections to the StorageGRID Grid Network and Client Network.
- Support Federal Information Processing Standard (FIPS) drives. When these drives are used with the Drive Security feature in SANtricity System Manager, unauthorized access to data is prevented.

### SG6060 overview

The StorageGRIDSG6060 appliance includes a compute controller and a storage controller shelf that contains two storage controllers and 60 drives. Optionally, 60-drive

expansion shelves can be added to the appliance.


### SG6060 components

The SG6060 appliance includes the following components:

Component	Description
Compute controller	<p>SG6000-CN controller, a one-rack unit (1U) server that includes:</p> <ul style="list-style-type: none"> <li>• 40 cores (80 threads)</li> <li>• 192 GB RAM</li> <li>• Up to 4 × 25 Gbps aggregate Ethernet bandwidth</li> <li>• 4 × 16 Gbps Fibre Channel (FC) interconnect</li> <li>• Baseboard management controller (BMC) that simplifies hardware management</li> <li>• Redundant power supplies</li> </ul>
Storage controller shelf	<p>E-Series E2860 controller shelf (storage array), a 4U shelf that includes:</p> <ul style="list-style-type: none"> <li>• Two E-Series E2800 controllers (duplex configuration) to provide storage controller failover support</li> <li>• Five-drawer drive shelf that holds sixty 3.5-inch drives (2 solid-state drives, or SSDs, and 58 NL-SAS drives)</li> <li>• Redundant power supplies and fans</li> </ul>
<p>Optional: Storage expansion shelves</p> <p><b>Note:</b> Expansion shelves can be installed during initial deployment or added later.</p>	<p>E-Series DE460C enclosure, a 4U shelf that includes:</p> <ul style="list-style-type: none"> <li>• Two input/output modules (IOMs)</li> <li>• Five drawers, each holding 12 NL-SAS drives, for a total of 60 drives</li> <li>• Redundant power supplies and fans</li> </ul> <p>Each SG6060 appliance can have one or two expansion shelves for a total of 180 drives.</p>


### SG6060 diagrams

This figure shows the front of the SG6060, which includes a 1U compute controller and a 4U shelf containing two storage controllers and 60 drives in five drive drawers.




	Description
1	SG6000-CN compute controller with front bezel
2	E2860 controller shelf with front bezel (optional expansion shelf looks identical)
3	SG6000-CN compute controller with front bezel removed
4	E2860 controller shelf with front bezel removed (optional expansion shelf looks identical)

This figure shows the back of the SG6060, including the compute and storage controllers, fans, and power supplies.



	Description
1	Power supply (1 of 2) for SG6000-CN compute controller
2	Connectors for SG6000-CN compute controller
3	Fan (1 of 2) for E2860 controller shelf
4	E-Series E2800 storage controller (1 of 2) and connectors
5	Power supply (1 of 2) for E2860 controller shelf

This figure shows the back of the optional expansion shelf for the SG6060, including the input/output modules (IOMs), fans, and power supplies. Each SG6060 can be installed with one or two expansion shelves, which can be included in the initial installation or added later.



	Description
1	Fan (1 of 2) for expansion shelf
2	IOM (1 of 2) for expansion shelf
3	Power supply (1 of 2) for expansion shelf

### SGF6024 overview

The StorageGRIDSGF6024 includes a compute controller and a storage controller shelf that holds 24 solid state drives.

### SGF6024 components


The SGF6024 appliance includes the following components:

Component	Description
Compute controller	<p>SG6000-CN controller, a one-rack unit (1U) server that includes:</p> <ul style="list-style-type: none"> <li>• 40 cores (80 threads)</li> <li>• 192 GB RAM</li> <li>• Up to 4 × 25 Gbps aggregate Ethernet bandwidth</li> <li>• 4 × 16 Gbps Fibre Channel (FC) interconnect</li> <li>• Baseboard management controller (BMC) that simplifies hardware management</li> <li>• Redundant power supplies</li> </ul>

Component	Description
Flash array (controller shelf)	<p>E-Series EF570 flash array (also known as a controller shelf), a 2U shelf that includes:</p> <ul style="list-style-type: none"> <li>• Two E-Series EF570 controllers (duplex configuration) to provide storage controller failover support</li> <li>• 24 solid state drives (also known as SSDs or flash drives)</li> <li>• Redundant power supplies and fans</li> </ul>


### SGF6024 diagrams

This figure shows the front of the SGF6024, which includes a 1U compute controller and a 2U enclosure containing two storage controllers and 24 flash drives.



	Description
1	SG6000-CN compute controller with front bezel
2	EF570 flash array with front bezel
3	SG6000-CN compute controller with front bezel removed
4	EF570 flash array with front bezel removed

This figure shows the back of the SGF6024, including the compute and storage controllers, fans, and power supplies.



	Description
1	Power supply (1 of 2) for SG6000-CN compute controller
2	Connectors for SG6000-CN compute controller
3	Power supply (1 of 2) for EF570 flash array
4	E-Series EF570 storage controller (1 of 2) and connectors

### Controllers in the SG6000 appliances

Each model of the StorageGRIDSG6000 appliance includes an SG6000-CN compute controller in a 1U enclosure and duplex E-Series storage controllers in a 2U or 4U enclosure, depending on the model. Review the diagrams to learn more about each type of controller.

#### All appliances: SG6000-CN compute controller


- Provides compute resources for the appliance.
- Includes the StorageGRID Appliance Installer.



StorageGRID software is not preinstalled on the appliance. This software is retrieved from the Admin Node when you deploy the appliance.

- Can connect to all three StorageGRID networks, including the Grid Network, the Admin Network, and the Client Network.
- Connects to the E-Series storage controllers and operates as the initiator.

This figure shows the connectors on the back of the SG6000-CN.




	Port	Type	Use
1	Interconnect ports 1-4	16-Gb/s Fibre Channel (FC), with integrated optics	Connect the SG6000-CN controller to the E2800 controllers (two connections to each E2800).
2	Network ports 1-4	10-GbE or 25-GbE, based on cable or SFP transceiver type, switch speed, and configured link speed	Connect to the Grid Network and the Client Network for StorageGRID.
3	BMC management port	1-GbE (RJ-45)	Connect to the SG6000-CN baseboard management controller.
4	Diagnostic and support ports	<ul style="list-style-type: none"> <li>• VGA</li> <li>• Serial, 115200 8-N-1</li> <li>• USB</li> </ul>	Reserved for technical support use.
5	Admin Network port 1	1-GbE (RJ-45)	Connect the SG6000-CN to the Admin Network for StorageGRID.

	Port	Type	Use
6	Admin Network port 2	1-GbE (RJ-45)	<p>Options:</p> <ul style="list-style-type: none"> <li>• Bond with management port 1 for a redundant connection to the Admin Network for StorageGRID.</li> <li>• Leave unwired and available for temporary local access (IP 169.254.0.1).</li> <li>• During installation, use port 2 for IP configuration if DHCP-assigned IP addresses are not available.</li> </ul>

#### SG6060: E2800 storage controllers

- Two controllers for failover support.
- Manage the storage of data on the drives.
- Function as standard E-Series controllers in a duplex configuration.
- Include SANtricity OS Software (controller firmware).
- Include SANtricity System Manager for monitoring storage hardware and for managing alerts, the AutoSupport feature, and the Drive Security feature.
- Connect to the SG6000-CN controller and provide access to the storage.

This figure shows the connectors on the back of each of the E2800 controllers.




	Port	Type	Use
1	Interconnect ports 1 and 2	16-Gb/s FC optical SFPa	Connect each of the E2800 controllers to the SG6000-CN controller.  There are four connections to the SG6000-CN controller (two from each E2800).
2	Management ports 1 and 2	1-Gb (RJ-45) Ethernet	<ul style="list-style-type: none"> <li>• Port 1 connects to the network where you access SANtricity System Manager on a browser.</li> <li>• Port 2 is reserved for technical support use.</li> </ul>
3	Diagnostic and support ports	<ul style="list-style-type: none"> <li>• RJ-45 serial port</li> <li>• Micro USB serial port</li> <li>• USB port</li> </ul>	Reserved for technical support use.
4	Drive expansion ports 1 and 2	12Gb/s SAS	Connect the ports to the drive expansion ports on the IOMs in the expansion shelf.

#### SGF6024: EF570 storage controllers

- Two controllers for failover support.
- Manage the storage of data on the drives.
- Function as standard E-Series controllers in a duplex configuration.
- Include SANtricity OS Software (controller firmware).
- Include SANtricity System Manager for monitoring storage hardware and for managing alerts, the AutoSupport feature, and the Drive Security feature.
- Connect to the SG6000-CN controller and provide access to the flash storage.


This figure shows the connectors on the back of each of the EF570 controllers.



	Port	Type	Use
1	Interconnect ports 1 and 2	16-Gb/s FC optical SFPa	Connect each of the EF570 controllers to the SG6000-CN controller.  There are four connections to the SG6000-CN controller (two from each EF570).
2	Diagnostic and support ports	<ul style="list-style-type: none"> <li>• RJ-45 serial port</li> <li>• Micro USB serial port</li> <li>• USB port</li> </ul>	Reserved for technical support use.
3	Drive expansion ports	12Gb/s SAS	Not used. The SGF6024 appliance does not support expansion drive shelves.
4	Management ports 1 and 2	1-Gb (RJ-45) Ethernet	<ul style="list-style-type: none"> <li>• Port 1 connects to the network where you access SANtricity System Manager on a browser.</li> <li>• Port 2 is reserved for technical support use.</li> </ul>

**SG6060: Input/output modules for optional expansion shelves**

The expansion shelf contains two input/output modules (IOMs) that connect to the storage controllers or to other expansion shelves.



	Port	Type	Use
1	Drive expansion ports 1-4	12Gb/s SAS	Connect each port to the storage controllers or additional expansion shelf (if any).

## Installation and deployment overview

You can install one or more StorageGRID storage appliances when you first deploy StorageGRID, or you can add appliance Storage Nodes later as part of an expansion. You might also need to install an appliance Storage Node as part of a recovery operation.

### What you'll need

Your StorageGRID system is using the required version of StorageGRID software.

Appliance	Required StorageGRID version
SG6060 with no expansion shelves	11.1.1 or later
SG6060 with expansion shelves (one or two)	11.3 or later <b>Note:</b> If you add expansion shelves after the initial deployment, you must use version 11.4 or later.
SGF6024	11.3 or later

### Installation and deployment tasks

Adding a StorageGRID storage appliance to a StorageGRID system includes four primary steps:

1. Preparing for installation:
  - Preparing the installation site
  - Unpacking the boxes and checking the contents
  - Obtaining additional equipment and tools
  - Gathering IP addresses and network information
  - Optional: Configuring an external key management server (KMS) if you plan to encrypt all appliance data. See details about external key management in the instructions for administering StorageGRID.
2. Installing the hardware:
  - Registering the hardware
  - Installing the appliance into a cabinet or rack
  - Installing the drives
  - Installing optional expansion shelves (model SG6060 only; maximum of two expansion shelves)
  - Cabling the appliance
  - Connecting the power cords and applying power
  - Viewing boot-up status codes
3. Configuring the hardware:
  - Accessing SANtricity System Manager to configure SANtricity System Manager settings
  - Accessing StorageGRID Appliance Installer, setting a static IP address for management port 1 on the storage controller, and configuring the link and network IP settings required to connect to StorageGRID networks

- Accessing the baseboard management controller (BMC) interface on the SG6000-CN controller
- Optional: Enabling node encryption if you plan to use an external KMS to encrypt appliance data.
- Optional: Changing the RAID mode.

#### 4. Deploying the appliance as a Storage Node:

Task	Instructions
Deploying an appliance Storage Node in a new StorageGRID system	<a href="#">Deploying an appliance Storage Node</a>
Adding an appliance Storage Node to an existing StorageGRID system	Instructions for expanding a StorageGRID system
Deploying an appliance Storage Node as part of a Storage Node recovery operation	Instructions for recovery and maintenance

#### Related information

[Preparing for installation](#)

[Installing the hardware](#)

[Configuring the hardware](#)

[Expand your grid](#)

[Maintain & recover](#)

[Administer StorageGRID](#)

### Preparing for installation

Preparing to install a StorageGRID appliance entails preparing the site and obtaining all required hardware, cables, and tools. You should also gather IP addresses and network information.

#### Steps

- [Preparing the site \(SG6000\)](#)
- [Unpacking the boxes \(SG6000\)](#)
- [Obtaining additional equipment and tools \(SG6000\)](#)
- [Web browser requirements](#)
- [Reviewing appliance network connections](#)
- [Gathering installation information \(SG6000\)](#)

#### Preparing the site (SG6000)

Before installing the appliance, you must make sure that the site and the cabinet or rack you plan to use meet the specifications for a StorageGRID appliance.

## Steps

1. Confirm that the site meets the requirements for temperature, humidity, altitude range, airflow, heat dissipation, wiring, power, and grounding. See the NetApp Hardware Universe for more information.
2. Confirm that your location provides 240-volt AC power for the SG6060 or 120-volt AC power for the SGF6024.
3. Obtain a 19-inch (48.3-cm) cabinet or rack to fit shelves of this size (without cables):

Type of shelf	Height	Width	Depth	Maximum weight
<b>E2860 controller shelf</b> for SG6060	6.87 in. (17.46 cm)	17.66 in. (44.86 cm)	38.25 in. (97.16 cm)	250 lb. (113 kg)
<b>Optional expansion shelf</b> for SG6060 (one or two)	6.87 in. (17.46 cm)	17.66 in. (44.86 cm)	38.25 in. (97.16 cm)	250 lb. (113 kg)
<b>EF570 controller shelf</b> for SGF6024	3.35 in. (8.50 cm)	17.66 in. (44.86 cm)	19.00 in. (48.26 cm)	51.74 lb. (23.47 kg)
<b>SG6000-CN controller</b> for each appliance	1.70 in. (4.32 cm)	17.32 in. (44.0 cm)	32.0 in. (81.3 cm)	39 lb. (17.7 kg)

4. Decide where you are going to install the appliance.



When installing the E2860 controller shelf or optional expansion shelves, install hardware from the bottom to the top of the rack or cabinet to prevent the equipment from tipping over. To ensure that the heaviest equipment is at the bottom of the cabinet or rack, install the SG6000-CN controller above the E2860 controller shelf and expansion shelves.



Before committing to the installation, verify that the 0.5m optic cables shipped with the appliance, or cables that you supply, are long enough for the planned layout.

## Related information

[NetApp Hardware Universe](#)

[NetApp Interoperability Matrix Tool](#)

## Unpacking the boxes (SG6000)


Before installing the StorageGRID appliance, unpack all boxes and compare the contents to the items on the packing slip.

### SG6060


- **SG6000-CN controller**



- E2860 controller shelf with no drives installed




- Two front bezels




- Two rail kits with instructions




- 60 drives (2 SSD and 58 NL-SAS)



- Four handles




- Back brackets and cage nuts for square-hole rack installation



#### SG6060 expansion shelf


- Expansion shelf with no drives installed



- Front bezel




- 60 NL-SAS drives




- One rail kit with instructions



- Four handles



- Back brackets and cage nuts for square-hole rack installation



SGF6024

- SG6000-CN controller



- EF570 flash array with 24 solid state (flash) drives installed



- Two front bezels



- Two rail kits with instructions



- Shelf endcaps



### Cables and connectors


The shipment for the StorageGRID appliance includes the following cables and connectors:

- **Four power cords for your country**



Your cabinet might have special power cords that you use instead of the power cords that ship with the appliance.

- **Optical cables and SFP transceivers**



Four optical cables for the FC interconnect ports

Four SFP+ transceivers, which support 16-Gb/s FC

- **Optional: Two SAS cables for connecting each SG6060 expansion shelf**



### Obtaining additional equipment and tools (SG6000)

Before installing the StorageGRID appliance, confirm you have all of the additional equipment and tools that you need.

You need the following additional equipment to install and configure the hardware:

- **Screwdrivers**




Phillips No. 2 screwdriver

Medium flat-blade screwdriver

- **ESD wrist strap**



- **Optical cables and SFP transceivers**



You need one of the following options:

- One to four TwinAx cables or optical cables for the 10/25-GbE ports you plan to use on the SG6000-CN controller
- One to four SFP+ transceivers for the 10/25-GbE ports if you will use optical cables and 10-GbE link speed
- One to four SFP28 transceivers for the 10/25-GbE ports if you will use optical cables and 25-GbE link speed

- **RJ-45 (Cat5/Cat5e/Cat6) Ethernet cables**



- **Service laptop**



Supported web browser

1-GbE (RJ-45) port

- **Optional tools**



Power drill with Phillips head bit

Flashlight

Mechanized lift for 60-drive shelves

### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

### Reviewing appliance network connections

Before installing the StorageGRID appliance, you should understand which networks can be connected to the appliance.

When you deploy a StorageGRID appliance as a Storage Node in a StorageGRID system, you can connect it to the following networks:

- **Grid Network for StorageGRID:** The Grid Network is used for all internal StorageGRID traffic. It provides connectivity between all nodes in the grid, across all sites and subnets. The Grid Network is required.
- **Admin Network for StorageGRID:** The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites. The Admin Network is optional.
- **Client Network for StorageGRID:** The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network provides client protocol access to the grid, so the

Grid Network can be isolated and secured. The Client Network is optional.

- **Management network for SANtricity System Manager:** This network provides access to SANtricity System Manager on the storage controller, allowing you to monitor and manage the hardware components in the storage controller shelf. This management network can be the same as the Admin Network for StorageGRID, or it can be an independent management network.
- **BMC management network for the SG6000-CN controller:** This network provides access to the baseboard management controller in the SG6000-CN, allowing you to monitor and manage the hardware components in the SG6000-CN controller. This management network can be the same as the Admin Network for StorageGRID, or it can be an independent management network.



For detailed information about StorageGRID networks, see the *Grid Primer*.

### Related information

[Gathering installation information \(SG6000\)](#)

[Cabling the appliance \(SG6000\)](#)

[Port bond modes for the SG6000-CN controller](#)

[Network guidelines](#)

### Port bond modes for the SG6000-CN controller

When configuring network links for the SG6000-CN, you can use port bonding for the 10/25-GbE ports that connect to the Grid Network and optional Client Network, and the 1-GbE management ports that connect to the optional Admin Network. Port bonding helps protect your data by providing redundant paths between StorageGRID networks and the appliance.

### Related information


[Configuring network links \(SG6000\)](#)

### Network bond modes for the 10/25-GbE ports

The 10/25-GbE networking ports on the SG6000-CN controller support Fixed port bond mode or Aggregate port bond mode for the Grid Network and Client Network connections.

### Fixed port bond mode

Fixed mode is the default configuration for the 10/25-GbE networking ports.



	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

When using Fixed port bond mode, the ports can be bonded using active-backup mode or Link Aggregation Control Protocol mode (LACP 802.3ad).

- In active-backup mode (default), only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Port 4 provides a backup path for port 2 (Grid Network), and port 3 provides a backup path for port 1 (Client Network).
- In LACP mode, each pair of ports forms a logical channel between the controller and the network, allowing for higher throughput. If one port fails, the other port continues to provide the channel. Throughput is reduced, but connectivity is not impacted.




If you do not need redundant connections, you can use only one port for each network. However, be aware that an alert will be triggered in the Grid Manager after StorageGRID is installed, indicating that the link is down. Because this port is disconnected on purpose, you can safely disable this alert.

From the Grid Manager, select **Alert > Rules**, select the rule, and click **Edit rule**. Then, uncheck the **Enabled** check box.

### Aggregate port bond mode

Aggregate port bond mode significantly increases the throughput for each StorageGRID network and provides additional failover paths.



	Which ports are bonded
1	All connected ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

If you plan to use aggregate port bond mode:

- You must use LACP network bond mode.
- You must specify a unique VLAN tag for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.
- The ports must be connected to switches that can support VLAN and LACP. If multiple switches are participating in the LACP bond, the switches must support multi-chassis link aggregation groups (MLAG), or equivalent.

- You must understand how to configure the switches to use VLAN, LACP, and MLAG, or equivalent.

If you do not want to use all four 10/25-GbE ports, you can use one, two, or three ports. Using more than one port maximizes the chance that some network connectivity will remain available if one of the 10/25-GbE ports fails.



If you choose to use fewer than four ports, be aware that one or more alarms will be raised in the Grid Manager after StorageGRID is installed, indicating that cables are unplugged. You can safely acknowledge the alarms to clear them.

### Network bond modes for the 1-GbE management ports

For the two 1-GbE management ports on the SG6000-CN controller, you can choose Independent network bond mode or Active-Backup network bond mode to connect to the optional Admin Network.

In Independent mode, only the management port on the left is connected to the Admin Network. This mode does not provide a redundant path. The management port on the right is unconnected and available for temporary local connections (uses IP address 169.254.0.1)

In Active-Backup mode, both management ports are connected to the Admin Network. Only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Bonding these two physical ports into one logical management port provides a redundant path to the Admin Network.



If you need to make a temporary local connection to the SG6000-CN controller when the 1-GbE management ports are configured for Active-Backup mode, remove the cables from both management ports, plug your temporary cable into the management port on the right, and access the appliance using IP address 169.254.0.1.



	Network bond mode
A	Both management ports are bonded into one logical management port connected to the Admin Network.
I	The port on the left is connected to the Admin Network. The port on the right is available for temporary local connections (IP address 169.254.0.1).

### Gathering installation information (SG6000)

As you install and configure the StorageGRID appliance, you must make decisions and gather information about Ethernet switch ports, IP addresses, and port and network bond modes.

#### About this task

You can use the following tables to record the required information for each network you connect to the appliance. These values are required to install and configure the hardware.

**Information needed to connect to SANtricity System Manager on the storage controllers**

You must connect both of the storage controllers in the appliance (either the E2800 controllers or the EF570 controllers) to the management network you will use for SANtricity System Manager. The controllers are located in each appliance as follows:

- SG6060: Controller A is on the top, and controller B is on the bottom.
- SGF6024: Controller A is on the left, and controller B is on the right.

Information needed	Your value for controller A	Your value for controller B
Ethernet switch port you will connect to management port 1 (labeled as P1 on the controller)		
MAC address for management port 1 (printed on a label near port P1)		
DHCP-assigned IP address for management port 1, if available after power on  <b>Note:</b> If the network you will connect to the storage controller includes a DHCP server, the network administrator can use the MAC address to determine the IP address that was assigned by the DHCP server.		
Static IP address you plan to use for the appliance on the management network	For IPv4: <ul style="list-style-type: none"> <li>• IPv4 address:</li> <li>• Subnet mask:</li> <li>• Gateway:</li> </ul> For IPv6: <ul style="list-style-type: none"> <li>• IPv6 address:</li> <li>• Routable IP address:</li> <li>• storage controller router IP address:</li> </ul>	For IPv4: <ul style="list-style-type: none"> <li>• IPv4 address:</li> <li>• Subnet mask:</li> <li>• Gateway:</li> </ul> For IPv6: <ul style="list-style-type: none"> <li>• IPv6 address:</li> <li>• Routable IP address:</li> <li>• storage controller router IP address:</li> </ul>
IP address format	Choose one: <ul style="list-style-type: none"> <li>• IPv4</li> <li>• IPv6</li> </ul>	Choose one: <ul style="list-style-type: none"> <li>• IPv4</li> <li>• IPv6</li> </ul>

Information needed	Your value for controller A	Your value for controller B
Speed and duplex mode  <b>Note:</b> You must make sure the Ethernet switch for the SANtricity System Manager management network is set to autonegotiate.	Must be: <ul style="list-style-type: none"> <li>Autonegotiate (default)</li> </ul>	Must be: <ul style="list-style-type: none"> <li>Autonegotiate (default)</li> </ul>

**Information needed to connect the SG6000-CN controller to the Admin Network**

The Admin Network for StorageGRID is an optional network, used for system administration and maintenance. The appliance connects to the Admin Network using the following 1-GbE management ports on the SG6000-CN controller.



Information needed	Your value
Admin Network enabled	Choose one: <ul style="list-style-type: none"> <li>No</li> <li>Yes (default)</li> </ul>
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>Independent (default)</li> <li>Active-Backup</li> </ul>
Switch port for the left port in the red circle in the diagram (default active port for Independent network bond mode)	
Switch port for the right port in the red circle in the diagram (Active-Backup network bond mode only)	

Information needed	Your value
<p>MAC address for the Admin Network port</p> <p><b>Note:</b> The MAC address label on the front of the SG6000-CN controller lists the MAC address for the BMC management port. To determine the MAC address for the Admin Network port, you must add <b>2</b> to the hexadecimal number on the label. For example, if the MAC address on the label ends in <b>09</b>, the MAC address for the Admin Port would end in <b>0B</b>. If the MAC address on the label ends in <b>(y)FF</b>, the MAC address for the Admin Port would end in <b>(y+1)01</b>. You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing <b>+ 2 =</b>.</p>	
<p>DHCP-assigned IP address for the Admin Network port, if available after power on</p> <p><b>Note:</b> You can determine the DHCP-assigned IP address by using the MAC address to look up the assigned IP.</p>	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
<p>Static IP address you plan to use for the appliance Storage Node on the Admin Network</p> <p><b>Note:</b> If your network does not have a gateway, specify the same static IPv4 address for the gateway.</p>	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Admin Network subnets (CIDR)	

**Information needed to connect and configure the 10/25-GbE ports on the SG6000-CN controller**

The four 10/25-GbE ports on the SG6000-CN controller connect to the StorageGRID Grid Network and the optional Client Network.

Information needed	Your value
Link speed	<p>Choose one:</p> <ul style="list-style-type: none"> <li>• Auto (default)</li> <li>• 10 GbE</li> <li>• 25 GbE</li> </ul>
Port bond mode	<p>Choose one:</p> <ul style="list-style-type: none"> <li>• Fixed (default)</li> <li>• Aggregate</li> </ul>

Information needed	Your value
Switch port for port 1 (Client Network for Fixed mode)	
Switch port for port 2 (Grid Network for Fixed mode)	
Switch port for port 3 (Client Network for Fixed mode)	
Switch port for port 4 (Grid Network for Fixed mode)	

#### Information needed to connect the SG6000-CN controller to the Grid Network

The Grid Network for StorageGRID is a required network, used for all internal StorageGRID traffic. The appliance connects to the Grid Network using the 10/25-GbE ports on the SG6000-CN controller.

Information needed	Your value
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>• Active-Backup (default)</li> <li>• LACP (802.3ad)</li> </ul>
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>
VLAN tag(if VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Grid Network, if available after power on	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Static IP address you plan to use for the appliance Storage Node on the Grid Network  <b>Note:</b> If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Grid Network subnets (CIDRs)	

#### Information needed to connect the SG6000-CN controller to the Client Network

The Client Network for StorageGRID is an optional network, typically used to provide client protocol access to the grid. The appliance connects to the Client Network using the 10/25-GbE ports on the SG6000-CN controller.

Information needed	Your value
Client Network enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>• Active-Backup (default)</li> <li>• LACP (802.3ad)</li> </ul>
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>
VLAN tag(If VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Client Network, if available after power on	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Static IP address you plan to use for the appliance Storage Node on the Client Network  <b>Note:</b> If the Client Network is enabled, the default route on the controller will use the gateway specified here.	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>

**Information needed to connect the SG6000-CN controller to the BMC management network**

You can access the BMC interface on the SG6000-CN controller using the following 1-GbE management port. This port supports remote management of the controller hardware over Ethernet using the Intelligent Platform Management Interface (IPMI) standard.



Information needed	Your value
Ethernet switch port you will connect to the BMC management port (circled in the diagram)	
DHCP-assigned IP address for the BMC management network, if available after power on	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>

Information needed	Your value
Static IP address you plan to use for the BMC management port	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>

### Related information

[Controllers in the SG6000 appliances](#)

[Reviewing appliance network connections](#)

[Port bond modes for the SG6000-CN controller](#)

[Cabling the appliance \(SG6000\)](#)

[Configuring StorageGRID IP addresses](#)

## Installing the hardware

Hardware installation entails installing the SG6000-CN controller and the storage controller shelf into a cabinet or rack, connecting the cables, and applying power.

### Steps

- [Registering the hardware](#)
- [SG6060: Installing 60-drive shelves into a cabinet or rack](#)
- [SG6060: Installing the drives](#)
- [SGF6024: Installing 24-drive shelves into a cabinet or rack](#)
- [SG6000-CN: Installing into a cabinet or rack](#)
- [Cabling the appliance \(SG6000\)](#)
- [SG6060: Cabling the optional expansion shelves](#)
- [Connecting power cords and applying power \(SG6000\)](#)
- [Viewing status indicators and buttons on the SG6000-CN controller](#)
- [Viewing boot-up status codes for the SG6000 storage controllers](#)

### Registering the hardware

Registering the appliance hardware provides support benefits.

### Steps

1. Locate the chassis serial number for the storage controller shelf.

You can find the number on the packing slip, in your confirmation email, or on the appliance after you unpack it.





There are several serial numbers on the storage appliance. The serial number on the storage controller shelf is the one that must be registered and used if you call for service or support on the appliance.

2. Go to the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com).
3. Determine whether you need to register the hardware:

If you are a...	Follow these steps...
Existing NetApp customer	<ol style="list-style-type: none"> <li>a. Sign in with your username and password.</li> <li>b. Select <b>Products &gt; My Products</b>.</li> <li>c. Confirm that the new serial number is listed.</li> <li>d. If it is not, follow the instructions for new NetApp customers.</li> </ol>
New NetApp customer	<ol style="list-style-type: none"> <li>a. Click <b>Register Now</b>, and create an account.</li> <li>b. Select <b>Products &gt; Register Products</b>.</li> <li>c. Enter the product serial number and requested details.</li> </ol> <p>After your registration is approved, you can download any required software. The approval process might take up to 24 hours.</p>

### SG6060: Installing 60-drive shelves into a cabinet or rack

You must install a set of rails for the E2860 controller shelf in your cabinet or rack, and then slide the controller shelf onto the rails. If you are installing 60-drive expansion shelves, the same procedure applies.

#### What you'll need

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.



Each 60-drive shelf weighs approximately 132 lb (60 kg) without drives installed. Four people or a mechanized lift are required to safely move the shelf.



To avoid damaging the hardware, never move the shelf if drives are installed. You must remove all drives before moving the shelf.



When installing the E2860 controller shelf or optional expansion shelves, install hardware from the bottom to the top of the rack or cabinet to prevent the equipment from tipping over. To ensure that the heaviest equipment is at the bottom of the cabinet or rack, install the SG6000-CN controller above the E2860 controller shelf and expansion shelves.



Before committing to the installation, verify that the 0.5m optic cables shipped with the appliance, or cables that you supply, are long enough for the planned layout.


## Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.

For square hole cabinets, you must first install the provided cage nuts to secure the front and rear of the shelf with screws.

2. Remove the outer packing box for the appliance. Then, fold down the flaps on the inner box.
3. If you are lifting the appliance by hand, attach the four handles to the sides of the chassis.

Push up on each handle until it clicks into place.



4. Place the back of the shelf (the end with the connectors) on the rails.
5. Supporting the shelf from the bottom, slide it into the cabinet. If you are using the handles, use the thumb latches to detach one handle at a time as you slide the shelf in.


To remove the handles, pull back on the release latch, push down, then pull away from the shelf.

6. Secure the shelf to the front of the cabinet.

Insert screws into the first and third holes from the top of the shelf on both sides.

7. Secure the shelf to the rear of the cabinet.

Place two back brackets on each side of the upper rear section of the shelf. Insert screws into the first and third holes of each bracket.



8. Repeat these steps for any expansion shelves.

### SG6060: Installing the drives

After installing the 60-drive shelf into a cabinet or rack, you must install all 60 drives into the shelf. The shipment for the E2860 controller shelf includes two SSD drives, which you should install in the top drawer of the controller shelf. Each optional expansion shelf includes 60 HDD drives and no SSD drives.

#### What you'll need

You have installed the E2860 controller shelf or optional expansion shelves (one or two) in the cabinet or rack.



To avoid damaging the hardware, never move the shelf if drives are installed. You must remove all drives before moving the shelf.


#### Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Remove the drives from their packaging.
3. Release the levers on the top drive drawer, and slide the drawer out using the levers.
4. Locate the two SSD drives.



Expansion shelves do not use SSD drives.

5. Raise each drive handle to a vertical position.
6. Install the two SSD drives in slots 0 and 1 (the first two slots along the lefthand side of the drawer).
7. Gently position each drive into its slot, and lower the raised drive handle until it clicks into place.



8. Install 10 HDD drives into the top drawer.
9. Slide the drawer back in by pushing on the center and closing both levers gently.



Stop pushing the drawer if you feel binding. Use the release levers at the front of the drawer to slide the drawer back out. Then, carefully reinsert the drawer into the slot.

10. Repeat these steps to install HDD drives into the other four drawers.



You must install all 60 drives to ensure correct operation.

11. Attach the front bezel to the shelf.
12. If you have expansion shelves, repeat these steps to install 12 HDD drives into each drawer of each expansion shelf.
13. Proceed to the instructions for installing the SG6000-CN into a cabinet or rack.

#### **SGF6024: Installing 24-drive shelves into a cabinet or rack**

You must install a set of rails for the EF570 controller shelf in your cabinet or rack, and then slide the array onto the rails.

## What you'll need

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.

## Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.

For square hole cabinets, you must first install the provided cage nuts to secure the front and rear of the shelf with screws.

2. Remove the outer packing box for the appliance. Then, fold down the flaps on the inner box.
3. Place the back of the shelf (the end with the connectors) on the rails.



A fully loaded shelf weighs approximately 52 lb (24 kg). Two persons are required to safely move the enclosure.

4. Carefully slide the enclosure all the way onto the rails.



You might need to adjust the rails to ensure that the enclosure slides all the way onto the rails.




Do not place additional equipment on the rails after you finish installing the enclosure. The rails are not designed to bear additional weight.




If applicable, you might need to remove the shelf end caps or the system bezel to secure the enclosure to the rack post; if so, you need to replace the end caps or bezel when you are done.

5. Secure the enclosure to the front of the cabinet or rack and rails by inserting two M5 screws through the mounting brackets (preinstalled on either side of the front of the enclosure), the holes on the rack or system cabinet, and the holes on the front of rails.



6. Secure the enclosure to the back of the rails by inserting two M5 screws through the brackets at the enclosure and the rail kit bracket.
7. If applicable, replace the shelf end caps or the system bezel.



### SG6000-CN: Installing into a cabinet or rack

You must install a set of rails for the SG6000-CN controller in your cabinet or rack, and then slide the controller onto the rails.


#### What you'll need

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.

- You have the instructions packaged with the rail kit.
- You have installed the E2860 controller shelf and drives or the EF570 controller shelf.


### Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.
2. On the two rails installed in the cabinet or rack, extend the movable parts of the rails until you hear a click.



3. Insert the SG6000-CN controller into the rails.
4. Slide the controller into the cabinet or rack.

When you cannot move the controller any further, pull the blue latches on both sides of the chassis to slide the controller all the way in.



Do not attach the front bezel until after you power on the controller.

5. Tighten the captive screws on the controller front panel to secure the controller in the rack.



### Cabling the appliance (SG6000)

You must connect the storage controllers to the SG6000-CN controller, connect the management ports on all three controllers, and connect the network ports on the SG6000-CN controller to the Grid Network and optional Client Network for StorageGRID.

#### What you'll need

- You have the four optical cables provided with the appliance for connecting the two storage controllers to the SG6000-CN controller.
- You have RJ-45 Ethernet cables (four minimum) for connecting the management ports.
- You have one of the following options for the network ports. These items are not provided with the appliance.
  - One to four TwinAx cables for connecting the four network ports.


- One to four SFP+ or SFP28 transceivers if you plan to use optical cables for the ports.




**Risk of exposure to laser radiation** — Do not disassemble or remove any part of an SFP transceiver. You might be exposed to laser radiation.

**About this task**

The following figure shows the three controllers in the SG6060 appliance, with the SG6000-CN compute controller on the top and the two E2800 storage controllers on the bottom.



The following figure shows the three controllers in the SGF6024 appliance, with the SG6000-CN compute controller on the top and the two EF570 storage controllers side by side below the compute controller.



	Port	Type of port	Function
1	BMC management port on the SG6000-CN controller	1-GbE (RJ-45)	Connects to the network where you access the BMC interface.
2	FC connection ports: <ul style="list-style-type: none"> <li>• 4 on the SG6000-CN controller</li> <li>• 2 on each storage controller</li> </ul>	16-Gb/s FC optical SFP+	Connect each storage controller to the SG6000-CN controller.
3	Four network ports on the SG6000-CN controller	10/25-GbE	Connect to the Grid Network and the Client Network for StorageGRID.
4	Admin Network port on the SG6000-CN controller (labelled P1 in the figure)	1-GbE (RJ-45) <b>Important:</b> This port operates only at 1000 baseT/full and does not support 10- or 100-megabit speeds.	Connects the SG6000-CN controller to the Admin Network for StorageGRID.
4	Rightmost RJ-45 port on the SG6000-CN controller	1-GbE (RJ-45) <b>Important:</b> This port operates only at 1000 baseT/full and does not support 10- or 100-megabit speeds.	<ul style="list-style-type: none"> <li>• Can be bonded with management port 1 if you want a redundant connection to the Admin Network.</li> <li>• Can be left unwired and available for temporary local access (IP 169.254.0.1).</li> <li>• During installation, can be used to connect the SG6000-CN controller to a service laptop if DHCP-assigned IP addresses are not available.</li> </ul>
5	Management port 1 on each storage controller	1-GbE (RJ-45)	Connects to the network where you access SANtricity System Manager.

	Port	Type of port	Function
5	Management port 2 on each storage controller	1-GbE (RJ-45)	Reserved for technical support.

### Steps

1. Connect the BMC management port on the SG6000-CN controller to the management network, using an Ethernet cable.

Although this connection is optional, it is recommended to facilitate support.

2. Connect the two FC ports on each storage controller to the FC ports on the SG6000-CN controller, using four optical cables and four SFP+ transceivers for the storage controllers.
3. Connect the network ports on the SG6000-CN controller to the appropriate network switches, using TwinAx cables or optical cables and SFP+ or SFP28 transceivers.



The four network ports must use the same link speed. Install SFP+ transceivers if you plan to use 10-GbE link speeds. Install SFP28 transceivers if you plan to use 25-GbE link speeds.

- If you plan to use Fixed port bond mode (default), connect the ports to the StorageGRID Grid and Client Networks, as shown in the table.

Port	Connects to...
Port 1	Client Network (optional)
Port 2	Grid Network
Port 3	Client Network (optional)
Port 4	Grid Network

- If you plan to use the Aggregate port bond mode, connect one or more of the network ports to one or more switches. You should connect at least two of the four ports to avoid having a single point of failure. If you use more than one switch for a single LACP bond, the switches must support MLAG or equivalent.
4. If you plan to use the Admin Network for StorageGRID, connect the Admin Network port on the SG6000-CN controller to the Admin Network, using an Ethernet cable.
  5. Connect management port 1 (P1) on each storage controller (the RJ-45 port on the left) to the management network for SANtricity System Manager, using an Ethernet cable.

Do not use management port 2 (P2) on the storage controllers (the RJ-45 port on the right). This port is reserved for technical support.

### Related information

[Port bond modes for the SG6000-CN controller](#)

[Reinstalling the SG6000-CN controller into a cabinet or rack](#)

## **SG6060: Cabling the optional expansion shelves**

If you are using expansion shelves, you must connect them to the E2860 controller shelf. You can have a maximum of two expansion shelves for each SG6060 appliance.

### **What you'll need**


- You have the two SAS cables shipped with each expansion shelf.
- You have installed the expansion shelves in the cabinet or rack that contains the E2860 controller shelf.

[SG6060: Installing 60-drive shelves into a cabinet or rack](#)

### **Step**

Connect each expansion shelf to the E2860 controller shelf as shown in the diagram.

This drawing shows two expansion shelves. If you have only one, connect IOM A to controller A and connect IOM B to controller B.



	Description
1	SG6000-CN

	Description
2	E2860 controller shelf
3	Controller A
4	Controller B
5	Expansion shelf 1
6	IOM A for expansion shelf 1
7	IOM B for expansion shelf 1
8	Expansion shelf 2
9	IOM A for expansion shelf 2
10	IOM B for expansion shelf 2

### Connecting power cords and applying power (SG6000)

After connecting the network cables, you are ready to apply power to the SG6000-CN controller and to the two storage controllers or optional expansion shelves.

#### Steps

1. Confirm that both controllers in the storage controller shelf are off.



**Risk of electrical shock** — Before connecting the power cords, make sure that the power switches for each of the two storage controllers are off.

2. If you have expansion shelves, confirm that both of the IOM power switches are off.



**Risk of electrical shock** — Before connecting the power cords, make sure that the two power switches for each of the expansion shelves are off.

3. Connect a power cord to each of the two power supply units in the SG6000-CN controller.
4. Connect these two power cords to two different power distribution units (PDUs) in the cabinet or rack.
5. Connect a power cord to each of the two power supply units in the storage controller shelf.
6. If you have expansion shelves, connect a power cord to each of the two power supply units in each expansion shelf.
7. Connect the two power cords in each storage shelf (including the optional expansion shelves) to two different PDUs in the cabinet or rack.
8. If the power button on the front of the SG6000-CN controller is not currently illuminated blue, press the button to turn on power to the controller.

Do not press the power button again during the power-on process.

9. Turn on the two power switches on the back of the storage controller shelf. If you have expansion shelves, turn on the two power switches for each shelf.
  - Do not turn off the power switches during the power-on process.
  - The fans in the storage controller shelf and optional expansion shelves might be very loud when they first start up. The loud noise during start-up is normal.
10. After the components have booted up, check their status.
  - Check the seven-segment display on the back of each storage controller. Refer to the article about viewing boot-up status codes for more information.
  - Verify that the power button on the front of the SG6000-CN controller is lit.
11. If errors occur, correct any issues.
12. Attach the front bezel to the SG6000-CN controller.

### Related information

[Viewing boot-up status codes for the SG6000 storage controllers](#)

[Viewing status indicators and buttons on the SG6000-CN controller](#)

[Reinstalling the SG6000-CN controller into a cabinet or rack](#)

### Viewing status indicators and buttons on the SG6000-CN controller

The SG6000-CN controller includes indicators that help you determine the status of the controller, including the following indicators and buttons.



	Display	Description
1	Power button	<ul style="list-style-type: none"><li>• Blue: The controller is powered on.</li><li>• Off: The controller is powered off.</li></ul>
2	Reset button	<i>No indicator</i>  Use this button to perform a hard reset of the controller.

	Display	Description
3	Identify button	<ul style="list-style-type: none"> <li>• Blinking or solid blue: Identifies the controller in the cabinet or rack.</li> <li>• Off: The controller is not visually identifiable in the cabinet or rack.</li> </ul> <p>This button can be set to Blink, On (Solid), or Off.</p>
4	Alarm LED	<ul style="list-style-type: none"> <li>• Amber: An error has occurred.</li> </ul> <p><b>Note:</b> To view the boot-up and error codes, you must access the BMC interface.</p> <ul style="list-style-type: none"> <li>• Off: No errors are present.</li> </ul>

### General boot-up codes

During boot-up or after a hard reset of the SG6000-CN controller, the following occurs:

1. The baseboard management controller (BMC) logs codes for the boot-up sequence, including any errors that occur.
2. The power button lights up.
3. If any errors occur during boot-up, the alarm LED lights up.

To view the boot-up and error codes, you must access the BMC interface.

### Related information

[Troubleshooting the hardware installation](#)

[Configuring the BMC interface](#)

[Powering on the SG6000-CN controller and verifying operation](#)

### Viewing boot-up status codes for the SG6000 storage controllers

Each storage controller has a seven-segment display that provides status codes as the controller powers up. The status codes are the same for both the E2800 controller and the EF570 controller.

### About this task

For descriptions of these codes, see the E-Series system monitoring information for your storage controller type.

### Steps


1. During boot-up, monitor progress by viewing the codes shown on the seven-segment display for each

storage controller.

The seven-segment display on each storage controller shows the repeating sequence **OS**, **Sd**, **blank** to indicate that the controller is performing start-of-day processing.

2. After the controllers have booted up, confirm that each storage controller shows 99, which is the default ID for an E-Series controller shelf.

Make sure this value is displayed on both storage controllers, as shown in this example E2800 controller.



3. If one or both controllers show other values, see the information about troubleshooting the hardware installation, and confirm you completed the installation steps correctly. If you are unable to resolve the problem, contact technical support.

#### Related information

[E5700 and E2800 System Monitoring Guide](#)

[Troubleshooting the hardware installation](#)

[NetApp Support](#)

[Powering on the SG6000-CN controller and verifying operation](#)

## Configuring the hardware

After applying power to the appliance, you must configure the network connections that will be used by StorageGRID. You must configure SANtricity System Manager, which is the software you will use to monitor the storage controllers and other hardware in the controller shelf. You must also ensure that you can access the BMC interface for the SG6000-CN controller.

#### Steps

- [Configuring StorageGRID connections](#)
- [Accessing and Configuring SANtricity System Manager](#)
- [Configuring the BMC interface](#)
- [Optional: Enabling node encryption](#)
- [Optional: Changing the RAID mode \(SG6000 only\)](#)
- [Optional: Remapping network ports for the appliance](#)

## Configuring StorageGRID connections

Before you can deploy a StorageGRID appliance as a Storage Node in a StorageGRID system, you must configure the connections between the appliance and the networks you plan to use. You can configure networking by browsing to the StorageGRID Appliance Installer, which is pre-installed on the SG6000-CN controller (the compute controller).

### Steps

- [Accessing the StorageGRID Appliance Installer](#)
- [Verifying and upgrading the StorageGRID Appliance Installer version](#)
- [Configuring network links \(SG6000\)](#)
- [Configuring StorageGRID IP addresses](#)
- [Verifying network connections](#)
- [Verifying port-level network connections](#)

### Accessing the StorageGRID Appliance Installer

You must access the StorageGRID Appliance Installer to verify the installer version and configure the connections between the appliance and the three StorageGRID networks: the Grid Network, the Admin Network (optional), and the Client Network (optional).

### What you'll need

- You are using any management client that can connect to the StorageGRID Admin Network, or you have a service laptop.
- The client or service laptop has a supported web browser.
- The SG6000-CN controller is connected to all of the StorageGRID networks you plan to use.
- You know the IP address, gateway, and subnet for the SG6000-CN controller on these networks.
- You have configured the network switches you plan to use.

### About this task

To initially access the StorageGRID Appliance Installer, you can use the DHCP-assigned IP address for the Admin Network port on the SG6000-CN controller (assuming the controller is connected to the Admin Network), or you can connect a service laptop directly to the SG6000-CN controller.

### Steps

1. If possible, use the DHCP address for the Admin Network port on the SG6000-CN controller to access the StorageGRID Appliance Installer.



- a. Locate the MAC address label on the front of the SG6000-CN controller, and determine the MAC address for the Admin Network port.

The MAC address label lists the MAC address for the BMC management port.

To determine the MAC address for the Admin Network port, you must add **2** to the hexadecimal number

on the label. For example, if the MAC address on the label ends in **09**, the MAC address for the Admin Port would end in **0B**. If the MAC address on the label ends in **(y)FF**, the MAC address for the Admin Port would end in **(y+1)01**. You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing **+ 2 =**.

b. Provide the MAC address to your network administrator, so they can look up the DHCP address for the appliance on the Admin Network.

c. From the client, enter this URL for the StorageGRID Appliance Installer:

**`https://Appliance_Controller_IP:8443`**

For *SG6000-CN\_Controller\_IP*, use the DHCP address.

d. If you are prompted with a security alert, view and install the certificate using the browser's installation wizard.

The alert will not appear the next time you access this URL.

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected to StorageGRID networks. Error messages might appear that will be resolved in later steps.

## Home

**i** The installation is ready to be started. Review the settings below, and then click Start Installation.

### This Node

Node type

Storage ▾

Node name

MM-2-108-SGA-lab25

Cancel

Save

### Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

172.16.1.178

Connection state

Connection to 172.16.1.178 ready

Cancel

Save

### Installation

Current state

Ready to start installation of MM-2-108-SGA-lab25 into grid with Admin Node 172.16.1.178 running StorageGRID 11.2.0, using StorageGRID software downloaded from the Admin Node.

Start Installation

2. If you cannot obtain an IP address using DHCP, you can use a link-local connection.
  - a. Connect a service laptop directly to the rightmost RJ-45 port on the SG6000-CN controller, using an Ethernet cable.



- b. Open a web browser on the service laptop.
- c. Enter this URL for the StorageGRID Appliance Installer:

**https://169.254.0.1:8443**

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected.



If you cannot access the Home page over a link-local connection, configure the service laptop IP address as 169.254.0.2, and try again.

### After you finish

After accessing the StorageGRID Appliance Installer:

- Verify that the StorageGRID Appliance Installer version on the appliance matches the software version installed on your StorageGRID system. Upgrade StorageGRID Appliance Installer, if necessary.

#### [Verifying and upgrading the StorageGRID Appliance Installer version](#)

- Review any messages displayed on the StorageGRID Appliance Installer Home page and configure the link configuration and the IP configuration, as required.

### Related information

[Web browser requirements](#)

#### Verifying and upgrading the StorageGRID Appliance Installer version

The StorageGRID Appliance Installer version on the appliance must match the software version installed on your StorageGRID system to ensure that all StorageGRID features are supported.

### What you'll need

You have accessed the StorageGRID Appliance Installer.

### About this task

StorageGRID appliances come from the factory preinstalled with the StorageGRID Appliance Installer. If you are adding an appliance to a recently upgraded StorageGRID system, you might need to manually upgrade the StorageGRID Appliance Installer before installing the appliance as a new node.

The StorageGRID Appliance Installer automatically upgrades when you upgrade to a new StorageGRID version. You do not need to upgrade the StorageGRID Appliance Installer on installed appliance nodes. This procedure is only required when you are installing an appliance that contains an earlier version of the StorageGRID Appliance Installer.

### Steps

1. From the StorageGRID Appliance Installer, select **Advanced > Upgrade Firmware**.
2. Compare the Current Firmware version to the software version installed on your StorageGRID system (from the Grid Manager select **Help > About**).

The second digit in the two versions should match. For example, if your StorageGRID system is running version 11.5.x.y, the StorageGRID Appliance Installer version should be 3.5.z.

3. If the appliance has a down-level version of the StorageGRID Appliance Installer, go to the NetApp Downloads page for StorageGRID.

## NetApp Downloads: StorageGRID

Sign in with the username and password for your NetApp account.

4. Download the appropriate version of the **Support file for StorageGRID Appliances** and the corresponding checksum file.

The Support file for StorageGRID Appliances file is a .zip archive that contains the current and previous firmware versions for all StorageGRID appliance models, in subdirectories for each controller type.

After downloading the Support file for StorageGRID Appliances file, extract the .zip archive and see the README file for important information about installing the StorageGRID Appliance Installer.

5. Follow the instructions on the Upgrade Firmware page of the StorageGRID Appliance Installer to perform these steps:
  - a. Upload the appropriate support file (firmware image) for your controller type and the checksum file.
  - b. Upgrade the inactive partition.
  - c. Reboot and swap partitions.
  - d. Upgrade the second partition.

### Related information

[Accessing the StorageGRID Appliance Installer](#)

### Configuring network links (SG6000)

You can configure network links for the ports used to connect the appliance to the Grid Network, the Client Network, and the Admin Network. You can set the link speed as well as the port and network bond modes.

### What you'll need

If you are cloning an appliance node, configure network links for the target appliance for all links used by the source appliance node.

If you plan to use the 25-GbE link speed:


- You are using SFP28 TwinAx cables, or you have installed SFP28 transceivers in the network ports you plan to use.
- You have connected the network ports to switches that can support these features.
- You understand how to configure the switches to use this higher speed.

If you plan to use Aggregate port bond mode, LACP network bond mode, or VLAN tagging:

- You have connected the network ports on the appliance to switches that can support VLAN and LACP.
- If multiple switches are participating in the LACP bond, the switches support multi-chassis link aggregation groups (MLAG), or equivalent.
- You understand how to configure the switches to use VLAN, LACP, and MLAG or equivalent.
- You know the unique VLAN tag to use for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.


### About this task

This figure shows how the four network ports are bonded in fixed port bond mode (default configuration).



	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

This figure shows how the four network ports are bonded in aggregate port bond mode.



	Which ports are bonded
1	All four ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

The table summarizes the options for configuring the four network ports. The default settings are shown in bold. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

• **Fixed (default) port bond mode**

Network bond mode	Client Network disabled (default)	Client Network enabled
<b>Active-Backup (default)</b>	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an active-backup bond for the Grid Network.</li> <li>• Ports 1 and 3 are not used.</li> <li>• A VLAN tag is optional.</li> </ul>	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an active-backup bond for the Grid Network.</li> <li>• Ports 1 and 3 use an active-backup bond for the Client Network.</li> <li>• VLAN tags can be specified for both networks for the convenience of the network administrator.</li> </ul>

Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad)	<ul style="list-style-type: none"> <li>Ports 2 and 4 use an LACP bond for the Grid Network.</li> <li>Ports 1 and 3 are not used.</li> <li>A VLAN tag is optional.</li> </ul>	<ul style="list-style-type: none"> <li>Ports 2 and 4 use an LACP bond for the Grid Network.</li> <li>Ports 1 and 3 use an LACP bond for the Client Network.</li> <li>VLAN tags can be specified for both networks for the convenience of the network administrator.</li> </ul>

• **Aggregate port bond mode**

Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad) only	<ul style="list-style-type: none"> <li>Ports 1-4 use a single LACP bond for the Grid Network.</li> <li>A single VLAN tag identifies Grid Network packets.</li> </ul>	<ul style="list-style-type: none"> <li>Ports 1-4 use a single LACP bond for the Grid Network and the Client Network.</li> <li>Two VLAN tags allow Grid Network packets to be segregated from Client Network packets.</li> </ul>

See “Network port connections for the SG6000-CN controller” for more information about port bond and network bond modes.

This figure shows how the two 1-GbE management ports on the SG6000-CN controller are bonded in Active-Backup network bond mode for the Admin Network.




**Steps**

1. From the StorageGRID Appliance Installer, click **Configure Networking > Link Configuration**.

The Network Link Configuration page displays a diagram of your appliance with the network and management ports numbered.

## Network Link Configuration



⚠ You might lose your connection if you make changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL using one of the other IP addresses assigned to the appliance.

The Link Status table lists the link state (up/down) and speed (1/10/25/40/100 Gbps) of the numbered ports.

### Link Status

Link	State	Speed (Gbps)
1	Up	10
2	Up	10
3	Down	N/A
4	Down	N/A
5	Up	1
6	Up	1

The first time you access this page:

- **Link Speed** is set to **10GbE**.
- **Port bond mode** is set to **Fixed**.
- **Network bond mode** is set to **Active-Backup** for the Grid Network.
- The **Admin Network** is enabled, and the network bond mode is set to **Independent**.
- The **Client Network** is disabled.

## Link Settings

Link speed

Port bond mode  Fixed  Aggregate

Choose Fixed port bond mode if you want to use ports 2 and 4 for the Grid Network and ports 1 and 3 for the Client Network (if enabled). Choose Aggregate port bond mode if you want all connected ports to share a single LACP bond for both the Grid and Client Networks.

## Grid Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

Enable VLAN (802.1q) tagging

MAC Addresses 50:6b:4b:42:d7:00 50:6b:4b:42:d7:01 50:6b:4b:42:d7:24 50:6b:4b:42:d7:25

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

## Admin Network

Enable network

Network bond mode  Independent  Active-Backup

Connect the Admin Network to port 5. Leave port 6 unconnected. If necessary, you can make a temporary direct Ethernet connection to port 6 and use link-local IP address 169.254.0.1 for access.

MAC Addresses d8:c4:97:2a:e4:95

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

## Client Network

Enable network

Enabling the Client Network causes the default gateway for this node to move to the Client Network. Before enabling the Client Network, ensure that you've added all necessary subnets to the Grid Network Subnet List. Otherwise, the connection to the node might be lost.

2. If you plan to use the 25-GbE link speed for the network ports, select **25GbE** from the Link speed drop-down list.

The network switches you are using for the Grid Network and the Client Network must also support and be configured for this speed. You must use SFP28 TwinAx cables or optical cables and SFP28 transceivers.

3. Enable or disable the StorageGRID networks you plan to use.

The Grid Network is required. You cannot disable this network.

- a. If the appliance is not connected to the Admin Network, unselect the **Enable network** check box for the Admin Network.

#### **Admin Network**

---

Enable network

- b. If the appliance is connected to the Client Network, select the **Enable network** check box for the Client Network.

The Client Network settings for the network ports are now shown.

4. Refer to the table, and configure the port bond mode and the network bond mode.

This example shows:

- **Aggregate** and **LACP** selected for the Grid and the Client networks. You must specify a unique VLAN tag for each network. You can select values between 0 and 4095.
- **Active-Backup** selected for the Admin Network.

### Link Settings

Link speed

Port bond mode  Fixed  Aggregate

Choose Fixed port bond mode if you want to use ports 2 and 4 for the Grid Network and ports 1 and 3 for the Client Network (if enabled). Choose Aggregate port bond mode if you want all connected ports to share a single LACP bond for both the Grid and Client Networks.

### Grid Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

### Admin Network

Enable network

Network bond mode  Independent  Active-Backup

Connect the Admin Network to ports 5 and 6. If necessary, you can make a temporary direct Ethernet connection by disconnecting ports 5 and 6, then connecting to port 6 and using link-local IP address 169.254.0.1 for access.

### Client Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

5. When you are satisfied with your selections, click **Save**.



You might lose your connection if you made changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:

**`https://SG6000-CN_Controller_IP:8443`**

### Related information

[Port bond modes for the SG6000-CN controller](#)

[Configuring StorageGRID IP addresses](#)

## Configuring StorageGRID IP addresses

You use the StorageGRID Appliance Installer to configure the IP addresses and routing information used for the appliance Storage Node on the StorageGRID Grid, Admin, and Client Networks.

### About this task

You must either assign a static IP for the appliance on each connected network or assign a permanent lease for the address on the DHCP server.

If you want to change the link configuration, see the instructions for changing the link configuration of the SG6000-CN controller.

### Steps

1. In the StorageGRID Appliance Installer, select **Configure Networking > IP Configuration**.

The IP Configuration page appears.

2. To configure the Grid Network, select either **Static** or **DHCP** in the **Grid Network** section of the page.

## Grid Network

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

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

 All required Grid Network subnets must also be defined in the Grid Network Subnet List on the Primary Admin Node before starting installation.

Subnets (CIDR)  



MTU  

3. If you selected **Static**, follow these steps to configure the Grid Network:

- Enter the static IPv4 address, using CIDR notation.
- Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

- If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

**`https://services_appliance_IP:8443`**

e. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

f. Click **Save**.

4. If you selected **DHCP**, follow these steps to configure the Grid Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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

d. Click **Save**.

5. To configure the Admin Network, select either **Static** or **DHCP** in the **Admin Network** section of the page.



To configure the Admin Network, you must enable the Admin Network on the Link Configuration page.

## Admin Network

The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites.

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

Subnets (CIDR)  +

MTU

6. If you selected **Static**, follow these steps to configure the Admin Network:

a. Enter the static IPv4 address, using CIDR notation, for Management Port 1 on the appliance.

Management Port 1 is the left of the two 1-GbE RJ45 ports on the right end of the appliance.

b. Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as

9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

**`https://services_appliance:8443`**

e. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route cannot be made to use the Admin Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

f. Click **Save**.

7. If you selected **DHCP**, follow these steps to configure the Admin Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route cannot be made to use the Admin Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

8. To configure the Client Network, select either **Static** or **DHCP** in the **Client Network** section of the page.



To configure the Client Network, you must enable the Client Network on the Link Configuration page.

## Client Network

The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network enables grid nodes to communicate with any subnet reachable through the Client Network gateway. The Client Network does not become operational until you complete the StorageGRID configuration steps.

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

MTU

9. If you selected **Static**, follow these steps to configure the Client Network:

- Enter the static IPv4 address, using CIDR notation.
- Click **Save**.
- Confirm that the IP address for the Client Network gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and cannot be moved to another interface while the Client Network is enabled.

- If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

- Click **Save**.

10. If you selected **DHCP**, follow these steps to configure the Client Network:

- After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address** and **Gateway** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and cannot be moved to another interface while the Client Network is enabled.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

## Related information

[Changing the link configuration of the SG6000-CN controller](#)

## Verifying network connections

You should confirm you can access the StorageGRID networks you are using from the appliance. To validate routing through network gateways, you should test connectivity between the StorageGRID Appliance Installer and IP addresses on different subnets. You can also verify the MTU setting.

## Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Ping and MTU Test**.

The Ping and MTU Test page appears.

### Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

#### Ping and MTU Test

Network: Grid

Destination IPv4 Address or FQDN: [ ]

Test MTU:

Test Connectivity

2. From the **Network** drop-down box, select the network you want to test: Grid, Admin, or Client.

3. Enter the IPv4 address or fully qualified domain name (FQDN) for a host on that network.

For example, you might want to ping the gateway on the network or the primary Admin Node.

4. Optionally, select the **Test MTU** check box to verify the MTU setting for the entire path through the network to the destination.

For example, you can test the path between the appliance node and a node at a different site.

5. Click **Test Connectivity**.

If the network connection is valid, the "Ping test passed" message appears, with the ping command output listed.

### Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

#### Ping and MTU Test

Network	<input type="text" value="Grid"/>
Destination IPv4 Address or FQDN	<input type="text" value="10.96.104.223"/>
Test MTU	<input checked="" type="checkbox"/>
<input type="button" value="Test Connectivity"/>	

Ping test passed

#### Ping command output

```
PING 10.96.104.223 (10.96.104.223) 1472(1500) bytes of data.  
1480 bytes from 10.96.104.223: icmp_seq=1 ttl=64 time=0.318 ms  
  
--- 10.96.104.223 ping statistics ---  
1 packets transmitted, 1 received, 0% packet loss, time 0ms  
rtt min/avg/max/mdev = 0.318/0.318/0.318/0.000 ms  
  
Found MTU 1500 for 10.96.104.223 via br0
```

### Related information

[Configuring network links \(SG6000\)](#)

[Changing the MTU setting](#)

### Verifying port-level network connections

To ensure that access between the StorageGRID Appliance Installer and other nodes is

not obstructed by firewalls, confirm that the StorageGRID Appliance Installer can connect to a specific TCP port or set of ports at the specified IP address or range of addresses.

### About this task

Using the list of ports provided in the StorageGRID Appliance Installer, you can test the connectivity between the appliance and the other nodes in your Grid Network.

Additionally, you can test connectivity on the Admin and Client Networks and on UDP ports, such as those used for external NFS or DNS servers. For a list of these ports, see the port reference in the StorageGRID networking guidelines.



The Grid Network ports listed in the port connectivity table are valid only for StorageGRID version 11.5.0. To verify which ports are correct for each node type, you should always consult the networking guidelines for your version of StorageGRID.

### Steps

1. From the StorageGRID Appliance Installer, click **Configure Networking > Port Connectivity Test (nmap)**.

The Port Connectivity Test page appears.

The port connectivity table lists node types that require TCP connectivity on the Grid Network. For each node type, the table lists the Grid Network ports that should be accessible to your appliance.

The following node types require TCP connectivity on the Grid Network.

Node Type	Grid Network Ports
Admin Node	22,80,443,1504,1505,1506,1508,7443,9999
Storage Node without ADC	22,1139,1502,1506,1511,7001,9042,9999,18002,18017,18019,18082,18083,18200
Storage Node with ADC	22,1139,1501,1502,1506,1511,7001,9042,9999,18000,18001,18002,18003,18017,18019,18082,18083,18200,19000
API Gateway	22,1506,1507,9999
Archive Node	22,1506,1509,9999,11139

You can test the connectivity between the appliance ports listed in the table and the other nodes in your Grid Network.

2. From the **Network** drop-down, select the network you want to test: **Grid**, **Admin**, or **Client**.
3. Specify a range of IPv4 addresses for the hosts on that network.

For example, you might want to probe the gateway on the network or the primary Admin Node.

Specify a range using a hyphen, as shown in the example.

4. Enter a TCP port number, a list of ports separated by commas, or a range of ports.

The following node types require TCP connectivity on the Grid Network.

Node Type	Grid Network Ports
Admin Node	22,80,443,1504,1505,1506,1508,7443,9999
Storage Node without ADC	22,1139,1502,1506,1511,7001,9042,9999,18002,18017,18019,18082,18083,18200
Storage Node with ADC	22,1139,1501,1502,1506,1511,7001,9042,9999,18000,18001,18002,18003,18017,18019,18082,18083,18200,19000
API Gateway	22,1506,1507,9999
Archive Node	22,1506,1509,9999,11139

### Port Connectivity Test

Network

IPv4 Address Ranges

Port Ranges

Protocol  TCP  UDP

### 5. Click **Test Connectivity**.

- If the selected port-level network connections are valid, the “Port connectivity test passed” message appears in a green banner. The nmap command output is listed below the banner.

Port connectivity test passed

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Fri Nov 13 18:32:03 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,2022 10.224.6.160-161
Nmap scan report for 10.224.6.160
Host is up (0.00072s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp   open  down

Nmap scan report for 10.224.6.161
Host is up (0.00060s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp   open  down

# Nmap done at Fri Nov 13 18:32:04 2020 -- 2 IP addresses (2 hosts up) scanned in 0.55 seconds
```

- If a port-level network connection is made to the remote host, but the host is not listening on one or more of the selected ports, the “Port connectivity test failed” message appears in a yellow banner. The nmap command output is listed below the banner.

Any remote port the host is not listening to has a state of “closed.” For example, you might see this yellow banner when the node you are trying to connect to is in a pre-installed state and the StorageGRID NMS service is not yet running on that node.

 Port connectivity test failed  
Connection not established. Services might not be listening on target ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:07:02 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,80,443,1504,1505,1506,1508,7443,9999
Nmap scan report for 172.16.4.71
Host is up (0.00020s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
80/tcp    open  http
443/tcp   open  https
1504/tcp   closed evb-elm
1505/tcp   open  funkproxy
1506/tcp   open  utcd
1508/tcp   open  diagmond
7443/tcp   open  oracleas-https
9999/tcp   open  abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:07:03 2020 -- 1 IP address (1 host up) scanned in 0.59 seconds
```

- If a port-level network connection cannot be made for one or more selected ports, the “Port connectivity test failed” message appears in a red banner. The nmap command output is listed below the banner.

The red banner indicates that a TCP connection attempt to a port on the remote host was made, but nothing was returned to the sender. When no response is returned, the port has a state of “filtered” and is likely blocked by a firewall.



Ports with “closed” are also listed.

 Port connectivity test failed  
Connection failed to one or more ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:11:01 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,79,80,443,1504,1505,1506,1508,7443,9999 172.16.4.71
Nmap scan report for 172.16.4.71
Host is up (0.00029s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
79/tcp    filtered finger
80/tcp    open  http
443/tcp   open  https
1504/tcp   closed evb-elm
1505/tcp   open  funkproxy
1506/tcp   open  utcd
1508/tcp   open  diagmond
7443/tcp   open  oracleas-https
9999/tcp   open  abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:11:02 2020 -- 1 IP address (1 host up) scanned in 1.60 seconds
```

## Related information

[Network guidelines](#)

## Accessing and Configuring SANtricity System Manager

You can use SANtricity System Manager to monitor the status of the storage controllers, storage disks, and other hardware components in the storage controller shelf. You can

also configure a proxy for E-Series AutoSupport that enables you to send AutoSupport messages from the appliance without the use of the management port.

### Steps

- [Setting up and Accessing SANtricity System Manager](#)
- [Reviewing the hardware status in SANtricity System Manager](#)
- [Setting the IP addresses for the storage controllers using the StorageGRID Appliance Installer](#)

### Setting up and Accessing SANtricity System Manager

You might need to access SANtricity System Manager on the storage controller to monitor the hardware in the storage controller shelf or to configure E-Series AutoSupport.

### What you'll need

- You are using a supported web browser.
- To access SANtricity System Manager through Grid Manager, you must have installed StorageGRID, and you must have the Storage Appliance Administrator permission or Root Access permission.
- To access SANtricity System Manager using the StorageGRID Appliance Installer, you must have the SANtricity System Manager administrator username and password.
- To access SANtricity System Manager directly using a web browser, you must have the SANtricity System Manager administrator username and password.



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager or the StorageGRID Appliance Installer. You can check your firmware version by using the StorageGRID Appliance Installer and selecting **Help > About**.



Accessing SANtricity System Manager from the Grid Manager or from the Appliance Installer is generally meant only for monitoring your hardware and configuring E-Series AutoSupport. Many features and operations within SANtricity System Manager such as upgrading firmware do not apply to monitoring your StorageGRID appliance. To avoid issues, always follow the hardware installation and maintenance instructions for your appliance.

### About this task

There are three ways to access SANtricity System Manager, depending upon what stage of the installation and configuration process you are in:

- If the appliance has not yet been deployed as a node in your StorageGRID system, you should use the Advanced tab in the StorageGRID Appliance Installer.



Once the node is deployed, you can no longer use the StorageGRID Appliance Installer to access SANtricity System Manager.

- If the appliance has been deployed as a node in your StorageGRID system, use the SANtricity System Manager tab on the Nodes page in Grid Manager.
- If you cannot use the StorageGRID Appliance Installer or Grid Manager, you can access SANtricity System Manager directly using a web browser connected to the management port.

This procedure includes steps for your initial access to SANtricity System Manager. If you have already set up SANtricity System Manager, go to the [configure hardware alerts](#) step.



Using either the Grid Manager or the StorageGRID Appliance Installer enables you to access SANtricity System Manager without having to configure or connect the management port of the appliance.

You use SANtricity System Manager to monitor the following:

- Performance data such as storage array level performance, I/O latency, CPU utilization, and throughput
- Hardware component status
- Support functions including viewing diagnostic data

You can use SANtricity System Manager to configure the following settings:

- Email alerts, SNMP alerts, or syslog alerts for the components in the storage controller shelf
- E-Series AutoSupport settings for the components in the storage controller shelf.

For additional details on E-Series AutoSupport, see the E-Series documentation center.

[NetApp E-Series Systems Documentation Site](#)

- Drive Security keys, which are needed to unlock secured drives (this step is required if the Drive Security feature is enabled)
- Administrator password for accessing SANtricity System Manager

## Steps

1. Do one of the following:

- Use the StorageGRID Appliance Installer and select **Advanced > SANtricity System Manager**
- Use the Grid Manager and select **Nodes > appliance Storage Node > SANtricity System Manager**



If these options are not available or the login page does not appear, you must use the IP address of the storage controller. Access SANtricity System Manager by browsing to the storage controller IP:

**`https://Storage_Controller_IP`**

The login page for SANtricity System Manager appears.

2. Set or enter the administrator password.



SANtricity System Manager uses a single administrator password that is shared among all users.

The Set Up wizard appears.

## 1 Welcome

## 2 Verify Hardware

## 3 Verify Hosts

## 4 Select Applications

## 5 Define Workloads

## 6 Acc

Welcome to the SANtricity® System Manager! With System Manager, you can...

- Configure your storage array and set up alerts.
- Monitor and troubleshoot any problems when they occur.
- Keep track of how your system is performing in real time.

Cancel


Next &gt;

3. Select **Cancel** to close the wizard.



Do not complete the Set Up wizard for a StorageGRID appliance.

The SANtricity System Manager home page appears.



4. Configure hardware alerts.

- a. Select **Help** to access the online help for SANtricity System Manager.
  - b. Use the **Settings > Alerts** section of the online help to learn about alerts.
  - c. Follow the “How To” instructions to set up email alerts, SNMP alerts, or syslog alerts.
5. Manage AutoSupport for the components in the storage controller shelf.
- a. Select **Help** to access the online help for SANtricity System Manager.
  - b. Use the **Support > Support Center** section of the online help to learn about the AutoSupport feature.
  - c. Follow the “How To” instructions to manage AutoSupport.

For specific instructions on configuring a StorageGrid proxy for sending E-Series AutoSupport messages without using the management port, go to the instructions for administering StorageGRID and search for "proxy settings for E-Series AutoSupport."

### [Administer StorageGRID](#)

6. If the Drive Security feature is enabled for the appliance, create and manage the security key.
- a. Select **Help** to access the online help for SANtricity System Manager.
  - b. Use the **Settings > System > Security key management** section of the online help to learn about Drive Security.
  - c. Follow the “How To” instructions to create and manage the security key.
7. Optionally, change the administrator password.
- a. Select **Help** to access the online help for SANtricity System Manager.
  - b. Use the **Home > Storage array administration** section of the online help to learn about the administrator password.
  - c. Follow the “How To” instructions to change the password.

### **Related information**

[Web browser requirements](#)

[Setting the IP addresses for the storage controllers using the StorageGRID Appliance Installer](#)

### **Reviewing the hardware status in SANtricity System Manager**

You can use SANtricity System Manager to monitor and manage the individual hardware components in the storage controller shelf and to review hardware diagnostic and environmental information, such as component temperatures, as well as issues related to the drives.

### **What you'll need**

- You are using a supported web browser.
- To access SANtricity System Manager through Grid Manager, you must have the Storage Appliance Administrator permission or Root Access permission.
- To access SANtricity System Manager using the StorageGRID Appliance Installer, you must have the SANtricity System Manager administrator username and password.
- To access SANtricity System Manager directly using a web browser, you must have the SANtricity System Manager administrator username and password.



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager or the StorageGRID Appliance Installer.



Accessing SANtricity System Manager from the Grid Manager or from the Appliance Installer is generally meant only for monitoring your hardware and configuring E-Series AutoSupport. Many features and operations within SANtricity System Manager such as upgrading firmware do not apply to monitoring your StorageGRID appliance. To avoid issues, always follow the hardware installation and maintenance instructions for your appliance.


## Steps

1. Access SANtricity System Manager.


### Setting up and Accessing SANtricity System Manager

2. Enter the administrator username and password if required.
3. Click **Cancel** to close the Set Up wizard and to display the SANtricity System Manager home page.

The SANtricity System Manager home page appears. In SANtricity System Manager, the controller shelf is referred to as a storage array.



4. Review the information displayed for appliance hardware and confirm that all hardware components have a status of Optimal.
  - a. Click the **Hardware** tab.
  - b. Click **Show back of shelf**.



From the back of the shelf, you can view both storage controllers, the battery in each storage controller, the two power canisters, the two fan canisters, and expansion shelves (if any). You can also view component temperatures.

- c. To see the settings for each storage controller, select the controller, and select **View settings** from the context menu.
- d. To see the settings for other components in the back of the shelf, select the component you want to view.
- e. Click **Show front of shelf**, and select the component you want to view.

From the front of the shelf, you can view the drives and the drive drawers for the storage controller shelf or the expansion shelves (if any).

If the status of any component is Needs Attention, follow the steps in the Recovery Guru to resolve the issue or contact technical support.

#### Setting the IP addresses for the storage controllers using the StorageGRID Appliance Installer

Management port 1 on each storage controller connects the appliance to the management network for SANtricity System Manager. If you cannot access to the SANtricity System Manager from the StorageGRID Appliance Installer, you must set a static IP address for each storage controller to ensure that you do not lose your management connection to the hardware and the controller firmware in the controller shelf.

#### What you'll need

- You are using any management client that can connect to the StorageGRID Admin Network, or you have a service laptop.
- The client or service laptop has a supported web browser.

## About this task

DHCP-assigned addresses can change at any time. Assign static IP addresses to the controllers to ensure consistent accessibility.



Follow this procedure only if you do not have access to SANtricity System Manager from the StorageGRID Appliance Installer (**Advanced > SANtricity System Manager**) or Grid Manager (**Nodes > SANtricity System Manager**).

## Steps

1. From the client, enter the URL for the StorageGRID Appliance Installer:  
**`https://Appliance_Controller_IP:8443`**

For *Appliance\_Controller\_IP*, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > Storage Controller Network Configuration**.

The Storage Controller Network Configuration page appears.

3. Depending on your network configuration, select **Enabled** for IPv4, IPv6, or both.
4. Make a note of the IPv4 address that is automatically displayed.

DHCP is the default method for assigning an IP address to the storage controller management port.



It might take a few minutes for the DHCP values to appear.

IPv4 Address Assignment	<input type="radio"/> Static	<input checked="" type="radio"/> DHCP
IPv4 Address (CIDR)	<input type="text" value="10.224.5.166/21"/>	
Default Gateway	<input type="text" value="10.224.0.1"/>	

5. Optionally, set a static IP address for the storage controller management port.



You should either assign a static IP for the management port or assign a permanent lease for the address on the DHCP server.

- a. Select **Static**.
- b. Enter the IPv4 address, using CIDR notation.
- c. Enter the default gateway.

IPv4 Address Assignment     Static     DHCP

IPv4 Address (CIDR)	10.224.2.200/21
Default Gateway	10.224.0.1

d. Click **Save**.

It might take a few minutes for your changes to be applied.

When you connect to SANtricity System Manager, you will use the new static IP address as the URL:  
**`https://Storage_Controller_IP`**

## Configuring the BMC interface

The user interface for the baseboard management controller (BMC) on the SG6000-CN controller provides status information about the hardware and allows you to configure SNMP settings and other options for the SG6000-CN controller.

### Steps

- [Changing the root password for the BMC interface](#)
- [Setting the IP address for the BMC management port](#)
- [Accessing the BMC interface](#)
- [Configuring SNMP settings for the SG6000-CN controller](#)
- [Setting up email notifications for alerts](#)

### Changing the root password for the BMC interface

For security, you must change the password for the BMC's root user.

### What you'll need

- The management client is using a supported web browser.

### About this task

When you first install the appliance, the BMC uses a default password for the root user (`root/calvin`). You must change the password for the root user to secure your system.


### Steps

1. From the client, enter the URL for the StorageGRID Appliance Installer:  
**`https://Appliance_Controller_IP:8443`**

For `Appliance_Controller_IP`, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > BMC Configuration**.



The Baseboard Management Controller Configuration page appears.

3. Enter a new password for the root account in the two fields provided.

## Baseboard Management Controller Configuration

### User Settings

Root Password	<input type="password" value="....."/>
Confirm Root Password	<input type="password" value="....."/>

4. Click **Save**.

### Setting the IP address for the BMC management port

Before you can access the BMC interface, you must configure the IP address for the BMC management port on the SG6000-CN controller.

### What you'll need

- The management client is using a supported web browser.
- You are using any management client that can connect to a StorageGRID network.
- The BMC management port is connected to the management network you plan to use.



### About this task

For support purposes, the BMC management port allows low-level hardware access.



You should only connect this port to a secure, trusted, internal management network. If no such network is available, leave the BMC port unconnected or blocked, unless a BMC connection is requested by technical support.

### Steps


1. From the client, enter the URL for the StorageGRID Appliance Installer:

`https://SG6000-CN_Controller_IP:8443`

For SG6000-CN\_Controller\_IP, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware** > **BMC Configuration**.



The Baseboard Management Controller Configuration page appears.

3. Make a note of the IPv4 address that is automatically displayed.

DHCP is the default method for assigning an IP address to this port.



It might take a few minutes for the DHCP values to appear.

#### Baseboard Management Controller Configuration

##### LAN IP Settings

IP Assignment	<input type="radio"/> Static	<input checked="" type="radio"/> DHCP
MAC Address	<input type="text" value="d8:c4:97:28:50:62"/>	
IPv4 Address (CIDR)	<input type="text" value="10.224.3.225/21"/>	
Default gateway	<input type="text" value="10.224.0.1"/>	

4. Optionally, set a static IP address for the BMC management port.



You should either assign a static IP for the BMC management port or assign a permanent lease for the address on the DHCP server.

- a. Select **Static**.
- b. Enter the IPv4 address, using CIDR notation.
- c. Enter the default gateway.

## Baseboard Management Controller Configuration

### LAN IP Settings

IP Assignment	<input checked="" type="radio"/> Static <input type="radio"/> DHCP
MAC Address	d8:c4:97:28:50:62
IPv4 Address (CIDR)	10.224.3.225/21
Default gateway	10.224.0.1

d. Click **Save**.

It might take a few minutes for your changes to be applied.

### Accessing the BMC interface

You can access the BMC interface on the SG6000-CN controller using the DHCP or static IP address for the BMC management port.

### What you'll need

- The BMC management port on the SG6000-CN controller is connected to the management network you plan to use.



- The management client is using a supported web browser.

### Steps

1. Enter the URL for the BMC interface:

**`https://BMC_Port_IP`**

For *BMC\_Port\_IP*, use the DHCP or static IP address for the BMC management port.

The BMC sign-in page appears.

2. Enter the root username and password, using the password you set when you changed the default root password:

**`root`**

**`password`**



# NetApp®

root

.....

Remember Username

Sign me in

[I forgot my password](#)

### 3. Select **Sign me in**.

The BMC dashboard appears.

### 4. Optionally, create additional users by selecting **Settings > User Management** and clicking on any “disabled” user.



When users sign in for the first time, they might be prompted to change their password for increased security.

#### Related information

[Changing the root password for the BMC interface](#)

## Configuring SNMP settings for the SG6000-CN controller

If you are familiar with configuring SNMP for hardware, you can use the BMC interface to configure the SNMP settings for the SG6000-CN controller. You can provide secure community strings, enable SNMP Trap, and specify up to five SNMP destinations.

### What you'll need

- You know how to access the BMC dashboard.
- You have experience in configuring SNMP settings for SNMPv1-v2c equipment.

### Steps

1. From the BMC dashboard, select **Settings > SNMP Settings**.
2. On the SNMP Settings page, select **Enable SNMP V1/V2**, and then provide a Read-Only Community String and a Read-Write Community String.

The Read-Only Community String is like a user ID or password. You should change this value to prevent intruders from getting information about your network setup. The Read-Write Community String protects the device against unauthorized changes.

3. Optionally, select **Enable Trap**, and enter the required information.



Enter the Destination IP for each SNMP trap using an IP address. Fully qualified domain names are not supported.

Enable traps if you want the SG6000-CN controller to send immediate notifications to an SNMP console when it is in an unusual state. Traps might indicate hardware failures of various components or temperature thresholds being exceeded.

4. Optionally, click **Send Test Trap** to test your settings.
5. If the settings are correct, click **Save**.

### Setting up email notifications for alerts


If you want email notifications to be sent when alerts occur, you must use the BMC interface to configure SMTP settings, users, LAN destinations, alert policies, and event filters.

### What you'll need

You know how to access the BMC dashboard.

### About this task

In the BMC interface, you use the **SMTP Settings**, **User Management**, and **Platform Event Filters** options on the Settings page to configure email notifications.



## Steps

1. Configure the SMTP settings.

- a. Select **Settings > SMTP Settings**.
- b. For Sender Email ID, enter a valid email address.

This email address is provided as the From address when the BMC sends email.

2. Set up users to receive alerts.

- a. From the BMC dashboard, select **Settings > User Management**.
- b. Add at least one user to receive alert notifications.

The email address you configure for a user is the address the BMC sends alert notifications to. For example, you could add a generic user, such as “notification-user,” and use the email address of a technical support team email distribution list.

3. Configure the LAN destination for alerts.

- a. Select **Settings > Platform Event Filters > LAN Destinations**.
- b. Configure at least one LAN destination.
  - Select **Email** as the Destination Type.
  - For BMC Username, select a user name that you added earlier.
  - If you added multiple users and want all of them to receive notification emails, you must add a LAN Destination for each user.

c. Send a test alert.

4. Configure alert policies so you can define when and where the BMC sends alerts.

- a. Select **Settings > Platform Event Filters > Alert Policies**.
- b. Configure at least one alert policy for each LAN destination.
  - For Policy Group Number, select **1**.
  - For Policy Action, select **Always send alert to this destination**.
  - For LAN Channel, select **1**.
  - In the Destination Selector, select the LAN destination for the policy.

5. Configure event filters to direct alerts for different event types to the appropriate users.

- a. Select **Settings > Platform Event Filters > Event Filters**.
- b. For Alert Policy Group Number, enter **1**.
- c. Create filters for every event you want the Alert Policy Group to be notified about.
  - You can create event filters for power actions, specific sensor events, or all events.
  - If you are uncertain which events to monitor, select **All Sensors** for Sensor Type and **All Events** for Event Options. If you receive unwanted notifications, you can change your selections later.

### Optional: Enabling node encryption

If you enable node encryption, the disks in your appliance can be protected by secure key management server (KMS) encryption against physical loss or removal from the site. You must select and enable node encryption during appliance installation and cannot unselect node encryption once the KMS encryption process starts.

#### What you'll need

Review the information about KMS in the instructions for administering StorageGRID.

#### About this task

An appliance that has node encryption enabled connects to the external key management server (KMS) that is configured for the StorageGRID site. Each KMS (or KMS cluster) manages the encryption keys for all appliance nodes at the site. These keys encrypt and decrypt the data on each disk in an appliance that has node encryption enabled.

A KMS can be set up in Grid Manager before or after the appliance is installed in StorageGRID. See the information about KMS and appliance configuration in the instructions for administering StorageGRID for additional details.

- If a KMS is set up before installing the appliance, KMS-controlled encryption begins when you enable node encryption on the appliance and add it to a StorageGRID site where KMS is configured.
- If a KMS is not set up before you install the appliance, KMS-controlled encryption is performed on each appliance that has node encryption enabled as soon as a KMS is configured and available for the site that contains the appliance node.



Any data that exists before an appliance that has node encryption enabled connects to the configured KMS is encrypted with a temporary key that is not secure. The appliance is not protected from removal or theft until the key is set to a value provided by the KMS.

Without the KMS key needed to decrypt the disk, data on the appliance cannot be retrieved and is effectively lost. This is the case whenever the decryption key cannot be retrieved from the KMS. The key becomes inaccessible if a customer clears the KMS configuration, a KMS key expires, connection to the KMS is lost, or the appliance is removed from the StorageGRID system where its KMS keys are installed.

#### Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.



After the appliance has been encrypted with a KMS key, the appliance disks cannot be decrypted without using the same KMS key.

## 2. Select **Configure Hardware > Node Encryption**.

The screenshot shows the NetApp StorageGRID Appliance Installer interface. The top navigation bar includes 'Home', 'Configure Networking', 'Configure Hardware', 'Monitor Installation', and 'Advanced'. The 'Configure Hardware' section is active, showing the 'Node Encryption' configuration page. The page title is 'Node Encryption'. Below the title, there is a descriptive paragraph: 'Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.' Below this is the 'Encryption Status' section, which contains a yellow warning box: 'You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.' Underneath the warning box, there is a checkbox labeled 'Enable node encryption' which is checked. A blue 'Save' button is located below the checkbox. At the bottom of the page, there is a section titled 'Key Management Server Details'.

## 3. Select **Enable node encryption**.

You can unselect **Enable node encryption** without risk of data loss until you select **Save** and the appliance node accesses the KMS encryption keys in your StorageGRID system and begins disk encryption. You are not able to disable node encryption after the appliance is installed.



After you add an appliance that has node encryption enabled to a StorageGRID site that has a KMS, you cannot stop using KMS encryption for the node.

## 4. Select **Save**.

## 5. Deploy the appliance as a node in your StorageGRID system.

KMS-controlled encryption begins when the appliance accesses the KMS keys configured for your StorageGRID site. The installer displays progress messages during the KMS encryption process, which might take a few minutes depending on the number of disk volumes in the appliance.



Appliances are initially configured with a random non-KMS encryption key assigned to each disk volume. The disks are encrypted using this temporary encryption key, that is not secure, until the appliance that has node encryption enabled accesses the KMS keys configured for your StorageGRID site.

### After you finish

You can view node-encryption status, KMS details, and the certificates in use when the appliance node is in maintenance mode.

### Related information

[Administer StorageGRID](#)

[Monitoring node encryption in maintenance mode](#)

## Optional: Changing the RAID mode (SG6000 only)

You can change to a different RAID mode on the appliance to accommodate your storage and recovery requirements. You can only change the mode before deploying the appliance Storage Node.

### What you'll need

- You are using any client that can connect to StorageGRID.
- The client has a supported web browser.

### About this task

Before deploying the appliance as a Storage Node, you can choose one of the following volume configuration options:

- **DDP**: This mode uses two parity drives for every eight data drives. This is the default and recommended mode for all appliances. When compared to RAID6, DDP delivers better system performance, reduced rebuild times after drive failures, and ease of management. DDP also provides drawer loss protection in 60-drive appliances.
- **DDP16**: This mode uses two parity drives for every 16 data drives, which results in higher storage efficiency compared to DDP. When compared to RAID6, DDP16 delivers better system performance, reduced rebuild times after drive failures, ease of management, and comparable storage efficiency. To use DDP16 mode, your configuration must contain at least 20 drives. DDP16 does not provide drawer loss protection.
- **RAID6**: This mode uses two parity drives for every 16 or more data drives. To use RAID 6 mode, your configuration must contain at least 20 drives. Although RAID6 can increase storage efficiency of the appliance when compared to DDP, it is not recommended for most StorageGRID environments.



If any volumes have already been configured or if StorageGRID was previously installed, changing the RAID mode causes the volumes to be removed and replaced. Any data on those volumes will be lost.

### Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

2. Select **Advanced > RAID Mode**.
3. On the **Configure RAID Mode** page, select the desired RAID mode from the Mode drop-down list.
4. Click **Save**.

### Related information

[NetApp E-Series Systems Documentation Site](#)

## Optional: Remapping network ports for the appliance

You might need to remap the internal ports on the appliance Storage Node to different external ports. For example, you might need to remap ports because of a firewall issue.

### What you'll need

- You have previously accessed the StorageGRID Appliance Installer.
- You have not configured and do not plan to configure load balancer endpoints.



If you remap any ports, you cannot use the same ports to configure load balancer endpoints. If you want to configure load balancer endpoints and have already remapped ports, follow the steps in the recovery and maintenance instructions for removing port remaps.

### Steps

1. From the StorageGRID Appliance Installer, click **Configure Networking > Remap Ports**.

The Remap Port page appears.

2. From the **Network** drop-down box, select the network for the port you want to remap: Grid, Admin, or Client.
3. From the **Protocol** drop-down box, select the IP protocol: TCP or UDP.
4. From the **Remap Direction** drop-down box, select which traffic direction you want to remap for this port: Inbound, Outbound, or Bi-directional.
5. For **Original Port**, enter the number of the port you want to remap.
6. For **Mapped-To Port**, enter the number of the port you want to use instead.
7. Click **Add Rule**.

The new port mapping is added to the table, and the remapping takes effect immediately.

### Remap Ports

If required, you can remap the internal ports on the appliance Storage Node to different external ports. For example, you might need to remap ports because of a firewall issue.

Network	Protocol	Remap Direction	Original Port	Mapped-To Port
<input type="radio"/> Grid	TCP	Bi-directional	1800	1801

8. To remove a port mapping, select the radio button for the rule you want to remove, and click **Remove Selected Rule**.

## Deploying an appliance Storage Node

After installing and configuring the storage appliance, you can deploy it as a Storage Node in a StorageGRID system. When you deploy an appliance as a Storage Node, you use the StorageGRID Appliance Installer included on the appliance.

### What you'll need

- If you are cloning an appliance node, continue following the process in recovery and maintenance.

#### Maintain & recover

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.
- 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.
- The primary Admin Node for the StorageGRID system has been deployed.
- 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 a service laptop with a supported web browser.

### About this task

Each storage appliance functions as a single Storage Node. Any appliance can connect to the Grid Network, the Admin Network, and the Client Network

To deploy an appliance Storage Node in a StorageGRID system, you access the StorageGRID Appliance Installer and perform the following steps:

- You specify or confirm the IP address of the primary Admin Node and the name of the Storage Node.
- You start the deployment and wait as volumes are configured and the software is installed.
- When the installation pauses partway through the appliance installation tasks, you resume the installation by signing into the Grid Manager, approving all grid nodes, and completing the StorageGRID installation and deployment processes.



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.

- If you are performing an expansion or recovery operation, follow the appropriate instructions:
  - To add an appliance Storage Node to an existing StorageGRID system, see the instructions for expanding a StorageGRID system.
  - To deploy an appliance Storage Node as part of a recovery operation, see instructions for recovery and maintenance.

### 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.

## 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

Connection state

Connection to 172.16.4.210 ready



### Node name

Node name




### Installation

Current state

Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

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	<ol style="list-style-type: none"> <li>Unselect the <b>Enable Admin Node discovery</b> check box.</li> <li>Enter the IP address manually.</li> <li>Click <b>Save</b>.</li> <li>Wait for the connection state for the new IP address to become ready.</li> </ol>
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"> <li>Select the <b>Enable Admin Node discovery</b> check box.</li> <li>Wait for the list of discovered IP addresses to be displayed.</li> <li>Select the primary Admin Node for the grid where this appliance Storage Node will be deployed.</li> <li>Click <b>Save</b>.</li> <li>Wait for the connection state for the new IP address to become ready.</li> </ol>

- In the **Node name** field, enter the name you want to use for this appliance node, and click **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.

- 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 installation and maintenance instructions for your appliance.



If you are deploying the Storage Node appliance as a node cloning target, stop the deployment process here and continue the node cloning procedure in recovery and maintenance.

[Maintain & recover](#)

- 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**.

- If your grid includes multiple appliance Storage Nodes, repeat these 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. This script applies only to Storage Nodes.

## Related information

[Expand your grid](#)

[Maintain & recover](#)

## Monitoring the storage 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**.

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 are not re-run. If you are re-running an installation, any tasks that do not need to be re-run are shown with a green status bar and a status of "Skipped."

2. Review the progress of the first two installation stages.

### 1. Configure storage

During this stage, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity software 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. Go to the next step.

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...

```

- Go to the Grid Manager, approve the pending storage node, and complete the StorageGRID installation process.

When you click **Install** from the Grid Manager, Stage 3 completes and stage 4, **Finalize Installation**, begins. When stage 4 completes, the controller is rebooted.

## Automating appliance installation and configuration

You can automate the installation and configuration of your appliances and configuration of the whole StorageGRID system.

### About this task

Automating installation and configuration can be useful for deploying multiple StorageGRID instances or one large, complex StorageGRID instance.

To automate installation and configuration, use one or more of the following options:

- Create a JSON file that specifies the configuration settings for your appliances. Upload the JSON file using the StorageGRID Appliance Installer.



You can use the same file to configure more than one appliance.

- Use the `StorageGRIDconfigure-sga.py` Python script to automate the configuration of your appliances.
- Use additional Python scripts to configure other components of the whole StorageGRID system (the "grid").



You can use StorageGRID automation Python scripts directly, or you can use them as examples of how to use the StorageGRID Installation REST API in grid deployment and configuration tools you develop yourself. See the information about downloading and extracting the StorageGRID installation files in the Recovery and Maintenance instructions.

## Automating appliance configuration using the StorageGRID Appliance Installer

You can automate the configuration of an appliance by using a JSON file that contains the configuration information. You upload the file using the StorageGRID Appliance Installer.

### What you'll need

- Your appliance must be on the latest firmware compatible with StorageGRID 11.5 or higher.
- You must be connected to the StorageGRID Appliance Installer on the appliance you are configuring using a supported browser.

### About this task

You can automate appliance configuration tasks such as configuring the following:

- Grid Network, Admin Network, and Client Network IP addresses
- BMC interface
- Network links
  - Port bond mode
  - Network bond mode
  - Link speed

Configuring your appliance using an uploaded JSON file is often more efficient than performing the configuration manually using multiple pages in the StorageGRID Appliance Installer, especially if you have to configure many nodes. You must apply the configuration file for each node one at a time.



Experienced users who want to automate both the installation and configuration of their appliances can use the `configure-sga.py` script.

[Automating installation and configuration of appliance nodes using the `configure-sga.py` script](#)

## Steps

1. Generate the JSON file using one of the following methods:

- The ConfigBuilder application

[ConfigBuilder.netapp.com](https://configbuilder.netapp.com)

- The `configure-sga.py` appliance configuration script. You can download the script from StorageGRID Appliance Installer (**Help > Appliance Configuration Script**). See the instructions on automating the configuration using the `configure-sga.py` script.

[Automating installation and configuration of appliance nodes using the configure-sga.py script](#)

The node names in the JSON file must follow these requirements:

- Must be a valid hostname containing at least 1 and no more than 32 characters
- Can use letters, numbers, and hyphens are allowed
- Cannot start or end with a hyphen or contain only numbers



Ensure that the node names (the top-level names) in the JSON file are unique, or you will not be able to configure more than one node using the JSON file.

2. Select **Advanced > Update Appliance Configuration**.

The Update Appliance Configuration page appears.

### Update Appliance Configuration

Use a JSON file to update this appliance's configuration. You can generate the JSON file from the [ConfigBuilder](#) application or from the [appliance configuration script](#).

**⚠** You might lose your connection if the applied configuration from the JSON file includes "link\_config" and/or "networks" sections. If you are not reconnected within 1 minute, re-enter the URL using one of the other IP addresses assigned to the appliance.

#### Upload JSON

JSON configuration	<input type="button" value="Browse"/>
Node name	<input type="button" value="-- Upload a file"/>
<input type="button" value="Apply JSON configuration"/>	

3. Select the JSON file with the configuration you want to upload.

- a. Select **Browse**.
- b. Locate and select the file.
- c. Select **Open**.

The file is uploaded and validated. When the validation process is complete, the file name is shown

next to a green check mark.



You might lose connection to the appliance if the configuration from the JSON file includes sections for "link\_config", "networks", or both. If you are not reconnected within 1 minute, re-enter the appliance URL using one of the other IP addresses assigned to the appliance.

#### Upload JSON

JSON configuration	<input type="button" value="Browse"/>	✓ appliances.orig.json
Node name	-- Select a node ▼	
<input type="button" value="Apply JSON configuration"/>		

The **Node name** drop down is populated with the top-level node names defined in the JSON file.



If the file is not valid, the file name is shown in red and an error message is displayed in a yellow banner. The invalid file is not applied to the appliance. You can use ConfigBuilder to ensure you have a valid JSON file.

4. Select a node from the list in the **Node name** drop down.

The **Apply JSON configuration** button is enabled.

#### Upload JSON

JSON configuration	<input type="button" value="Browse"/>	✓ appliances.orig.json
Node name	Lab-80-1000 ▼	
<input type="button" value="Apply JSON configuration"/>		

5. Select **Apply JSON configuration**.

The configuration is applied to the selected node.

### Automating installation and configuration of appliance nodes using the `configure-sga.py` script

You can use the `configure-sga.py` script to automate many of the installation and configuration tasks for StorageGRID appliance nodes, including installing and configuring a primary Admin Node. This script can be useful if you have a large number of appliances to configure. You can also use the script to generate a JSON file that contains appliance configuration information.

## What you'll need

- 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 primary Admin Node using the StorageGRID Appliance Installer.
- If you are installing the primary Admin Node, you know its IP address.
- If you are installing and configuring other nodes, the primary Admin Node has been deployed, and you know its IP address.
- For all nodes other than the primary Admin Node, 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 downloaded the `configure-sga.py` file. The file is included in the installation archive, or you can access it by clicking **Help > Appliance Installation Script** in the StorageGRID Appliance Installer.



This procedure is for advanced users with some experience using command-line interfaces. Alternatively, you can also use the StorageGRID Appliance Installer to automate the configuration.

[Automating appliance configuration using the StorageGRID Appliance Installer](#)

## Steps

1. Log in to the Linux machine you are using to run the Python script.
2. For general help with the script syntax and to see a list of the available parameters, enter the following:

```
configure-sga.py --help
```

The `configure-sga.py` script uses five subcommands:

- `advanced` for advanced StorageGRID appliance interactions, including BMC configuration and creating a JSON file containing the current configuration of the appliance
- `configure` for configuring the RAID mode, node name, and networking parameters
- `install` for starting a StorageGRID installation
- `monitor` for monitoring a StorageGRID installation
- `reboot` for rebooting the appliance

If you enter a subcommand (`advanced`, `configure`, `install`, `monitor`, or `reboot`) argument followed by the `--help` option you will get a different help text providing more detail on the options available within that subcommand:

```
configure-sga.py subcommand --help
```

3. To confirm the current configuration of the appliance node, enter the following where `SGA-install-ip` is any one of the IP addresses for the appliance node:

```
configure-sga.py configure SGA-INSTALL-IP
```

The results show current IP information for the appliance, including the IP address of the primary Admin Node and information about the Admin, Grid, and Client Networks.

```

Connecting to +https://10.224.2.30:8443+ (Checking version and
connectivity.)
2021/02/25 16:25:11: Performing GET on /api/versions... Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/system-info... Received
200
2021/02/25 16:25:11: Performing GET on /api/v2/admin-connection...
Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/link-config... Received
200
2021/02/25 16:25:11: Performing GET on /api/v2/networks... Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/system-config... Received
200

```

#### StorageGRID Appliance

```

Name:          LAB-SGA-2-30
Node type:     storage

```

#### StorageGRID primary Admin Node

```

IP:           172.16.1.170
State:        unknown
Message:      Initializing...
Version:      Unknown

```

#### Network Link Configuration

##### Link Status

Link	State	Speed (Gbps)
----	-----	-----
1	Up	10
2	Up	10
3	Up	10
4	Up	10
5	Up	1
6	Down	N/A

##### Link Settings

```

Port bond mode:    FIXED
Link speed:        10GBE

Grid Network:      ENABLED
  Bonding mode:    active-backup
  VLAN:            novlan
  MAC Addresses:   00:a0:98:59:8e:8a  00:a0:98:59:8e:82

Admin Network:     ENABLED
  Bonding mode:    no-bond
  MAC Addresses:   00:80:e5:29:70:f4

```

```
Client Network:      ENABLED
    Bonding mode:    active-backup
    VLAN:            novlan
    MAC Addresses:   00:a0:98:59:8e:89  00:a0:98:59:8e:81

Grid Network
  CIDR:             172.16.2.30/21 (Static)
  MAC:              00:A0:98:59:8E:8A
  Gateway:          172.16.0.1
  Subnets:         172.17.0.0/21
                   172.18.0.0/21
                   192.168.0.0/21
  MTU:              1500

Admin Network
  CIDR:             10.224.2.30/21 (Static)
  MAC:              00:80:E5:29:70:F4
  Gateway:          10.224.0.1
  Subnets:         10.0.0.0/8
                   172.19.0.0/16
                   172.21.0.0/16
  MTU:              1500

Client Network
  CIDR:             47.47.2.30/21 (Static)
  MAC:              00:A0:98:59:8E:89
  Gateway:          47.47.0.1
  MTU:              2000

#####
#####  If you are satisfied with this configuration,      #####
##### execute the script with the "install" sub-command. #####
#####
```

4. If you need to change any of the values in the current configuration, use the `configure` subcommand to update them. For example, if you want to change the IP address that the appliance uses for connection to the primary Admin Node to `172.16.2.99`, enter the following:

```
configure-sga.py configure --admin-ip 172.16.2.99 SGA-INSTALL-IP
```

5. If you want to back up the appliance configuration to a JSON file, use the `advanced` and `backup-file` subcommands. For example, if you want to back up the configuration of an appliance with IP address `SGA-INSTALL-IP` to a file named `appliance-SG1000.json`, enter the following:

```
configure-sga.py advanced --backup-file appliance-SG1000.json SGA-INSTALL-IP
```

The JSON file containing the configuration information is written to the same directory you executed the script from.



Check that the top-level node name in the generated JSON file matches the appliance name. Do not make any changes to this file unless you are an experienced user and have a thorough understanding of StorageGRID APIs.

- When you are satisfied with the appliance configuration, use the `install` and `monitor` subcommands to install the appliance:

```
configure-sga.py install --monitor SGA-INSTALL-IP
```

- If you want to reboot the appliance, enter the following:

```
configure-sga.py reboot SGA-INSTALL-IP
```

## Automating the configuration of StorageGRID

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

### What you'll need

- 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>	Sample 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 sample 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

- 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`.

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

For example:

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

## After you finish

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, you need to extract 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.
```

## Overview of installation REST APIs

StorageGRID provides two REST APIs for performing installation tasks: the StorageGRID Installation API and the StorageGRID Appliance Installer API.

Both APIs use 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 (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs 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 in the event that 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.

### StorageGRID Appliance Installer API

The StorageGRID Appliance Installer API can be accessed over HTTPS from `Controller_IP:8443`.

To access the API documentation, go to the StorageGRID Appliance Installer on the appliance and select **Help > API Docs** from the menu bar.

The StorageGRID Appliance Installer API includes the following sections:

- **clone** — Operations to configure and control node cloning.
- **encryption** — Operations to manage encryption and view encryption status.
- **hardware configuration** — Operations to configure system settings on attached hardware.
- **installation** — Operations for starting the appliance installation and for monitoring installation status.
- **networking** — Operations related to the Grid, Admin, and Client Network configuration for a StorageGRID appliance and appliance port settings.
- **setup** — Operations to help with initial appliance installation setup including requests to get information about the system and update the primary Admin Node IP.
- **support** — Operations for rebooting the controller and getting logs.
- **upgrade** — Operations related to upgrading appliance firmware.
- **uploadsg** — Operations for uploading StorageGRID installation files.

## Troubleshooting the hardware installation

If you encounter issues during the installation, you might find it helpful to review troubleshooting information related to hardware setup and connectivity issues.

### Related information

[Hardware setup appears to hang](#)

## Viewing boot-up codes for the SG6000-CN controller

When you apply power to the appliance, the BMC logs a series of boot-up codes for the SG6000-CN controller. You can view these codes in several ways.

### What you'll need

- You know how to access the BMC dashboard.
- If you want to use a kernel-based virtual machine (KVM), you have experience deploying and using KVM applications.
- If you want to use serial-over-LAN (SOL), you have experience using IPMI SOL console applications.

### Steps

1. Select one of the following methods for viewing the boot-up codes for the appliance controller, and gather the required equipment.

Method	Required equipment
VGA console	<ul style="list-style-type: none"><li>• VGA-capable monitor</li><li>• VGA cable</li></ul>
KVM	<ul style="list-style-type: none"><li>• KVM application</li><li>• RJ-45 cable</li></ul>
Serial port	<ul style="list-style-type: none"><li>• DB-9 serial cable</li><li>• Virtual serial terminal</li></ul>
SOL	<ul style="list-style-type: none"><li>• Virtual serial terminal</li></ul>

2. If you are using a VGA console, perform these steps:
  - a. Connect a VGA-capable monitor to the VGA port on the back of the appliance.
  - b. View the codes displayed on the monitor.
3. If you are using BMC KVM, perform these steps:
  - a. Connect to the BMC management port and log into the BMC web interface.
  - b. Select **Remote Control**.
  - c. Launch the KVM.
  - d. View the codes on the virtual monitor.
4. If you are using a serial port and terminal, perform these steps:
  - a. Connect to the DB-9 serial port on the back of the appliance.
  - b. Use settings `115200 8-N-1`.
  - c. View the codes printed over the serial terminal.
5. If you are using SOL, perform these steps:

a. Connect to the IPMI SOL using the BMC IP address and login credentials.

```
ipmitool -I lanplus -H 10.224.3.91 -U root -P calvin sol activate
```

b. View the codes on the virtual serial terminal.

6. Use the table to look up the codes for your appliance.

Code	Indicates
HI	The master boot script has started.
HP	The system is checking to see if the network interface card (NIC) firmware needs to be updated.
RB	The system is rebooting after applying firmware updates.
FP	The hardware subsystem firmware update checks have been completed. Inter-controller communication services are starting.
HE	<p>For an appliance Storage Node only:</p> <p>The system is awaiting connectivity with the storage controllers and synchronizing with the SANtricity operating system.</p> <p><b>Note:</b> If the boot-up procedure does not progress past this stage, perform these steps:</p> <ol style="list-style-type: none"><li>Confirm that the four interconnect cables between the SG6000-CN controller and the two storage controllers are securely connected.</li><li>As required, replace one or more of the cables, and try again.</li><li>If this does not resolve the issue, contact technical support.</li></ol>
HC	The system is checking for existing StorageGRID installation data.
HO	The StorageGRID Appliance Installer is running.
HA	StorageGRID is running.

### Viewing error codes for the SG6000-CN controller

If a hardware error occurs when the SG6000-CN controller is booting up, the BMC logs an error code. As required, you can view these error codes using the BMC interface, and

then work with technical support to resolve the issue.

### What you'll need

- You know how to access the BMC dashboard.

### Steps

1. From the BMC dashboard, select **BIOS POST Code**.
2. Review the information displayed for Current Code and the Previous Code.

If any of the following error codes are shown, work with technical support to resolve the issue.

Code	Indicates
0x0E	Microcode not found
0x0F	Microcode not loaded
0x50	Memory initialization error. Invalid memory type or incompatible memory speed.
0x51	Memory initialization error. SPD reading has failed.
0x52	Memory initialization error. Invalid memory size or memory modules do not match.
0x53	Memory initialization error. No usable memory detected.
0x54	Unspecified memory initialization error
0x55	Memory not installed
0x56	Invalid CPU type or speed
0x57	CPU mismatch
0x58	CPU self-test failed, or possible CPU cache error
0x59	CPU micro-code is not found, or micro-code update failed
0x5A	Internal CPU error
0x5B	Reset PPI is not available
0x5C	PEI phase BMC self-test failure

<b>Code</b>	<b>Indicates</b>
0xD0	CPU initialization error
0xD1	North bridge initialization error
0xD2	South bridge initialization error
0xD3	Some architectural protocols are not available
0xD4	PCI resource allocation error. Out of resources.
0xD5	No space for legacy option ROM
0xD6	No console output devices are found
0xD7	No console input devices are found
0xD8	Invalid password
0xD9	Error loading boot option (LoadImage returned error)
0xDA	Boot option failed (StartImage returned error)
0xDB	Flash update failed
0xDC	Reset protocol is not available
0xDD	DXE phase BMC self-test failure
0xE8	MRC: ERR_NO_MEMORY
0xE9	MRC: ERR_LT_LOCK
0xEA	MRC: ERR_DDR_INIT
0xEB	MRC: ERR_MEM_TEST
0xEC	MRC: ERR_VENDOR_SPECIFIC
0xED	MRC: ERR_DIMM_COMPAT
0xEE	MRC: ERR_MRC_COMPATIBILITY

Code	Indicates
0xEF	MRC: ERR_MRC_STRUCT
0xF0	MRC: ERR_SET_VDD
0xF1	MRC: ERR_IOT_MEM_BUFFER
0xF2	MRC: ERR_RC_INTERNAL
0xF3	MRC: ERR_INVALID_REG_ACCESS
0xF4	MRC: ERR_SET_MC_FREQ
0xF5	MRC: ERR_READ_MC_FREQ
0x70	MRC: ERR_DIMM_CHANNEL
0x74	MRC: ERR_BIST_CHECK
0xF6	MRC: ERR_SMBUS
0xF7	MRC: ERR_PCU
0xF8	MRC: ERR_NGN
0xF9	MRC: ERR_INTERLEAVE_FAILURE

### Hardware setup appears to hang

The StorageGRID Appliance Installer might not be available if hardware faults or cabling errors prevent the storage controllers or the SG6000-CN controller from completing their boot-up processing.

#### Steps

1. For the storage controllers, watch the codes on the seven-segment displays.

While the hardware is initializing during power up, the two seven-segment displays show a sequence of codes. When the hardware boots successfully, both seven-segment displays show 99.

2. Review the LEDs on the SG6000-CN controller and the boot-up and error codes displayed in the BMC.
3. If you need help resolving an issue, contact technical support.

#### Related information

[Viewing boot-up status codes for the SG6000 storage controllers](#)

[E5700 and E2800 System Monitoring Guide](#)

[Viewing status indicators and buttons on the SG6000-CN controller](#)

[Viewing boot-up codes for the SG6000-CN controller](#)

[Viewing error codes for the SG6000-CN controller](#)

## Troubleshooting connection issues

If you encounter connection issues during the StorageGRID appliance installation, you should perform the corrective action steps listed.

### Unable to connect to the appliance

If you cannot connect to the appliance, there might be a network issue, or the hardware installation might not have been completed successfully.

### Steps

1. If you are unable to connect to SANtricity System Manager:
  - a. Try to ping the appliance using the IP address for either storage controller on the management network for SANtricity System Manager:  
**ping *Storage\_Controller\_IP***
  - b. If you receive no response from the ping, confirm you are using the correct IP address.  
  
Use the IP address for management port 1 on either storage controller.
  - c. If the IP address is correct, check appliance cabling and the network setup.  
  
If that does not resolve the issue, contact technical support.
  - d. If the ping was successful, open a web browser.
  - e. Enter the URL for SANtricity System Manager:  
**https://*Storage\_Controller\_IP***  
  
The log in page for SANtricity System Manager appears.
2. If you are unable to connect to the SG6000-CN controller:
  - a. Try to ping the appliance using the IP address for the SG6000-CN controller:  
**ping *SG6000-CN\_Controller\_IP***
  - b. If you receive no response from the ping, confirm you are using the correct IP address.  
  
You can use the IP address of the appliance on the Grid Network, the Admin Network, or the Client Network.
  - c. If the IP address is correct, check appliance cabling, SFP transceivers, and the network setup.  
  
If that does not resolve the issue, contact technical support.
  - d. If the ping was successful, open a web browser.
  - e. Enter the URL for the StorageGRID Appliance Installer:  
**https://*SG6000-CN\_Controller\_IP*:8443**

The Home page appears.

### Expansion shelves do not appear in Appliance Installer

If you have installed expansion shelves for the SG6060 and they do not appear in the StorageGRID Appliance Installer, you should verify that the shelves have been completely installed and powered on.

#### About this task

You can verify that the expansion shelves are connected to the appliance by viewing the following information in the StorageGRID Appliance Installer:

- The **Home** page contains a message about expansion shelves.



**i** The storage system contains 2 expansion shelves.

- The **Advanced > RAID Mode** page indicates by number of drives whether or not the appliance includes expansion shelves. For example, in the following screen shot, two SSDs and 178 HDDs are shown. An SG6060 with two expansion shelves contains 180 total drives.

### Configure RAID Mode

This appliance contains the following drives.

Type	Size	Number of drives
SSD	800 GB	2
HDD	11.8 TB	178

If the StorageGRID Appliance Installer pages do not indicate that expansion shelves are present, follow this procedure.

#### Steps

1. Verify that all required cables have been firmly connected.
2. Verify that you have powered on the expansion shelves.
3. If you need help resolving an issue, contact technical support.

#### Related information

[SG6060: Cabling the optional expansion shelves](#)

[Connecting power cords and applying power \(SG6000\)](#)

### Rebooting the SG6000-CN controller while the StorageGRID Appliance Installer is running


You might need to reboot the SG6000-CN controller while the StorageGRID Appliance Installer is running. For example, you might need to reboot the controller if the installation fails.

#### About this task

This procedure only applies when the SG6000-CN controller is running the StorageGRID Appliance Installer. Once the installation is completed, this step no longer works because the StorageGRID Appliance Installer is no longer available.

## Steps

1. From the StorageGRID Appliance Installer, click **Advanced** > **Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
  - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



The SG6000-CN controller is rebooted.

## Maintaining the SG6000 appliance

You might need to perform maintenance procedures on the SG6000 appliance. The procedures in this section assume that the appliance has already been deployed as a Storage Node in a StorageGRID system.

### Steps

- [Placing an appliance into maintenance mode](#)
- [Upgrading SANtricity OS on the storage controllers](#)
- [Upgrading drive firmware using SANtricity System Manager](#)
- [Adding an expansion shelf to a deployed SG6060](#)
- [Turning the controller identify LED on and off](#)
- [Locating the controller in a data center](#)
- [Replacing a storage controller](#)
- [Replacing hardware components in the storage controller shelf](#)
- [Replacing hardware components in the optional 60-drive expansion shelf](#)
- [Shutting down the SG6000-CN controller](#)

- Powering on the SG6000-CN controller and verifying operation
- Replacing the SG6000-CN controller
- Replacing a power supply in the SG6000-CN controller
- Removing the SG6000-CN controller from a cabinet or rack
- Reinstalling the SG6000-CN controller into a cabinet or rack
- Removing the SG6000-CN controller cover
- Reinstalling the SG6000-CN controller cover
- Replacing the Fibre Channel HBA in the SG6000-CN controller
- Changing the link configuration of the SG6000-CN controller
- Changing the MTU setting
- Checking the DNS server configuration
- Monitoring node encryption in maintenance mode

### Placing an appliance into maintenance mode

You must place the appliance into maintenance mode before performing specific maintenance procedures.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission. For details, see the instructions for administering StorageGRID.

#### About this task

Placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.



The password and host key for a StorageGRID appliance in maintenance mode remain the same as they were when the appliance was in service.

#### Steps

1. From the Grid Manager, select **Nodes**.
2. From the tree view of the Nodes page, select the appliance Storage Node.
3. Select **Tasks**.

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

Places the appliance's compute controller into maintenance mode.

Maintenance Mode

#### 4. Select **Maintenance Mode**.

A confirmation dialog box appears.

### Enter Maintenance Mode on SGA-106-15

You must place the appliance's compute controller into maintenance mode to perform certain maintenance procedures on the appliance.

Attention: All StorageGRID services on this node will be shut down. Wait a few minutes for the node to reboot into maintenance mode.

If you are ready to start, enter the provisioning passphrase and click OK.

Provisioning Passphrase

Cancel

OK

#### 5. Enter the provisioning passphrase, and select **OK**.

A progress bar and a series of messages, including including "Request Sent," "Stopping StorageGRID," and "Rebooting," indicate that the appliance is completing the steps for entering maintenance mode.

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

**Attention:** Your request has been sent, but the appliance might take 10-15 minutes to enter maintenance mode. Do not perform maintenance procedures until this tab indicates maintenance mode is ready, or data could become corrupted.



Request Sent

When the appliance is in maintenance mode, a confirmation message lists the URLs you can use to access the StorageGRID Appliance Installer.

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

This node is currently in maintenance mode. Navigate to one of the URLs listed below and perform any necessary maintenance procedures.

- <https://172.16.2.106:8443>
- <https://10.224.2.106:8443>
- <https://47.47.2.106:8443>
- <https://169.254.0.1:8443>

When you are done with any required maintenance procedures, you must exit maintenance mode by clicking Reboot Controller from the StorageGRID Appliance Installer.

6. To access the StorageGRID Appliance Installer, browse to any of the URLs displayed.

If possible, use the URL containing the IP address of the appliance's Admin Network port.




Accessing <https://169.254.0.1:8443> requires a direct connection to the local management port.


7. From the StorageGRID Appliance Installer, confirm that the appliance is in maintenance mode.

⚠ This node is in maintenance mode. Perform any required maintenance procedures. If you want to exit maintenance mode manually to resume normal operation, go to **Advanced > Reboot Controller** to **reboot** the controller.

8. Perform any required maintenance tasks.
9. After completing maintenance tasks, exit maintenance mode and resume normal node operation. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Upgrading SANtricity OS on the storage controllers

To ensure optimal functioning of the storage controller, you must upgrade to the latest maintenance release of the SANtricity OS that is qualified for your StorageGRID appliance. Consult the NetApp Interoperability Matrix Tool (IMT) to determine which version you should be using. If you need assistance, contact technical support.

Use one of the following procedures based on the version of SANtricity OS currently installed:

- If the storage controller is using SANtricity OS 08.42.20.00 (11.42) or newer, use the Grid Manager to perform the upgrade.

#### [Upgrading SANtricity OS on the storage controllers using the Grid Manager](#)

- If the storage controller is using a SANtricity OS version older than 08.42.20.00 (11.42), use maintenance mode to perform the upgrade.

#### [Upgrading SANtricity OS on the storage controllers using maintenance mode](#)



When upgrading the SANtricity OS for your storage appliance, you must follow the instructions in the StorageGRID documentation. If you use any other instructions, your appliance could become inoperable.

#### **Related information**

[NetApp Interoperability Matrix Tool](#)

[NetApp Downloads: SANtricity OS](#)

[Monitor & troubleshoot](#)

#### **Upgrading SANtricity OS on the storage controllers using the Grid Manager**

For storage controllers currently using SANtricity OS 08.42.20.00 (11.42) or newer, you must use the Grid Manager to apply an upgrade.

#### **What you'll need**

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- You must have the Maintenance permission.
- You must be signed in to the Grid Manager using a supported browser.
- You must have the provisioning passphrase.
- You must have access to the NetApp downloads page for SANtricity OS.

#### **About this task**

You cannot perform other software updates (StorageGRID software upgrade or a hotfix) until you have completed the SANtricity OS upgrade process. If you attempt to start a hotfix or a StorageGRID software upgrade before the SANtricity OS upgrade process has finished, you are redirected to the SANtricity OS upgrade page.

The procedure will not be complete until the SANtricity OS upgrade has been successfully applied to all applicable nodes. It might take more than 30 minutes to load the SANtricity OS on each node and up to 90 minutes to reboot each StorageGRID storage appliance.



The following steps are only applicable when you are using the Grid Manager to perform the upgrade. The storage controllers in the SG6000 series appliances cannot be upgraded using the Grid Manager when the controllers are using SANtricity OS older than 08.42.20.00 (11.42).



This procedure will automatically upgrade the NVSRAM to the most recent version associated with the SANtricity OS upgrade. You do not need to apply a separate NVSRAM upgrade file.

## Steps

1. From a service laptop, download the new SANtricity OS Software file from the NetApp support site.

Be sure to choose the correct SANtricity OS version for the storage controllers in your appliance. The SG6060 uses the E2800 controller, and the SGF6024 uses the EF570 controller.

### [NetApp Downloads: SANtricity OS](#)

2. Sign in to the Grid Manager using a supported browser.
3. Select **Maintenance**. Then, in the System section of the menu, select **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.

- To perform a major version upgrade of StorageGRID, see the [instructions for upgrading StorageGRID](#), and then select **StorageGRID Upgrade**.
- To apply a hotfix to all nodes in your system, see "Hotfix procedure" in the [recovery and maintenance instructions](#), and then select **StorageGRID Hotfix**.
- To upgrade SANtricity OS software on a storage controller, see "Upgrading SANtricity OS Software on the storage controllers" in the installation and maintenance instructions for your storage appliance, and then select **SANtricity OS**:

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)



4. Click **SANtricity OS**.

The SANtricity OS page appears.

## SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

### SANtricity OS Upgrade File

---

SANtricity OS Upgrade File



Browse

### Passphrase

---

Provisioning Passphrase



Start

5. Select the SANtricity OS upgrade file you downloaded from the NetApp support site.
  - a. Click **Browse**.
  - b. Locate and select the file.
  - c. Click **Open**.

The file is uploaded and validated. When the validation process is done, the file name is shown in the Details field.



Do not change the file name since it is part of the verification process.

## SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

### SANtricity OS Upgrade File

SANtricity OS Upgrade File



Browse



RC\_XXXXXXXXXX\_XXXXXX.dlp

Details



RC\_XXXXXXXXXX\_XXXXXX.dlp

### Passphrase

Provisioning Passphrase



Start

6. Enter the provisioning passphrase.

The **Start** button is enabled.

## SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

### SANtricity OS Upgrade File

SANtricity OS Upgrade File



Browse



RC\_XXXXXXXXXX\_XXXXXX.dlp

Details



RC\_XXXXXXXXXX\_XXXXXX.dlp

### Passphrase

Provisioning Passphrase



Start

7. Click **Start**.

A warning box appears stating that your browser's connection might be lost temporarily as services on nodes that are upgraded are restarted.

**Warning**

**Nodes can disconnect and services might be affected**

The node will be automatically rebooted at the end of upgrade and services will be affected. Are you sure you want to start the SANtricity OS upgrade?

8. Click **OK** to stage the SANtricity OS upgrade file to the primary Admin Node.

When the SANtricity OS upgrade starts:

a. The health check is run. This process checks that no nodes have the status of Needs Attention.

 If any errors are reported, resolve them and click **Start** again.

b. The SANtricity OS Upgrade Progress table appears. This table shows all Storage Nodes in your grid and the current stage of the upgrade for each node.

 The table shows all Storage Nodes, including software-based Storage Nodes. You must approve the upgrade for all Storage Nodes, even though a SANtricity OS upgrade has no effect on software-based Storage Nodes. The upgrade message returned for software-based Storage Nodes is "SANtricity OS upgrade is not applicable to this node."

#### SANtricity OS Upgrade Progress

 **Storage Nodes** - 0 out of 4 completed

Search  

Site	Name	Progress	Stage	Details	Action
RTP Lab 1	DT-10-224-1-181-S1		Waiting for you to approve		<input type="button" value="Approve"/>
RTP Lab 1	DT-10-224-1-182-S2		Waiting for you to approve		<input type="button" value="Approve"/>
RTP Lab 1	DT-10-224-1-183-S3		Waiting for you to approve		<input type="button" value="Approve"/>
RTP Lab 1	NetApp-SGA-Lab2-002-024		Waiting for you to approve		<input type="button" value="Approve"/>

9. Optionally, sort the list of nodes 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.

You can scroll through the list of nodes by using the left and right arrows at the bottom right corner of the section.

10. Approve the grid nodes you are ready to add to the upgrade queue. Approved nodes of the same type are upgraded one at a time.



Do not approve the SANtricity OS upgrade for an appliance storage node unless you are sure the node is ready to be stopped and rebooted. When the SANtricity OS upgrade is approved on a node, the services on that node are stopped. Later, when the node is upgraded, the appliance node is rebooted. These operations might cause service interruptions for clients that are communicating with the node.

- Click either of the **Approve All** buttons to add all Storage Nodes to the SANtricity OS upgrade queue.



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).

- Click one or more **Approve** buttons to add one or more nodes to the SANtricity OS upgrade queue.



You can delay applying a SANtricity OS upgrade to a node, but the SANtricity OS upgrade process will not be complete until you approve the SANtricity OS upgrade on all the listed Storage Nodes.

After you click **Approve**, the upgrade process determines if the node can be upgraded. If a node can be upgraded, it is added to the upgrade queue. +

For some nodes, the selected upgrade file is intentionally not applied and you can complete the upgrade process without upgrading these specific nodes. For nodes intentionally not upgraded, the process will show stage of Complete with one of the following messages in the Details column:

- Storage Node was already upgraded.
- SANtricity OS upgrade is not applicable to this node.
- SANtricity OS file is not compatible with this node.

The message “SANtricity OS upgrade is not applicable to this node” indicates that the node does not have a storage controller that can be managed by the StorageGRID system. This message will appear for non-appliance Storage Nodes. You can complete the SANtricity OS upgrade process without upgrading the node displaying this message.

The message “SANtricity OS file is not compatible with this node” indicates that the node requires a SANtricity OS file different than the one the process is attempting to install. After you complete the current SANtricity OS upgrade, download the SANtricity OS appropriate for the node and repeat the upgrade process.

11. If you need to remove a node or all nodes from the SANtricity OS upgrade queue, click **Remove** or **Remove All**.

As shown in the example, when the stage progresses beyond Queued, the **Remove** button is hidden and you can no longer remove the node from the SANtricity OS upgrade process.

Storage Nodes - 1 out of 9 completed Approve All Remove All

Search

Site	Name	Progress	Stage	Details	Action
Raleigh	RAL-S1-101-196	<div style="width: 0%;"></div>	Queued		Remove
Raleigh	RAL-S2-101-197	<div style="width: 100%; background-color: green;"></div>	Complete		
Raleigh	RAL-S3-101-198	<div style="width: 0%;"></div>	Queued		Remove
Sunnyvale	SVL-S1-101-199	<div style="width: 0%;"></div>	Queued		Remove
Sunnyvale	SVL-S2-101-93	<div style="width: 0%;"></div>	Waiting for you to approve		Approve
Sunnyvale	SVL-S3-101-94	<div style="width: 0%;"></div>	Waiting for you to approve		Approve
Vancouver	VTC-S1-101-193	<div style="width: 0%;"></div>	Waiting for you to approve		Approve
Vancouver	VTC-S2-101-194	<div style="width: 0%;"></div>	Waiting for you to approve		Approve
Vancouver	VTC-S3-101-195	<div style="width: 0%;"></div>	Waiting for you to approve		Approve

12. Wait while the SANtricity OS upgrade is applied to each approved grid node.



If any node shows a stage of Error while the SANtricity OS upgrade is being applied, the upgrade has failed for that node. The appliance might need to be placed in maintenance mode to recover from the failure. Contact technical support before continuing.

If the firmware on the node is too old to be upgraded with the Grid Manager, the node shows a stage of Error with the details: "You must use maintenance mode to upgrade SANtricity OS on this node. See the installation and maintenance instructions for your appliance. After the upgrade, you can use this utility for future upgrades." To resolve the error, do the following:

- a. Use maintenance mode to upgrade SANtricity OS on the node that shows a stage of Error.
- b. Use the Grid Manager to re-start and complete the SANtricity OS upgrade.

When the SANtricity OS upgrade is complete on all approved nodes, the SANtricity OS Upgrade Progress table closes and a green banner shows the date and time the SANtricity OS upgrade was completed.

SANtricity OS upgrade completed at 2020-04-07 13:26:02 EDT.

**SANtricity OS Upgrade File**

SANtricity OS Upgrade File

**Passphrase**

Provisioning Passphrase

13. Repeat this upgrade procedure for any nodes with a stage of Complete that require a different SANtricity OS upgrade file.



For any nodes with a status of Needs Attention, use maintenance mode to perform the upgrade.

### Related information

[NetApp Interoperability Matrix Tool](#)

[Upgrading SANtricity OS on the storage controllers using maintenance mode](#)

### Upgrading SANtricity OS on the storage controllers using maintenance mode

For storage controllers currently using SANtricity OS older than 08.42.20.00 (11.42), you must use the maintenance mode procedure to apply an upgrade.

### What you'll need

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- If the StorageGRID appliance is running in a StorageGRID system, the SG6000-CN controller has been placed in maintenance mode.



Maintenance mode interrupts the connection to the storage controller.

[Placing an appliance into maintenance mode](#)

### About this task

Do not upgrade the SANtricity OS or NVSRAM in the E-Series controller on more than one StorageGRID appliance at a time.



Upgrading more than one StorageGRID appliance at a time might cause data unavailability, depending on your deployment model and ILM policies.

### Steps

1. From a service laptop, access SANtricity System Manager and sign in.
2. Download the new SANtricity OS Software file and NVSRAM file to the management client.



The NVSRAM is specific to the StorageGRID appliance. Do not use the standard NVSRAM download.


3. Follow the instructions in the *Upgrading SANtricity OS* guide or the SANtricity System Manager online help to upgrade the firmware and NVSRAM.




Activate the upgrade files immediately. Do not defer activation.

4. Once the upgrade operation has completed, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
  - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in

maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[NetApp Interoperability Matrix Tool](#)

[Upgrading SANtricity OS on the storage controllers using the Grid Manager](#)

## Upgrading drive firmware using SANtricity System Manager

You upgrade your drive firmware to make sure you have all the latest features and bug fixes.

### What you'll need

- The storage appliance has an Optimal status.
- All drives have an Optimal status.
- You have the latest version of SANtricity System Manager installed that is compatible with your StorageGRID version.
- You have placed the StorageGRID appliance in maintenance mode.

## Placing an appliance into maintenance mode



Maintenance mode interrupts the connection to the storage controller, stopping all I/O activity and placing all drives offline.



Do not upgrade the drive firmware on more than one StorageGRID appliance at a time. Doing so might cause data unavailability, depending on your deployment model and ILM policies.

### Steps

1. Access SANtricity System Manager using one of these methods:
  - Use the StorageGRID Appliance Installer and select **Advanced > SANtricity System Manager**
  - Use the Grid Manager and select **Nodes > appliance Storage Node > SANtricity System Manager**



If these options are not available or the SANtricity System Manager login page does not appear, access SANtricity System Manager by browsing to the storage controller IP:  
**`https://Storage_Controller_IP`**

2. Enter the SANtricity System Manager administrator username and password, if required.
3. Verify the drive firmware version currently installed in the storage appliance:
  - a. From SANtricity System Manager, select **Support > Upgrade Center**.
  - b. Under Drive Firmware upgrade, select **Begin Upgrade**.

The Upgrade Drive Firmware displays the drive firmware files currently installed.

- c. Note the current drive firmware revisions and drive identifiers in the Current Drive Firmware column.

Current Drive Firmware	Associated Drives
MS02, KPM51VUG800G	View drives

In this example:

- The drive firmware revision is **MS02**.
- The drive identifier is **KPM51VUG800G**.

Select **View drives** in the Associated Drives column to display where these drives are installed in your storage appliance.

d. Close the Upgrade Drive Firmware window.

4. Download and prepare the available drive firmware upgrade:

a. Under Drive Firmware upgrade, select **NetApp Support**.

b. On the NetApp Support web site, select the **Downloads** tab, and then select **E-Series Disk Drive Firmware**.

The E-Series Disk Firmware page displays.

c. Search for each **Drive Identifier** installed in your storage appliance and verify that each drive identifier has the latest firmware revision.

- If the firmware revision is not a link, this drive identifier has the latest firmware revision.
- If one or more drive part numbers are listed for a drive identifier, a firmware upgrade is available for these drives. You can select any link to download the firmware file.

PRODUCTS ▾ SYSTEMS ▾ DOCS & KNOWLEDGEBASE ▾ COMMUNITY ▾ DOWNLOADS ▾ TOOLS ▾ CASES ▾ PARTS ▾

Downloads > Firmware > E-Series Disk Firmware

## E-Series Disk Firmware

[Download all current E-Series Disk Firmware](#)

Drive Part Number	Descriptions	Drive Identifier	Firmware Rev. (Download)	Notes and Config Info	Release Date
<input type="text" value="Drive Part Number"/>	<input type="text" value="Descriptions"/>	<input type="text" value="KPM51VUG800G"/>	<input type="text" value="Firmware Rev. (Download)"/>		
<a href="#">E-X4041C</a>	SSD, 800GB, SAS, PI	KPM51VUG800G	<a href="#">MS03</a>	<a href="#">MS02 Fixes Bug 1194908</a> <a href="#">MS03 Fixes Bug 1334862</a>	04-Sep-2020

d. If a later firmware revision is listed, select the link in the Firmware Rev. (Download) column to download a .zip archive containing the firmware file.

e. Extract (unzip) the drive firmware archive files you downloaded from the Support site.

5. Install the drive firmware upgrade:

a. From SANtricity System Manager, under Drive Firmware upgrade, select **Begin Upgrade**.

b. Select **Browse**, and select the new drive firmware files that you downloaded from the Support site.

Drive firmware files have a filename similar to

D\_HUC101212CSS600\_30602291\_MS01\_2800\_0002.dlp.

You can select up to four drive firmware files, one at a time. If more than one drive firmware file is compatible with the same drive, you get a file conflict error. Decide which drive firmware file you want to use for the upgrade and remove the other one.

c. Select **Next**.

**Select Drives** lists the drives that you can upgrade with the selected firmware files.

Only drives that are compatible appear.

The selected firmware for the drive appears in **Proposed Firmware**. If you must change this firmware, select **Back**.

- d. Select **Offline (parallel)** upgrade.

You can use the offline upgrade method because the appliance is in maintenance mode, where I/O activity is stopped for all drives and all volumes.

- e. In the first column of the table, select the drive or drives you want to upgrade.

The best practice is to upgrade all drives of the same model to the same firmware revision.

- f. Select **Start**, and confirm that you want to perform the upgrade.

If you need to stop the upgrade, select **Stop**. Any firmware downloads currently in progress complete. Any firmware downloads that have not started are canceled.



Stopping the drive firmware upgrade might result in data loss or unavailable drives.

- g. (Optional) To see a list of what was upgraded, select **Save Log**.

The log file is saved in the downloads folder for your browser with the name `latest-upgrade-log-timestamp.txt`.

If any of the following errors occur during the upgrade procedure, take the appropriate recommended action.

▪ **Failed assigned drives**

One reason for the failure might be that the drive does not have the appropriate signature. Make sure that the affected drive is an authorized drive. Contact technical support for more information.

When replacing a drive, make sure that the replacement drive has a capacity equal to or greater than the failed drive you are replacing.

You can replace the failed drive while the storage array is receiving I/O.

▪ **Check storage array**

- Make sure that an IP address has been assigned to each controller.
- Make sure that all cables connected to the controller are not damaged.
- Make sure that all cables are tightly connected.

▪ **Integrated hot spare drives**

This error condition must be corrected before you can upgrade the firmware.

▪ **Incomplete volume groups**

If one or more volume groups or disk pools are incomplete, you must correct this error condition before you can upgrade the firmware.

▪ **Exclusive operations (other than background media/parity scan) currently running on any**

## volume groups

If one or more exclusive operations are in progress, the operations must complete before the firmware can be upgraded. Use System Manager to monitor the progress of the operations.

- **Missing volumes**

You must correct the missing volume condition before the firmware can be upgraded.

- **Either controller in a state other than Optimal**

One of the storage array controllers needs attention. This condition must be corrected before the firmware can be upgraded.

- **Mismatched Storage Partition information between Controller Object Graphs**

An error occurred while validating the data on the controllers. Contact technical support to resolve this issue.

- **SPM Verify Database Controller check fails**

A storage partitions mapping database error occurred on a controller. Contact technical support to resolve this issue.

- **Configuration Database Validation (If supported by the storage array's controller version)**

A configuration database error occurred on a controller. Contact technical support to resolve this issue.

- **MEL Related Checks**

Contact technical support to resolve this issue.

- **More than 10 DDE Informational or Critical MEL events were reported in the last 7 days**

Contact technical support to resolve this issue.

- **More than 2 Page 2C Critical MEL Events were reported in the last 7 days**

Contact technical support to resolve this issue.

- **More than 2 Degraded Drive Channel Critical MEL events were reported in the last 7 days**

Contact technical support to resolve this issue.

- **More than 4 critical MEL entries in the last 7 days**

Contact technical support to resolve this issue.


6. Once the upgrade operation has completed, reboot the appliance. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in

maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Upgrading SANtricity OS on the storage controllers](#)

## Adding an expansion shelf to a deployed SG6060

To increase storage capacity, you can add one or two expansion shelves to an SG6060 that is deployed in a StorageGRID system.

### What you'll need

- You must have the provisioning passphrase.
- You must be running StorageGRID 11.4 or later.
- You have the expansion shelf and two SAS cables for each expansion shelf.
- You have physically located the storage appliance where you are adding the expansion shelf in the data center.

[Locating the controller in a data center](#)

## About this task

To add an expansion shelf, you perform these high-level steps:

- Install the hardware in the cabinet or rack.
- Place the SG6060 into maintenance mode.
- Connect the expansion shelf to the E2860 controller shelf or to another expansion shelf.
- Start the expansion using the StorageGRID Appliance Installer
- Wait until the new volumes are configured.

Completing the procedure for one or two expansion shelves should take one hour or less per appliance node. To minimize downtime, the following steps instruct you to install the new expansion shelves and drives before placing the SG6060 into maintenance mode. The remaining steps should take approximately 20 to 30 minutes per appliance node.

## Steps

1. Follow the instructions for installing 60-drive shelves into a cabinet or rack.

[SG6060: Installing 60-drive shelves into a cabinet or rack](#)

2. Follow the instructions for installing the drives.


[SG6060: Installing the drives](#)

3. From the Grid Manager, place the SG6000-CN controller into maintenance mode.

[Placing an appliance into maintenance mode](#)

4. Connect each expansion shelf to the E2860 controller shelf as shown in the diagram.

This drawing shows two expansion shelves. If you have only one, connect IOM A to controller A and connect IOM B to controller B.



	Description
1	SG6000-CN

	Description
2	E2860 controller shelf
3	Controller A
4	Controller B
5	Expansion shelf 1
6	IOM A for expansion shelf 1
7	IOM B for expansion shelf 1
8	Expansion shelf 2
9	IOM A for expansion shelf 2
10	IOM B for expansion shelf 2

5. Connect the power cords and apply power to the expansion shelves.
  - a. Connect a power cord to each of the two power supply units in each expansion shelf.
  - b. Connect the two power cords in each expansion shelf to two different PDUs in the cabinet or rack.
  - c. Turn on the two power switches for each expansion shelf.
    - Do not turn off the power switches during the power-on process.
    - The fans in the expansion shelves might be very loud when they first start up. The loud noise during start-up is normal.
6. Monitor the Home page of the StorageGRID Appliance Installer.

In approximately five minutes, the expansion shelves finish powering up and are detected by the system. The Home page shows the number of new expansion shelves detected, and the Start Expansion button is enabled.

The screenshot shows examples of the messages that could appear on the Home page, depending on the number of existing or new expansion shelves, as follows:

- The banner circled at the top of the page indicates the total number of expansion shelves detected.
  - The banner indicates the total number of expansion shelves, whether the shelves are configured and deployed or new and unconfigured.
  - If no expansion shelves are detected, the banner will not appear.
- The message circled at the bottom of the page indicates an expansion is ready to be started.
  - The message indicates the number of new expansion shelves StorageGRID detects. “Attached” indicates that the shelf is detected. “Unconfigured” indicates that the shelf is new and not yet configured using the StorageGRID Appliance Installer.



Expansion shelves that are already deployed are not included in this message. They are included in the count in the banner at the top of the page.

- The message will not appear if new expansion shelves are not detected.

The expansion is ready to be started. Make sure this page accurately indicates the number of new storage shelves you are trying to add, then click Start Expansion.

The storage system contains 2 expansion shelves.

**This Node**

Node type: Storage

Node name: NetApp-SGA

Cancel Save

**Primary Admin Node connection**

Enable Admin Node discovery:

Primary Admin Node IP: 172.16.4.71

Connection state: Connection to 172.16.4.71 ready

Cancel Save

**Installation**

Current state: Ready to start configuration of 1 attached but unconfigured expansion shelf.

Start Expansion

7. As necessary, resolve any issues described in the messages on the Home page.

For example, use SANtricity System Manager to resolve any storage hardware issues.

8. Verify that the number of expansion shelves displayed on the Home page matches the number of expansion shelves you are adding.



If the new expansion shelves have not been detected, verify that they are properly cabled and powered up.

9. Click **Start Expansion** to configure the expansion shelves and make them available for object storage.

10. Monitor the progress of the expansion shelf configuration.

Progress bars appear on the web page, just as they do during initial installation.

1. Configure storage <span style="float: right;">Running</span>		
Step	Progress	Status
Connect to storage controller	<div style="width: 100%; height: 10px; background-color: green;"></div>	Complete
Clear existing configuration	<div style="width: 100%; height: 10px; background-color: green;"></div>	Skipped
Configure volumes	<div style="width: 30%; height: 10px; background-color: blue;"></div>	Creating volume StorageGRID-obj-22
Configure caching	<div style="width: 0%; height: 10px; background-color: gray;"></div>	Pending
Configure host settings	<div style="width: 0%; height: 10px; background-color: gray;"></div>	Pending


2. Complete storage expansion <span style="float: right;">Pending</span>	

When configuration is complete, the appliance automatically reboots to exit maintenance mode and rejoin the grid. This process can take up to 20 minutes.



If the appliance does not rejoin the grid, go to the StorageGRID Appliance Installer Home page, select **Advanced > Reboot Controller**, and then select **Reboot into Maintenance Mode**.

When the reboot is complete, the **Tasks** tab looks like the following screenshot:



11. Verify the status of the appliance Storage Node and the new expansion shelves.
  - a. In the Grid Manager, select **Nodes** and verify that the appliance Storage Node has a green checkmark icon.
 

The green checkmark icon means that no alerts are active and the node is connected to the grid. For a description of node icons, see the instructions for monitoring and troubleshooting StorageGRID.
  - b. Select the **Storage** tab and confirm that 16 new object stores are shown in the Object Storage table for each expansion shelf you added.
  - c. Verify that each new expansion shelf has a shelf status of Nominal and a configuration status of Configured.

Storage Shelves												
Shelf Chassis Serial Number	Shelf ID	Shelf Status	IOM Status	Power Supply Status	Drawer Status	Fan Status	Drive Slots	Data Drives	Data Drive Size	Cache Drives	Cache Drive Size	Configuration Status
721924500063	99	Nominal	N/A	Nominal	Nominal	Nominal	60	58	9.80 TB	2	800.17 GB	Configured (in use)
721929500038	0	Nominal	Nominal	Nominal	Nominal	Nominal	60	60	9.80 TB	0	0 bytes	Configured (in use)
721929500039	1	Nominal	Nominal	Nominal	Nominal	Nominal	60	60	9.80 TB	0	0 bytes	Configured (in use)

## Related information

[Unpacking the boxes \(SG6000\)](#)

[SG6060: Installing 60-drive shelves into a cabinet or rack](#)

[SG6060: Installing the drives](#)

[Monitor & troubleshoot](#)

## Turning the controller identify LED on and off


The blue identify LED on the front and back of the controller can be turned on to help locate the appliance in a data center.

### What you'll need

You must have the BMC IP address of the controller you want to identify.

### Steps

1. Access the controller BMC interface.
2. Select **Server Identify**.
3. Select **ON** and then select **Perform Action**.



## Result

The blue identify LEDs light on the front (shown) and rear of the controller.



If a bezel is installed on the controller, it might be difficult to see the front identify LED.

### After you finish

To turn off the controller identify LED:

- Press the identify LED switch on the controller front panel.
- From the controller BMC interface, select **Server Identify**, select **OFF** and then select **Perform Action**.

The blue identify LEDs on the front and rear of the controller go off.



### Related information

[Verifying the Fibre Channel HBA to replace](#)

[Locating the controller in a data center](#)

[Accessing the BMC interface](#)

### Locating the controller in a data center

Locate the controller so that you can perform hardware maintenance or upgrades.

### What you'll need

- You have determined which controller requires maintenance.

(Optional) To help locate the controller in your data center, turn on the blue identify LED.

## Turning the controller identify LED on and off

### Steps

1. Find the controller requiring maintenance in the data center.
  - Look for a lit blue identify LED on the front or rear of the controller.

The front identify LED is behind the controller front bezel and might be difficult to see if the bezel is installed.



- Check the tags attached to the front of each controller for a matching part number.
2. Remove the controller front bezel, if one is installed, to access the front panel controls and indicators.
3. Optional: Turn off the blue identify LED if you used it to locate the controller.
  - Press the identify LED switch on the controller front panel.
  - Use the controller BMC interface.

## Turning the controller identify LED on and off

### Related information

[Removing the Fibre Channel HBA](#)

[Removing the SG6000-CN controller from a cabinet or rack](#)

[Shutting down the SG6000-CN controller](#)

### Replacing a storage controller

You might need to replace an E2800 controller or an EF570 controller if it is not functioning optimally or if it has failed.

### What you'll need

- You have a replacement controller with the same part number as the controller you are replacing.
- You have labels to identify each cable that is connected to the controller.
- You have an ESD wristband, or you have taken other antistatic precautions.
- You have a #1 Phillips screwdriver.
- You have the E-Series instructions for replacing a controller in duplex configuration.



Refer to the E-Series instructions only when directed or if you need more details to perform a specific step. Do not rely on the E-Series instructions to replace a controller in the StorageGRID appliance, because the procedures are not the same.

- You have physically located the storage appliance where you are replacing the controller in the data center.

#### [Locating the controller in a data center](#)

### About this task

You can determine if you have a failed controller in two ways:

- The Recovery Guru in SANtricity System Manager directs you to replace the controller.
- The amber Attention LED on the controller is on, indicating that the controller has a fault.



If both controllers in the shelf have their Attention LEDs on, contact technical support for assistance.

Because the storage controller shelf contains two storage controllers, you can replace one of the controllers while your appliance is powered on and performing read/write operations, as long as the following conditions are true:

- The second controller in the shelf has Optimal status.
- The “OK to remove” field in the Details area of the Recovery Guru in SANtricity System Manager displays Yes, indicating that it is safe to remove this component.



If the second controller canister in the shelf does not have Optimal status or if the Recovery Guru indicates that it is not OK to remove the controller canister, contact technical support.

When you replace a controller, you must remove the battery from the original controller and install it in the replacement controller.



The storage controllers in the appliance do not include host interface cards (HIC).


### Steps

1. Unpack the new controller, and set it on a flat, static-free surface.

Save the packing materials to use when shipping the failed controller.

2. Locate the MAC address and FRU part number labels on the back of the replacement controller.

This figure shows the E2800 controller. The procedure for replacing the EF570 controller is identical.



Label	Label	Description
1	MAC address	The MAC address for management port 1 (“P1”). If you used DHCP to obtain the original controller’s IP address, you will need this address to connect to the new controller.
2	FRU part number	The FRU part number. This number must match the replacement part number for the currently installed controller.

3. Prepare to remove the controller.

You use SANtricity System Manager to perform these steps. As needed for additional details, reference the E-Series instructions for replacing the storage controller.

- a. Confirm that the replacement part number for the failed controller is the same as the FRU part number for the replacement controller.

When a controller has a fault and needs to be replaced, the replacement part number is displayed in the Details area of the Recovery Guru. If you need to find this number manually, you can look on the **Base** tab for the controller.



**Possible loss of data access** -- If the two part numbers are not the same, do not attempt this procedure.

- b. Back up the configuration database.

If a problem occurs when you remove a controller, you can use the saved file to restore your configuration.

- c. Collect support data for the appliance.



Collecting support data before and after replacing a component ensures you can send a full set of logs to technical support in case the replacement does not resolve the problem.

- d. Take the controller you plan to replace offline.
- 4. Remove the controller from the appliance:
  - a. Put on an ESD wristband or take other antistatic precautions.
  - b. Label the cables and then disconnect the cables and SFPs.

 To prevent degraded performance, do not twist, fold, pinch, or step on the cables.

- c. Release the controller from the appliance by squeezing the latch on the cam handle until it releases, and then open the cam handle to the right.
- d. Using two hands and the cam handle, slide the controller out of the appliance.

 Always use two hands to support the weight of the controller.

- e. Place the controller on a flat, static-free surface with the removable cover facing up.
- f. Remove the cover by pressing down on the button and sliding the cover off.

- 5. Remove the battery from the failed controller, and install it into the replacement controller:

- a. Confirm that the green LED inside the controller (between the battery and the DIMMs) is off.


If this green LED is on, the controller is still using battery power. You must wait for this LED to go off before removing any components.



Item	Description
	Internal Cache Active LED

Item	Description
	Battery

- b. Locate the blue release latch for the battery.
- c. Unlatch the battery by pushing the release latch down and away from the controller.



Item	Description
	Battery release latch
	Battery

- d. Lift up on the battery, and slide it out of the controller.
- e. Remove the cover from the replacement controller.
- f. Orient the replacement controller so that the slot for the battery faces toward you.
- g. Insert the battery into the controller at a slight downward angle.

You must insert the metal flange at the front of the battery into the slot on the bottom of the controller, and slide the top of the battery beneath the small alignment pin on the left side of the controller.

- h. Move the battery latch up to secure the battery.

When the latch clicks into place, the bottom of the latch hooks into a metal slot on the chassis.

- i. Turn the controller over to confirm that the battery is installed correctly.



**Possible hardware damage** — The metal flange at the front of the battery must be completely inserted into the slot on the controller (as shown in the first figure). If the battery is not installed correctly (as shown in the second figure), the metal flange might contact the controller board, causing damage.

- **Correct** — The battery's metal flange is completely inserted in the slot on the controller:



- **Incorrect** — The battery's metal flange is not inserted into the slot on the controller:



- j. Replace the controller cover.

6. Install the replacement controller into the appliance.

- a. Turn the controller over, so that the removable cover faces down.
- b. With the cam handle in the open position, slide the controller all the way into the appliance.

- c. Move the cam handle to the left to lock the controller in place.
- d. Replace the cables and SFPs.
- e. If the original controller used DHCP for the IP address, locate the MAC address on the label on the back of the replacement controller. Ask your network administrator to associate the DNS/network and IP address for the controller you removed with the MAC address for the replacement controller.



If the original controller did not use DHCP for the IP address, the new controller will adopt the IP address of the controller you removed.

7. Bring the controller online using SANtricity System Manager:
  - a. Select **Hardware**.
  - b. If the graphic shows the drives, select **Show back of shelf**.
  - c. Select the controller you want to place online.
  - d. Select **Place Online** from the context menu, and confirm that you want to perform the operation.
  - e. Verify that the seven-segment display shows a state of 99.
8. Confirm that the new controller is Optimal, and collect support data.

#### Related information

[NetApp E-Series Systems Documentation Site](#)

### Replacing hardware components in the storage controller shelf

If a hardware problem occurs, you might need to replace a component in the storage controller shelf.

#### What you'll need

- You have the E-Series hardware replacement procedure.
- You have physically located the storage appliance where you are replacing storage shelf hardware components in the data center.

[Locating the controller in a data center](#)

#### About this task

To replace the battery in the storage controller, see the instructions in these instructions for replacing a storage controller. Those instructions describe how to remove a controller from the appliance, remove the battery from the controller, install the battery, and replace the controller.

For instructions for the other field replaceable units (FRUs) in the controller shelves, access the E-Series procedures for system maintenance.

FRU	See instructions
Battery	StorageGRID (these instructions): Replacing a storage controller

FRU	See instructions
Drive	E-Series: <ul style="list-style-type: none"> <li>• Replace drive (60-drive)</li> <li>• Replace drive (12-drive or 24-drive)</li> </ul>
Power canister	E-Series <ul style="list-style-type: none"> <li>• Replace power canister (60-drive)</li> <li>• Replace power supply (12-drive or 24-drive)</li> </ul>
Fan canister (60-drive shelves only)	E-Series: Replace fan canister (60-drive)
Drive drawer (60-drive shelves only)	E-Series: Replace drive drawer (60-drive)

**Related information**

[NetApp E-Series Systems Documentation Site](#)

[Replacing a storage controller](#)

**Replacing hardware components in the optional 60-drive expansion shelf**

You might need to replace an input/output module, a power supply, or a fan in the expansion shelf.

**What you'll need**

- You have the E-Series hardware replacement procedure.
- You have physically located the storage appliance where you are replacing expansion shelf hardware components in the data center.

[Locating the controller in a data center](#)

**About this task**

To replace an input/output module (IOM) in a 60-drive expansion shelf, see the instructions in these instructions for replacing a storage controller.

To replace a power supply or a fan in a 60-drive expansion shelf, access the E-Series procedures for maintaining 60-drive hardware.

FRU	See E-Series instructions for
Input/output module (IOM)	Replacing an IOM
Power canister	Replace power canister (60-drive)
Fan canister	Replace fan canister (60-drive)

## Shutting down the SG6000-CN controller

Shut down the SG6000-CN controller to perform hardware maintenance.

### What you'll need

- You have physically located the SG6000-CN controller requiring maintenance in the data center.

[Locating the controller in a data center](#)

- The appliance has been placed maintenance mode.

[Placing an appliance into maintenance mode](#)

### About this task

To prevent service interruptions, confirm that all other Storage Nodes are connected to the grid before shutting down the controller or shut down the controller during a scheduled maintenance window when periods of service disruption are normally expected. See the information about determining node connection states in the instructions for managing objects with information lifecycle management.



If you have ever used an ILM rule that creates only one copy of an object, you must shut down the controller during a scheduled maintenance window. Otherwise, you might temporarily lose access to those objects during this procedure. See information about managing objects with information lifecycle management.

### Steps

1. When the appliance has been placed maintenance mode, shut down the SG6000-CN controller:



You must perform a controlled shut down of the controller by entering the commands specified below. Shutting down the controller using the power switch will result in data loss.

- a. Log in to the grid node using PuTTY or another ssh client:
  - 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. Shut down the SG6000-CN controller:

**shutdown -h now**

This command might take up to 10 minutes to complete.


2. Use one of the following methods to verify that the SG6000-CN controller is powered off:
  - Look at the blue power LED on the front of the controller and confirm that it is off.



- Look at the green LEDs on both power supplies in the rear of the controller and confirm that they blink at a regular rate (approximately one blink per second).



- Use the controller BMC interface:
  - i. Access the controller BMC interface.  
[Accessing the BMC interface](#)
  - ii. Select **Power Control**.
  - iii. Verify that the Power Actions indicates that the host is currently off.



### Related information

[Removing the SG6000-CN controller from a cabinet or rack](#)

### Powering on the SG6000-CN controller and verifying operation

Power on the controller after completing maintenance.

#### What you'll need

- You have installed the controller in a cabinet or rack and connected the data and power cables.

[Reinstalling the SG6000-CN controller into a cabinet or rack](#)

- You have physically located the controller in the data center.

[Locating the controller in a data center](#)

#### Steps


1. Power on the SG6000-CN controller and monitor the controller LEDs and boot-up codes using one of the following methods:
  - Press the power switch on the front of the controller.



- Use the controller BMC interface:
  - i. Access the controller BMC interface.

[Accessing the BMC interface](#)

- ii. Select **Power Control**.
- iii. Select **Power On Server** and then select **Perform Action**.



Use the BMC interface to monitor start-up status.

2. Confirm that the appliance controller displays in the Grid Manager and with no alerts.

It might take up to 20 minutes for the controller to display in the Grid Manager.

3. Confirm that the new SG6000-CN controller is fully operational:
  - a. Log in to the grid node using PuTTY or another ssh client:
    - i. Enter the following command: `ssh admin@grid_node_IP`
    - ii. Enter the password listed in the `Passwords.txt` file.
    - iii. Enter the following command to switch to root: `su -`
    - iv. Enter the password listed in the `Passwords.txt` file.

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

- b. Enter the following command and verify that it returns the expected output:  
`cat /sys/class/fc_host/*/port_state`

Expected output:

```
Online
Online
Online
```

If the expected output is not returned, contact technical support.

- c. Enter the following command and verify that it returns the expected output:  
`cat /sys/class/fc_host/*/speed`

Expected output:

```
16 Gbit
16 Gbit
16 Gbit16 Gbit
16 Gbit
```

If the expected output is not returned, contact technical support.

- d. From the Nodes page in Grid Manager, make sure that the appliance node is connected to the grid and does not have any alerts.



Do not take another appliance node offline unless this appliance has a green icon.

4. Optional: Install the front bezel, if one was removed.

#### Related information

[Viewing status indicators and buttons on the SG6000-CN controller](#)

[Viewing boot-up status codes for the SG6000 storage controllers](#)

## Replacing the SG6000-CN controller

You might need to replace the SG6000-CN controller if it is not functioning optimally or if it has failed.

### What you'll need

- You have a replacement controller with the same part number as the controller you are replacing.
- You have labels to identify each cable that is connected to the controller.
- You have physically located the controller to replace in the data center.

#### [Locating the controller in a data center](#)

### About this task

The appliance Storage Node will not be accessible when you replace the SG6000-CN controller. If the SG6000-CN controller is functioning sufficiently, you can perform a controlled shutdown at the start of this procedure.



If you are replacing the controller before installing StorageGRID software, you might not be able to access the StorageGRID Appliance Installer immediately after completing this procedure. While you can access the StorageGRID Appliance Installer from other hosts on the same subnet as the appliance, you cannot access it from hosts on other subnets. This condition should resolve itself within 15 minutes (when any ARP cache entries for the original controller time out), or you can clear the condition immediately by purging any old ARP cache entries manually from the local router or gateway.

### Steps

1. If the SG6000-CN controller is functioning sufficiently to allow for a controlled shutdown, shut down the SG6000-CN controller.

#### [Shutting down the SG6000-CN controller](#)

The green Cache Active LED on the back of the E2800 controller is on when cached data needs to be written to the drives. You must wait for this LED to turn off.

2. Use one of two methods to verify that the power for the SG6000-CN controller is off:
  - The power indicator LED on the front of the controller is off.
  - The Power Control page of the BMC interface indicates that the controller is off.
3. If the StorageGRID networks attached to the controller use DHCP servers, update DNS/network and IP address settings.
  - a. Locate the MAC address label on the front of the SG6000-CN controller, and determine the MAC address for the Admin Network port.



The MAC address label lists the MAC address for the BMC management port. To determine the MAC address for the Admin Network port, you must add **2** to the hexadecimal number on the label. For example, if the MAC address on the label ends in **09**, the MAC address for the Admin Port would end in **0B**. If the MAC address on the label ends in **(y)FF**, the MAC address for the Admin Port would end in **(y+1)01**. You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing **+ 2 =**.

- b. Ask your network administrator to associate the DNS/network and IP address for the controller you removed with the MAC address for the replacement controller.



You must ensure that all IP addresses for the original controller have been updated before you apply power to the replacement controller. Otherwise, the controller will obtain new DHCP IP addresses when it boots up and might not be able to reconnect to StorageGRID. This step applies to all StorageGRID networks that are attached to the controller.



If the original controller used static IP address, the new controller will automatically adopt the IP addresses of the controller you removed.

4. Remove and replace the SG6000-CN controller:

- a. Label the cables and then disconnect the cables and any SFP+ or SFP28 transceivers.



To prevent degraded performance, do not twist, fold, pinch, or step on the cables.

- b. Remove the failed controller from the cabinet or rack.
- c. Install the replacement controller into the cabinet or rack.
- d. Replace the cables and any SFP+ or SFP28 transceivers.
- e. Power on the controller and monitor the controller LEDs and boot-up codes.

5. Confirm that the appliance Storage Node appears in the Grid Manager and that no alarms appear.

6. From the Grid Manager, select **Nodes**, and verify that the BMC IP address for the node controller is correct.

If the node controller IP address is not valid or is not in the expected range, reconfigure the IP address as described in the recovery and maintenance instructions.

### [Maintain & recover](#)

#### **Related information**

[SG6000-CN: Installing into a cabinet or rack](#)

[Viewing status indicators and buttons on the SG6000-CN controller](#)

[Viewing boot-up codes for the SG6000-CN controller](#)

#### **Replacing a power supply in the SG6000-CN controller**

The SG6000-CN controller has two power supplies for redundancy. If one of the power supplies fails, you must replace it as soon as possible to ensure that the compute controller has redundant power.

#### **What you'll need**

- You have unpacked the replacement power supply unit.
- You have physically located the controller where you are replacing the power supply in the data center.

[Locating the controller in a data center](#)

- You have confirmed that the other power supply is installed and running.


### About this task

The figure shows the two power supply units for the SG6000-CN controller, which are accessible from the back of the controller.




### Steps

1. Unplug the power cord from the power supply.
2. Lift the cam handle.




3. Press the blue latch and pull the power supply out.



- Slide the replacement power supply into the chassis.

Ensure that the blue latch is on the right side when you slide the unit in.



- Push the cam handle down to secure the power supply.
- Attach the power cord to the power supply, and ensure that the green LED comes on.

### Removing the SG6000-CN controller from a cabinet or rack

Remove the SG6000-CN controller from a cabinet or rack to access the top cover or to move the controller to a different location.

#### What you'll need

- You have labels to identify each cable that is connected to the SG6000-CN controller.
- You have physically located the SG6000-CN controller where you are performing maintenance in the data center.

#### [Locating the controller in a data center](#)

- You have shut down the SG6000-CN controller.

#### [Shutting down the SG6000-CN controller](#)



Do not shut down the controller using the power switch.

#### Steps

- Label and then disconnect the controller power cables.
- Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
- Label and then disconnect the controller data cables and any SFP+ or SFP28 transceivers.



To prevent degraded performance, do not twist, fold, pinch, or step on the cables.

- Loosen the two captive screws on the controller front panel.



5. Slide the SG6000-CN controller forward out of the rack until the mounting rails are fully extended and you hear the latches on both sides click.

The controller top cover is accessible.

6. Optional: If you are fully removing the controller from the cabinet or rack, follow the instructions for the rail kit to remove the controller from the rails.

### **Related information**

[Removing the SG6000-CN controller cover](#)

### **Reinstalling the SG6000-CN controller into a cabinet or rack**

Reinstall the controller into a cabinet or rack when hardware maintenance is complete.

### **What you'll need**


You have reinstalled the controller cover.

[Reinstalling the SG6000-CN controller cover](#)

### **Steps**

1. Press the blue rail releases both rack rails at the same time and slide the SG6000-CN controller into the rack until it is fully seated.

When you cannot move the controller any further, pull the blue latches on both sides of the chassis to slide the controller all the way in.



Do not attach the front bezel until after you power on the controller.

2. Tighten the captive screws on the controller front panel to secure the controller in the rack.



3. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
4. Reconnect the controller data cables and any SFP+ or SFP28 transceivers.



To prevent degraded performance, do not twist, fold, pinch, or step on the cables.

#### [Cabling the appliance \(SG6000\)](#)

5. Reconnect the controller power cables.

#### [Connecting power cords and applying power \(SG6000\)](#)

#### **After you finish**

The controller can be restarted.

#### [Powering on the SG6000-CN controller and verifying operation](#)

## Removing the SG6000-CN controller cover

Remove the controller cover to access internal components for maintenance.

### What you'll need

Remove the controller from the cabinet or rack to access the top cover.

### Removing the SG6000-CN controller from a cabinet or rack

#### Steps

1. Make sure that the SG6000-CN controller cover latch is not locked. If necessary, turn the blue plastic latch lock one-quarter turn in the unlock direction, as shown on the latch lock.
2. Rotate the latch up and back toward the rear of the SG6000-CN controller chassis until it stops; then, carefully lift the cover from the chassis and set it aside.



Wrap the strap end of an ESD wristband around your wrist and secure the clip end to a metal ground to prevent static discharge when working inside the SG6000-CN controller.

### Related information

[Removing the Fibre Channel HBA](#)

## Reinstalling the SG6000-CN controller cover

Reinstall the controller cover when internal hardware maintenance is complete.

### What you'll need

You have completed all maintenance procedures inside the controller.

#### Steps

1. With the cover latch open, hold the cover above the chassis and align the hole in the top cover latch with the pin in the chassis. When the cover is aligned, lower it onto the chassis.



2. Rotate the cover latch forward and down until it stops and the cover fully seats into the chassis. Verify that there are no gaps along the front edge of the cover.

If the cover is not fully seated, you might not be able to slide the SG6000-CN controller into the rack.

3. Optional: Turn the blue plastic latch lock one-quarter turn in the lock direction, as shown on the latch lock, to lock it.

#### After you finish

Reinstall the controller in the cabinet or rack.

[Reinstalling the SG6000-CN controller into a cabinet or rack](#)

#### Replacing the Fibre Channel HBA in the SG6000-CN controller

You might need to replace the Fibre Channel host bus adapter (HBA) in the SG6000-CN controller if it is not functioning optimally or if it has failed.

#### Verifying the Fibre Channel HBA to replace

If you are unsure which Fibre Channel host bus adapter (HBA) to replace, complete this procedure to identify it.

#### What you'll need

- You have the serial number of the storage appliance or SG6000-CN controller where the Fibre Channel HBA needs to be replaced.



If the serial number of the storage appliance containing the Fibre Channel HBA you are replacing starts with the letter Q, it will not be listed in the Grid Manager. You must check the tags attached to the front of each SG6000-CN controller in the data center until you find a match.

- You must be signed in to the Grid Manager using a supported browser.

#### Steps

1. From the Grid Manager, select **Nodes**.
2. From the tree view of the Nodes page, select an appliance Storage Node.
3. Select the **Hardware** tab.

Check the Storage Appliance Chassis Serial Number and the Compute Controller Serial Number in the StorageGRID Appliance section to see if one of these serial numbers matches the serial number of the storage appliance where you are replacing the Fibre Channel HBA. If either serial number matches, you have found the correct appliance.



- If the StorageGRID Appliance section does not display, the node selected is not a StorageGRID appliance. Select a different node from the tree view.
- If the Appliance Model is not SG6060, select a different node from the tree view.
- If the serial numbers do not match, select a different node from the tree view.

4. After you locate the node where the Fibre Channel HBA needs to be replaced, write down the Compute Controller BMC IP address listed the StorageGRID Appliance section.

You can use this IP address to turn on the compute controller identify LED, to help you locate the appliance in the data center.

### Turning the controller identify LED on and off

#### Related information

[Removing the Fibre Channel HBA](#)

#### Removing the Fibre Channel HBA

You might need to replace the Fibre Channel host bus adapter (HBA) in the SG6000-CN controller if it is not functioning optimally or if it has failed.

#### What you'll need

- You have the correct replacement Fibre Channel HBA.

- You have determined which SG6000-CN controller contains the Fibre Channel HBA to replace.

#### [Verifying the Fibre Channel HBA to replace](#)

- You have physically located the SG6000-CN controller where you are replacing the Fibre Channel HBA in the data center.

#### [Locating the controller in a data center](#)

- You have removed the controller cover.

#### [Removing the SG6000-CN controller cover](#)

### About this task


To prevent service interruptions, confirm that all other Storage Nodes are connected to the grid before starting the Fibre Channel HBA replacement or replace the adapter during a scheduled maintenance window when periods of service disruption are normally expected. See the information about determining node connection states in the instructions for managing objects with information lifecycle management.



If you have ever used an ILM rule that creates only one copy of an object, you must replace the Fibre Channel HBA during a scheduled maintenance window. Otherwise, you might temporarily lose access to those objects during this procedure. See information about managing objects with information lifecycle management.

### Steps

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
2. Locate the riser assembly at the rear of the controller that contains the Fibre Channel HBA.



3. Grasp the riser assembly through the blue-marked holes and carefully lift it upwards. Move the riser assembly toward the front of the chassis as you lift it to allow the external connectors in its installed

adapters to clear the chassis.

4. Place the riser card on a flat anti-static surface with the metal frame side down to access the adapters.



There are two adapters in the riser assembly: a Fibre Channel HBA and an Ethernet network adapter. The Fibre Channel HBA is indicated in the illustration.

5. Open the blue adapter latch (circled) and carefully remove the Fibre Channel HBA from the riser assembly. Rock the adapter slightly to help remove the adapter from its connector. Do not use excessive force.
6. Place the adapter on a flat anti-static surface.

#### **After you finish**

Install the replacement Fibre Channel HBA.

#### [Reinstalling the Fibre Channel HBA](#)

#### **Related information**

#### [Reinstalling the Fibre Channel HBA](#)

#### [Administer StorageGRID](#)

#### [Monitor & troubleshoot](#)

#### [Manage objects with ILM](#)

#### **Reinstalling the Fibre Channel HBA**

The replacement Fibre Channel HBA is installed into the same location as the one that was removed.

#### **What you'll need**

- You have the correct replacement Fibre Channel HBA.
- You have removed the existing Fibre Channel HBA.

#### [Removing the Fibre Channel HBA](#)

#### **Steps**

1. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.

2. Remove the replacement Fibre Channel HBA from its packaging.
3. With the blue adapter latch in the open position, align the Fibre Channel HBA with its connector on the riser assembly; then, carefully press the adapter into the connector until it is fully seated.



There are two adapters in the riser assembly: a Fibre Channel HBA and an Ethernet network adapter. The Fibre Channel HBA is indicated in the illustration.

4. Locate the alignment hole on the riser assembly (circled) that aligns with a guide pin on the system board to ensure correct riser assembly positioning.



5. Position the riser assembly in the chassis, making sure that it aligns with the connector and guide pin on the system board; then, insert the riser assembly.
6. Carefully press the riser assembly in place along its center line, next to the blue-marked holes, until it is fully seated.
7. Remove the protective caps from the Fibre Channel HBA ports where you will be reinstalling cables.

#### **After you finish**

If you have no other maintenance procedures to perform in the controller, reinstall the controller cover.

#### [Reinstalling the SG6000-CN controller cover](#)

#### **Changing the link configuration of the SG6000-CN controller**

You can change the Ethernet link configuration of the SG6000-CN controller. You can change the port bond mode, the network bond mode, and the link speed.

## What you'll need

The appliance has been placed in maintenance mode.

## Placing an appliance into maintenance mode


### About this task

Options for changing the Ethernet link configuration of the SG6000-CN controller include:

- Changing **Port bond mode** from Fixed to Aggregate, or from Aggregate to Fixed
- Changing **Network bond mode** from Active-Backup to LACP, or from LACP to Active-Backup
- Enabling or disabling VLAN tagging, or changing the value of a VLAN tag
- Changing the link speed.

### Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > Link Configuration**.



2. Make the desired changes to the link configuration.

For more information on the options, see [Configuring network links \(SG6000\)](#).

3. When you are satisfied with your selections, click **Save**.



You might lose your connection if you made changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:

**`https://Appliance_Controller_IP:8443`**

If you made changes to the VLAN settings, the subnet for the appliance might have changed. If you need to change the IP addresses for the appliance, follow the instructions for configuring IP addresses.


## Configuring StorageGRID IP addresses

4. Select **Configure Networking > Ping Test** from the menu.
5. Use the Ping Test tool to check connectivity to IP addresses on any networks that might have been affected by the link configuration changes you made in the [link configuration changes](#) step.


In addition to any other tests you choose to perform, confirm that you can ping the Grid Network IP address

of the primary Admin Node, and the Grid Network IP address of at least one other Storage Node. If necessary, return to the [link configuration changes](#) step and correct any link configuration issues.

- When you are satisfied that your link configuration changes are working, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
  - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Changing the MTU setting

You can change the MTU setting that you assigned when you configured IP addresses for the appliance node.

### What you'll need

The appliance has been placed maintenance mode.

**Steps**

1. From the StorageGRID Appliance Installer, select **Configure Networking > IP Configuration**.
2. Make the desired changes to the MTU settings for the Grid Network, Admin Network, and Client Network.

**Grid Network**

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

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

**⚠** All required Grid Network subnets must also be defined in the Grid Network Subnet List on the Primary Admin Node before starting installation.

Subnets (CIDR)	<input type="text" value="172.18.0.0/21"/>	✘
	<input type="text" value="172.18.0.0/21"/>	✘
	<input type="text" value="192.168.0.0/21"/>	+ ✘
MTU	<input type="text" value="1500"/>	
	<input type="button" value="Cancel"/>	
	<input type="button" value="Save"/>	



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.




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

- When you are satisfied with the settings, select **Save**.
- Reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
  - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Administer StorageGRID](#)

## Checking the DNS server configuration

You can check and temporarily change the domain name system (DNS) servers that are currently in use by this appliance node.

### What you'll need

The appliance has been placed maintenance mode.

### Placing an appliance into maintenance mode

#### About this task

You might need to change the DNS server settings if an encrypted appliance cannot connect to the key management server (KMS) or KMS cluster because the hostname for the KMS was specified as a domain name instead of an IP address. Any changes that you make to the DNS settings for the appliance are temporary and are lost when you exit maintenance mode. To make these changes permanent, specify the DNS servers in Grid Manager (**Maintenance > Network > DNS Servers**).

- Temporary changes to the DNS configuration are necessary only for node-encrypted appliances where the KMS server is defined using a fully qualified domain name, instead of an IP address, for the hostname.
- When a node-encrypted appliance connects to a KMS using a domain name, it must connect to one of the DNS servers defined for the grid. One of these DNS servers then translates the domain name into an IP address.
- If the node cannot reach a DNS server for the grid, or if you changed the grid-wide DNS settings when a node-encrypted appliance node was offline, the node is unable to connect to the KMS. Encrypted data on the appliance cannot be decrypted until the DNS issue is resolved.

To resolve a DNS issue preventing KMS connection, specify the IP address of one or more DNS servers in the StorageGRID Appliance Installer. These temporary DNS settings allow the appliance to connect to the KMS and decrypt data on the node.

For example, if the DNS server for the grid changes while an encrypted node was offline, the node will not be able to reach the KMS when it comes back online, since it is still using the previous DNS values. Entering the new DNS server IP address in the StorageGRID Appliance Installer allows a temporary KMS connection to decrypt the node data.

#### Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > DNS Configuration**.
2. Verify that the DNS servers specified are correct.

## DNS Servers

⚠ Configuration changes made on this page will not be passed to the StorageGRID software after appliance installation.

### Servers

Server 1	<input type="text" value="10.224.223.135"/>	✕
Server 2	<input type="text" value="10.224.223.136"/>	+ ✕
<input type="button" value="Cancel"/>		<input type="button" value="Save"/>

3. If required, change the DNS servers.



Changes made to the DNS settings are temporary and are lost when you exit maintenance mode.

4. When you are satisfied with the temporary DNS settings, select **Save**.

The node uses the DNS server settings specified on this page to reconnect to the KMS, allowing data on the node to be decrypted.

5. After node data is decrypted, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.

NetApp® StorageGRID® Appliance Installer

Home | Configure Networking ▾ | Configure Hardware ▾ | Monitor Installation | Advanced ▾

Reboot Controller  
Request a controller reboot.


RAID Mode  
Upgrade Firmware  
Reboot Controller

Reboot into StorageGRID | Reboot into Maintenance Mode



When the node reboots and rejoins the grid, it uses the system-wide DNS servers listed in the Grid Manager. After rejoining the grid, the appliance will no longer use the temporary DNS servers specified in the StorageGRID Appliance Installer while the appliance was in maintenance mode.

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



### Monitoring node encryption in maintenance mode

If you enabled node encryption for the appliance during installation, you can monitor the node-encryption status of each appliance node, including the node-encryption state and key management server (KMS) details.

#### What you'll need

- Node encryption must have been enabled for the appliance during installation. You cannot enable node encryption after the appliance is installed.
- The appliance has been placed into maintenance mode.

[Placing an appliance into maintenance mode](#)

#### Steps

1. From the StorageGRID Appliance Installer, select **Configure Hardware > Node Encryption**.

## Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

### Encryption Status

 You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

Save

### Key Management Server Details

View the status and configuration details for the KMS that manages the encryption key for this appliance. You must use the Grid Manager to make configuration changes.

KMS display name	thales
External key UID	41b0306abcce451facfe01b1b4870ae1c1ec6bd5e3849d790223766baf35c57
Hostnames	10.96.99.164 10.96.99.165
Port	5696

Server certificate >

Client certificate >

### Clear KMS Key

 Do not clear the KMS key if you need to access or preserve any data on this appliance.

If you want to reinstall this appliance node (for example, in another grid), you must clear the KMS key. When the KMS key is cleared, all data on this appliance is deleted.

Clear KMS Key and Delete Data

The Node Encryption page includes these three sections:

- Encryption Status shows whether node encryption is enabled or disabled for the appliance.
- Key Management Server Details shows information about the KMS being used to encrypt the appliance. You can expand the server and client certificate sections to view certificate details and status.
  - To address issues with the certificates themselves, such as renewing expired certificates, see the information about KMS in the instructions for administering StorageGRID.
  - If there are unexpected problems connecting to KMS hosts, verify that the domain name system (DNS) servers are correct and that appliance networking is correctly configured.  
[Checking the DNS server configuration](#)
  - If you are unable to resolve your certificate issues, contact technical support.
- Clear KMS Key disables node encryption for the appliance, removes the association between the appliance and the key management server that was configured for the StorageGRID site, and deletes


all data from the appliance. You must clear the KMS key before you can install the appliance into another StorageGRID system.

### Clearing the key management server configuration




Clearing the KMS configuration deletes data from the appliance, rendering it permanently inaccessible. This data is not recoverable.

- When you are done checking node-encryption status, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
  - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



### Related information

[Administer StorageGRID](#)

## Clearing the key management server configuration

Clearing the key management server (KMS) configuration disables node encryption on your appliance. After clearing the KMS configuration, the data on your appliance is permanently deleted and is no longer accessible. This data is not recoverable.

### What you'll need

If you need to preserve data on the appliance, you must perform a node decommission procedure before you clear the KMS configuration.



When KMS is cleared, data on the appliance will be permanently deleted and no longer accessible. This data is not recoverable.

Decommission the node to move any data it contains to other nodes in StorageGRID. See the recovery and maintenance instructions for grid node decommissioning.

### About this task

Clearing the appliance KMS configuration disables node encryption, removing the association between the appliance node and the KMS configuration for the StorageGRID site. Data on the appliance is then deleted and the appliance is left in a pre-install state. This process cannot be reversed.

You must clear the KMS configuration:

- Before you can install the appliance into another StorageGRID system, that does not use a KMS or that uses a different KMS.



Do not clear the KMS configuration if you plan to reinstall an appliance node in a StorageGRID system that uses the same KMS key.

- Before you can recover and reinstall a node where the KMS configuration was lost and the KMS key is not recoverable.
- Before returning any appliance that was previously in use at your site.
- After decommissioning a appliance that had node encryption enabled.



Decommission the appliance before clearing KMS to move its data to other nodes in your StorageGRID system. Clearing KMS before decommissioning the appliance will result in data loss and might render the appliance inoperable.

### Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.  
**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > Node Encryption**.

## Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

### Encryption Status

 You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

Save

### Key Management Server Details

View the status and configuration details for the KMS that manages the encryption key for this appliance. You must use the Grid Manager to make configuration changes.

KMS display name	thales
External key UID	41b0306abcce451facfe01b1b4870ae1c1ec6bd5e3849d790223766baf35c57
Hostnames	10.96.99.164 10.96.99.165
Port	5696

Server certificate >

Client certificate >

### Clear KMS Key

 Do not clear the KMS key if you need to access or preserve any data on this appliance.


If you want to reinstall this appliance node (for example, in another grid), you must clear the KMS key. When the KMS key is cleared, all data on this appliance is deleted.

Clear KMS Key and Delete Data



If the KMS configuration is cleared, data on the appliance will be permanently deleted. This data is not recoverable.

- At the bottom of the window, select **Clear KMS Key and Delete Data**.
- If you are sure that you want to clear the KMS configuration, type **clear** and select **Clear KMS Key and Delete Data**.



The KMS encryption key and all data are deleted from the node, and the appliance reboots. This can take up to 20 minutes.

5. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

6. Select **Configure Hardware > Node Encryption**.
7. Verify that node encryption is disabled and that the key and certificate information in **Key Management Server Details** and the **Clear KMS Key and Delete Data** control are removed from the window.

Node encryption cannot be reenabled on the appliance until it is reinstalled in a grid.

#### After you finish

After the appliance reboots and you have verified that KMS has been cleared and that the appliance is in a pre-install state, you can physically remove the appliance from your StorageGRID system. See the recovery and maintenance instructions for information about preparing an appliance for reinstallation.

#### Related information

[Administer StorageGRID](#)

[Maintain & recover](#)

## SG5700 storage appliances

Learn how to install and maintain StorageGRID SG5712 and SG5760 appliances.

- [StorageGRID appliance overview](#)

- [Installation and deployment overview](#)
- [Preparing for installation](#)
- [Installing the hardware](#)
- [Configuring the hardware](#)
- [Deploying an appliance Storage Node](#)
- [Monitoring the storage appliance installation](#)
- [Automating appliance installation and configuration](#)
- [Overview of installation REST APIs](#)
- [Troubleshooting the hardware installation](#)
- [Maintaining the SG5700 appliance](#)

## StorageGRID appliance overview

The SG5700 StorageGRID appliance is an integrated storage and computing platform that operates as a Storage Node in a StorageGRID grid. The appliance can be used in a hybrid grid environment that combines appliance Storage Nodes and virtual (software-based) Storage Nodes.

The StorageGRID SG5700 appliance provides the following features:

- Integrates the storage and computing elements for a StorageGRID Storage Node.
- Includes the StorageGRID Appliance Installer to simplify Storage Node deployment and configuration.
- Includes E-Series SANtricity System Manager for hardware management and monitoring.
- Supports up to four 10-GbE or 25-GbE connections to the StorageGRID Grid Network and Client Network.
- Supports Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. When these drives are used with the Drive Security feature in SANtricity System Manager, unauthorized access to data is prevented.

The SG5700 appliance is available in two models: the SG5712 and the SG5760. Both models include the following components:

Component	SG5712	SG5760
Compute controller	E5700SG controller	E5700SG controller
Storage controller	E-Series E2800 controller	E-Series E2800 controller
Chassis	E-Series DE212C enclosure, a two rack-unit (2U) enclosure	E-Series DE460C enclosure, a four rack-unit (4U) enclosure
Drives	12 NL-SAS drives (3.5-inch)	60 NL-SAS drives (3.5-inch)
Redundant power supplies and fans	Two power-fan canisters	Two power canisters and two fan canisters


The maximum raw storage available in the StorageGRID appliance is fixed, based on the number of drives in each enclosure. You cannot expand the available storage by adding a shelf with additional drives.

### Model SG5712

This figure shows the front and back of the SG5712 model, a 2U enclosure that holds 12 drives.




The SG5712 includes two controllers and two power-fan canisters.



	Description
1	E2800 controller (storage controller)
2	E5700SG controller (compute controller)
3	Power-fan canisters

### Model SG5760

This figure shows the front and back of the SG5760 model, a 4U enclosure that holds 60 drives in 5 drive drawers.



The SG5760 includes two controllers, two fan canisters, and two power canisters.

	Description
1	E2800 controller (storage controller)
2	E5700SG controller (compute controller)
3	Fan canister (1 of 2)
4	Power canister (1 of 2)

#### Related information

[NetApp E-Series Systems Documentation Site](#)

#### Controllers in the StorageGRID appliance

Both the SG5712 and SG5760 models of the StorageGRID appliance include an E5700SG controller and an E2800 controller. You should review the diagrams to learn the differences between the controllers.

#### E5700SG controller


- Operates as the compute server for the appliance.
- Includes the StorageGRID Appliance Installer.



StorageGRID software is not preinstalled on the appliance. This software is accessed from the Admin Node when you deploy the appliance.

- Can connect to all three StorageGRID networks, including the Grid Network, the Admin Network, and the Client Network.
- Connects to the E2800 controller and operates as the initiator.

This figure shows the connectors on the back of the E5700SG controller.




	Port	Type	Use
1	Interconnect ports 1 and 2	16Gb/s Fibre Channel (FC), optical SFPa	Connect the E5700SG controller to the E2800 controller.
2	Diagnostic and support ports	<ul style="list-style-type: none"> <li>• RJ-45 serial port</li> <li>• Micro USB serial port</li> <li>• USB port</li> </ul>	Reserved for technical support.
3	Drive expansion ports	12Gb/s SAS	Not used. StorageGRID appliances do not support expansion drive shelves.
4	Network ports 1-4	10-GbE or 25-GbE, based on SFP transceiver type, switch speed, and configured link speed	Connect to the Grid Network and the Client Network for StorageGRID.
5	Management port 1	1-Gb (RJ-45) Ethernet	Connect to the Admin Network for StorageGRID.

	Port	Type	Use
6	Management port 2	1-Gb (RJ-45) Ethernet	<p>Options:</p> <ul style="list-style-type: none"> <li>• Bond with management port 1 for a redundant connection to the Admin Network for StorageGRID.</li> <li>• Leave unwired and available for temporary local access (IP 169.254.0.1).</li> <li>• During installation, use port 2 for IP configuration if DHCP-assigned IP addresses are not available.</li> </ul>

### E2800 controller

- Operates as the storage controller for the appliance.
- Manages the storage of data on the drives.
- Functions as a standard E-Series controller in simplex mode.
- Includes SANtricity OS Software (controller firmware).
- Includes SANtricity System Manager for monitoring appliance hardware and for managing alerts, the AutoSupport feature, and the Drive Security feature.
- Connects to the E5700SG controller and operates as the target.

This figure shows the connectors on the back of the E2800 controller.



	Port	Type	Use
1	Interconnect ports 1 and 2	16Gb/s FC optical SFPa	Connect the E2800 controller to the E5700SG controller.
2	Management ports 1 and 2	1-Gb (RJ-45) Ethernet	<ul style="list-style-type: none"> <li>• Port 1 connects to the network where you access SANtricity System Manager on a browser.</li> <li>• Port 2 is reserved for technical support use.</li> </ul>

	Port	Type	Use
3	Diagnostic and support ports	<ul style="list-style-type: none"> <li>• RJ-45 serial port</li> <li>• Micro USB serial port</li> <li>• USB port</li> </ul>	Reserved for technical support use.
4	Drive expansion ports.	12Gb/s SAS	Not used. StorageGRID appliances do not support expansion drive shelves.

## Installation and deployment overview

You can install one or more StorageGRID appliances when you first deploy StorageGRID, or you can add appliance Storage Nodes later as part of an expansion. You might also need to install an appliance Storage Node as part of a recovery operation.

Adding a StorageGRID storage appliance to a StorageGRID system includes four primary steps:

### 1. Preparing for installation:

- Preparing the installation site
- Unpacking the boxes and checking the contents
- Obtaining additional equipment and tools
- Gathering IP addresses and network information
- Optional: Configuring an external key management server (KMS) if you plan to encrypt all appliance data. See details about external key management in the instructions for administering StorageGRID.

### 2. Installing the hardware:

- Registering the hardware
- Installing the appliance into a cabinet or rack
- Installing the drives (SG5760 only)
- Cabling the appliance
- Connecting the power cords and applying power
- Viewing boot-up status codes

### 3. Configuring the hardware:

- Accessing SANtricity System Manager, setting a static IP address for management port 1 on the E2800 controller, and configuring SANtricity System Manager settings
- Accessing StorageGRID Appliance Installer and configuring the link and network IP settings required to connect to StorageGRID networks
- Optional: Enabling node encryption if you plan to use an external KMS to encrypt appliance data.
- Optional: Changing the RAID mode.

### 4. Deploying the appliance as a Storage Node:

Task	Instructions
Deploying an appliance Storage Node in a new StorageGRID system	<a href="#">Deploying an appliance Storage Node</a>
Adding an appliance Storage Node to an existing StorageGRID system	Instructions for expanding a StorageGRID system
Deploying an appliance Storage Node as part of a Storage Node recovery operation	Instructions for recovery and maintenance

### Related information

[Preparing for installation](#)

[Installing the hardware](#)

[Configuring the hardware](#)

[Install VMware](#)

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

[Install Ubuntu or Debian](#)

[SG100 & SG1000 services appliances](#)

[Expand your grid](#)

[Maintain & recover](#)

[Administer StorageGRID](#)

## Preparing for installation

Preparing to install a StorageGRID appliance entails preparing the site and obtaining all required hardware, cables, and tools. You should also gather IP addresses and network information.

### Steps

- [Preparing the site \(SG5700\)](#)
- [Unpacking the boxes \(SG5700\)](#)
- [Obtaining additional equipment and tools \(SG5700\)](#)
- [Web browser requirements](#)
- [Reviewing appliance network connections](#)
- [Gathering installation information \(SG5700\)](#)

### Preparing the site (SG5700)

Before installing the appliance, you must make sure that the site and the cabinet or rack

you plan to use meet the specifications for a StorageGRID appliance.

### Steps

1. Confirm that the site meets the requirements for temperature, humidity, altitude range, airflow, heat dissipation, wiring, power, and grounding. See the NetApp Hardware Universe for more information.
2. If you are installing the SG5760 model, confirm that your location provides 240-volt AC power.
3. Obtain a 19-inch (48.3-cm) cabinet or rack to fit shelves of this size (without cables):

Appliance model	Height	Width	Depth	Maximum weight
SG5712 (12 drives)	3.41 in. (8.68 cm)	17.6 in. (44.7 cm)	21.1 in. (53.6 cm)	63.9 lb (29.0 kg)
SG5760 (60 drives)	6.87 in. (17.46 cm)	17.66 in. (44.86 cm)	38.25 in. (97.16 cm)	250 lb. (113 kg)

4. Install any required network switches. See the NetApp Interoperability Matrix Tool for compatibility information.

### Related information

[NetApp Hardware Universe](#)

[NetApp Interoperability Matrix Tool](#)


### Unpacking the boxes (SG5700)

Before installing the StorageGRID appliance, unpack all boxes and compare the contents to the items on the packing slip.

- **SG5712 appliance with 12 drives installed**



- **SG5760 appliance with no drives installed**




- **Front bezel for the appliance**




- Rail kit with instructions




- SG5760: Sixty drives



- SG5760: Handles



- SG5760: Back brackets and cage nuts for square-hole rack installation



#### Cables and connectors


The shipment for the StorageGRID appliance includes the following cables and connectors:

- Two power cords for your country



Your cabinet might have special power cords that you use instead of the power cords that ship with the appliance.

- **Optical cables and SFP transceivers**



Two optical cables for the FC interconnect ports


Eight SFP+ transceivers, compatible with both the four 16Gb/s FC interconnect ports and the four 10-GbE network ports

### Obtaining additional equipment and tools (SG5700)

Before installing the StorageGRID appliance, confirm you have all of the additional equipment and tools that you need.

You need the following additional equipment to install and configure the hardware:

- **Screwdrivers**




Phillips No. 2 screwdriver

Medium flat-blade screwdriver

- **ESD wrist strap**



- **Optical cables and SFP transceivers**



Optical cables for the 10/25-GbE ports you plan to use

Optional: SFP28 transceivers if you want to use 25-GbE link speed

- **Ethernet cables**



- **Service laptop**



Supported web browser

SSH client, such as PuTTY

1-Gb (RJ-45) Ethernet port

- **Optional tools**



Power drill with Phillips head bit

Flashlight

Mechanized lift for SG5760

### **Web browser requirements**

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## Reviewing appliance network connections

Before installing the StorageGRID appliance, you should understand which networks can be connected to the appliance and how the ports on each controller are used.

### StorageGRID appliance networks

When you deploy a StorageGRID appliance as a Storage Node in a StorageGRID grid, you can connect it to the following networks:


- **Grid Network for StorageGRID:** The Grid Network is used for all internal StorageGRID traffic. It provides connectivity between all nodes in the grid, across all sites and subnets. The Grid Network is required.
- **Admin Network for StorageGRID:** The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites. The Admin Network is optional.
- **Client Network for StorageGRID:** The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network provides client protocol access to the grid, so the Grid Network can be isolated and secured. The Client Network is optional.
- **Management network for SANtricity System Manager:** This network provides access to SANtricity System Manager on the E2800 controller, allowing you to monitor and manage the hardware components in the appliance. This management network can be the same as the Admin Network for StorageGRID, or it can be an independent management network.



For detailed information about StorageGRID networks, see the *Grid Primer*.

### StorageGRID appliance connections

When you install a StorageGRID appliance, you must connect the two controllers to each other and to the required networks. The figure shows the two controllers in the SG5760, with the E2800 controller on the top and the E5700SG controller on the bottom. In the SG5712, the E2800 controller is to the left of the E5700SG controller.



	Port	Type of port	Function
1	Two interconnect ports on each controller	16Gb/s FC optical SFP+	Connect the two controllers to each other.
2	Management port 1 on the E2800 controller	1-GbE (RJ-45)	Connects to the network where you access SANtricity System Manager. You can use the Admin Network for StorageGRID or an independent management network.
2	Management port 2 on the E2800 controller	1-GbE (RJ-45)	Reserved for technical support.
3	Management port 1 on the E5700SG controller	1-GbE (RJ-45)	Connects the E5700SG controller to the Admin Network for StorageGRID.
3	Management port 2 on the E5700SG controller	1-GbE (RJ-45)	<ul style="list-style-type: none"> <li>• Can be bonded with management port 1 if you want a redundant connection to the Admin Network.</li> <li>• Can be left unwired and available for temporary local access (IP 169.254.0.1).</li> <li>• During installation, can be used to connect the E5700SG controller to a service laptop if DHCP-assigned IP addresses are not available.</li> </ul>

	Port	Type of port	Function
4	10/25-GbE ports 1-4 on the E5700SG controller	10-GbE or 25-GbE  <b>Note:</b> The SFP+ transceivers included with the appliance support 10-GbE link speeds. If you want to use 25-GbE link speeds for the four network ports, you must provide SFP28 transceivers.	Connect to the Grid Network and the Client Network for StorageGRID. See “10/25-GbE port connections for the E5700SG controller.”

**Related information**

[Gathering installation information \(SG5700\)](#)

[Cabling the appliance \(SG5700\)](#)

[Port bond modes for E5700SG controller ports](#)

[Network guidelines](#)

[Install VMware](#)

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

[Install Ubuntu or Debian](#)

**Port bond modes for E5700SG controller ports**

When configuring network links for the E5700SG controller ports, you can use port bonding for the 10/25-GbE ports that connect to the Grid Network and optional Client Network, and the 1-GbE management ports that connect to the optional Admin Network. Port bonding helps protect your data by providing redundant paths between StorageGRID networks and the appliance.

**Related information**


[Configuring network links \(SG5700\)](#)

**Network bond modes for the 10/25-GbE ports**

The 10/25-GbE networking ports on the E5700SG controller support Fixed port bond mode or Aggregate port bond mode for the Grid Network and Client Network connections.

**Fixed port bond mode**

Fixed mode is the default configuration for the 10/25-GbE networking ports.



	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

When using Fixed port bond mode, you can use one of two network bond modes: Active-Backup or Link Aggregation Control Protocol (LACP).


- In Active-Backup mode (default), only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Port 4 provides a backup path for port 2 (Grid Network), and port 3 provides a backup path for port 1 (Client Network).
- In LACP mode, each pair of ports forms a logical channel between the controller and the network, allowing for higher throughput. If one port fails, the other port continues to provide the channel. Throughput is reduced, but connectivity is not impacted.



If you do not need redundant connections, you can use only one port for each network. However, be aware that an alarm will be raised in the Grid Manager after StorageGRID is installed, indicating that a cable is unplugged. You can safely acknowledge this alarm to clear it.

**Aggregate port bond mode**

Aggregate port bond mode significantly increases the throughput for each StorageGRID network and provides additional failover paths.



	Which ports are bonded
1	All connected ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

If you plan to use Aggregate port bond mode:

- You must use LACP network bond mode.
- You must specify a unique VLAN tag for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.
- The ports must be connected to switches that can support VLAN and LACP. If multiple switches are participating in the LACP bond, the switches must support multi-chassis link aggregation groups (MLAG), or equivalent.
- You must understand how to configure the switches to use VLAN, LACP, and MLAG, or equivalent.

If you do not want to use all four 10/25-GbE ports, you can use one, two, or three ports. Using more than one port maximizes the chance that some network connectivity will remain available if one of the 10/25-GbE ports fails.



If you choose to use fewer than four ports, be aware that one or more alarms will be raised in the Grid Manager after StorageGRID is installed, indicating that cables are unplugged. You can safely acknowledge the alarms to clear them.

### Network bond modes for the 1-GbE management ports


For the two 1-GbE management ports on the E5700SG controller, you can choose Independent network bond mode or Active-Backup network bond mode to connect to the optional Admin Network.

In Independent mode, only management port 1 is connected to the Admin Network. This mode does not provide a redundant path. Management port 2 is left unwired and available for temporary local connections (use IP address 169.254.0.1)

In Active-Backup mode, both management ports 1 and 2 are connected to the Admin Network. Only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Bonding these two physical ports into one logical management port provides a redundant path to the Admin Network.



If you need to make a temporary local connection to the E5700SG controller when the 1-GbE management ports are configured for Active-Backup mode, remove the cables from both management ports, plug your temporary cable into management port 2, and access the appliance using IP address 169.254.0.1.



### Gathering installation information (SG5700)

As you install and configure the StorageGRID appliance, you must make decisions and gather information about Ethernet switch ports, IP addresses, and port and network bond modes.

#### About this task

You can use the following tables to record the required information for each network you connect to the appliance. These values are required to install and configure the hardware.

**Information needed to connect to SANtricity System Manager on the E2800 controller**

You must connect the E2800 controller to the management network you will use for SANtricity System Manager.

Information needed	Your value
Ethernet switch port you will connect to management port 1	
MAC address for management port 1 (printed on a label near port P1)	
DHCP-assigned IP address for management port 1, if available after power on  <b>Note:</b> If the network you will connect to the E2800 controller includes a DHCP server, the network administrator can use the MAC address to determine the IP address that was assigned by the DHCP server.	
Speed and duplex mode  <b>Note:</b> You must make sure the Ethernet switch for the SANtricity System Manager management network is set to autonegotiate.	Must be: <ul style="list-style-type: none"> <li>• Autonegotiate (default)</li> </ul>
IP address format	Choose one: <ul style="list-style-type: none"> <li>• IPv4</li> <li>• IPv6</li> </ul>
Static IP address you plan to use for the appliance on the management network	For IPv4: <ul style="list-style-type: none"> <li>• IPv4 address:</li> <li>• Subnet mask:</li> <li>• Gateway:</li> </ul> For IPv6: <ul style="list-style-type: none"> <li>• IPv6 address:</li> <li>• Routable IP address:</li> <li>• E2800 controller router IP address:</li> </ul>

### Information needed to connect the E5700SG controller to the Admin Network

The Admin Network for StorageGRID is an optional network, used for system administration and maintenance. The appliance connects to the Admin Network using the 1-GbE management ports on the E5700SG controller.

Information needed	Your value
Admin Network enabled	Choose one: <ul style="list-style-type: none"><li>• No</li><li>• Yes (default)</li></ul>
Network bond mode	Choose one: <ul style="list-style-type: none"><li>• Independent</li><li>• Active-Backup</li></ul>
Switch port for port 1	
Switch port for port 2 (Active-Backup network bond mode only)	
DHCP-assigned IP address for management port 1, if available after power on  <b>Note:</b> If the Admin Network includes a DHCP server, the E5700SG controller displays the DHCP-assigned IP address on its seven-segment display after it boots up. You can also determine the DHCP-assigned IP address by using the MAC address to look up the assigned IP.	<ul style="list-style-type: none"><li>• IPv4 address (CIDR):</li><li>• Gateway:</li></ul>
Static IP address you plan to use for the appliance Storage Node on the Admin Network  <b>Note:</b> If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"><li>• IPv4 address (CIDR):</li><li>• Gateway:</li></ul>
Admin Network subnets (CIDR)	

### Information needed to connect and configure the 10/25-GbE ports on the E5700SG controller

The four 10/25-GbE ports on the E5700SG controller connect to the StorageGRID Grid Network and Client Network.



See "10/25-GbE port connections for the E5700SG controller" for more information about the options for these ports.

Information needed	Your value
Link speed  <b>Note:</b> If you select 25 GbE, you must install SPF28 transceivers. Auto-negotiation is not supported, so you must also configure the ports and the connected switches for 25GbE.	Choose one: <ul style="list-style-type: none"> <li>• 10 GbE (default)</li> <li>• 25 GbE</li> </ul>
Port bond mode	Choose one: <ul style="list-style-type: none"> <li>• Fixed (default)</li> <li>• Aggregate</li> </ul>
Switch port for port 1 (Client Network)	
Switch port for port 2 (Grid Network)	
Switch port for port 3 (Client Network)	
Switch port for port 4 (Grid Network)	

**Information needed to connect the E5700SG controller to the Grid Network**

The Grid Network for StorageGRID is a required network, used for all internal StorageGRID traffic. The appliance connects to the Grid Network using the 10/25-GbE ports on the E5700SG controller.



See "10/25-GbE port connections for the E5700SG controller" for more information about the options for these ports.

Information needed	Your value
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>• Active-Backup (default)</li> <li>• LACP (802.3ad)</li> </ul>
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>
VLAN tag(if VLAN tagging is enabled)	Enter a value between 0 and 4095:

Information needed	Your value
DHCP-assigned IP address for the Grid Network, if available after power on  <b>Note:</b> If the Grid Network includes a DHCP server, the E5700SG controller displays the DHCP-assigned IP address for the Grid Network on its seven-segment display after it boots up.	<ul style="list-style-type: none"> <li>IPv4 address (CIDR):</li> <li>Gateway:</li> </ul>
Static IP address you plan to use for the appliance Storage Node on the Grid Network  <b>Note:</b> If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> <li>IPv4 address (CIDR):</li> <li>Gateway:</li> </ul>
Grid Network subnets (CIDR)  <b>Note:</b> If the Client Network is not enabled, the default route on the controller will use the gateway specified here.	

**Information needed to connect the E5700SG controller to the Client Network**

The Client Network for StorageGRID is an optional network, typically used to provide client protocol access to the grid. The appliance connects to the Client Network using the 10/25-GbE ports on the E5700SG controller.



See "10/25-GbE port connections for the E5700SG controller" for more information about the options for these ports.

Information needed	Your value
Client Network enabled	Choose one: <ul style="list-style-type: none"> <li>No (default)</li> <li>Yes</li> </ul>
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>Active-Backup (default)</li> <li>LACP (802.3ad)</li> </ul>
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> <li>No (default)</li> <li>Yes</li> </ul>

Information needed	Your value
VLAN tag  (if VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Client Network, if available after power on	<ul style="list-style-type: none"> <li>IPv4 address (CIDR):</li> <li>Gateway:</li> </ul>
Static IP address you plan to use for the appliance Storage Node on the Client Network  <b>Note:</b> If the Client Network is enabled, the default route on the controller will use the gateway specified here.	<ul style="list-style-type: none"> <li>IPv4 address (CIDR):</li> <li>Gateway:</li> </ul>

### Related information

[Reviewing appliance network connections](#)

[Port bond modes for E5700SG controller ports](#)

[Configuring the hardware](#)

## Installing the hardware

Hardware installation entails installing the appliance into a cabinet or rack, connecting the cables, and applying power.

### Steps

- [Registering the hardware](#)
- [Installing the appliance in a cabinet or rack \(SG5700\)](#)
- [Cabling the appliance \(SG5700\)](#)
- [Connecting power cords and applying power \(SG5700\)](#)
- [Viewing SG5700 boot-up status codes](#)

### Registering the hardware

Registering the appliance hardware provides support benefits.

### Steps

1. Locate the chassis serial number.

You can find the number on the packing slip, in your confirmation email, or on the appliance after you unpack it.



2. Go to the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com).

### 3. Determine whether you need to register the hardware:

If you are a...	Follow these steps...
Existing NetApp customer	<ol style="list-style-type: none"><li>Sign in with your username and password.</li><li>Select <b>Products &gt; My Products</b>.</li><li>Confirm that the new serial number is listed.</li><li>If it is not, follow the instructions for new NetApp customers.</li></ol>
New NetApp customer	<ol style="list-style-type: none"><li>Click <b>Register Now</b>, and create an account.</li><li>Select <b>Products &gt; Register Products</b>.</li><li>Enter the product serial number and requested details.</li></ol> <p>After your registration is approved, you can download any required software. The approval process might take up to 24 hours.</p>

### Installing the appliance in a cabinet or rack (SG5700)

You must install rails in your cabinet or rack and then slide the appliance onto the rails. If you have an SG5760, you must also install the drives after installing the appliance.

#### What you'll need

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.
- You have the *Installation and Setup Instructions* for the appliance.



Install hardware from the bottom of the rack or cabinet or rack up to prevent the equipment from tipping over.



The SG5712 weighs approximately 64 lb (29 kg) when fully loaded with drives. Two people or a mechanized lift are required to safely move the SG5712.



The SG5760 weighs approximately 132 lb (60 kg) with no drives installed. Four people or a mechanized lift are required to safely move an empty SG5760.



To avoid damaging the hardware, never move an SG5760 if drives are installed. You must remove all drives before moving the shelf.

#### Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.
2. If you have an SG5760, follow these steps to prepare for moving the appliance.

- a. Remove the outer packing box. Then, fold down the flaps on the inner box.
- b. If you are lifting the SG5760 by hand, attach the four handles to the sides of the chassis.

You remove these handles as you slide the appliance onto the rails.


3. See the *Installation and Setup Instructions*, and slide the appliance in the cabinet or rack.
4. See the *Installation and Setup Instructions*, and secure the appliance to the cabinet or rack.

If you have an SG5760, use the back brackets to secure the appliance to the rear of the rack or cabinet. Use the cage nuts if your rack or cabinet has square holes.

5. If you have an SG5760, install 12 drives in each of the 5 drive drawers.

You must install all 60 drives to ensure correct operation.

- a. Put on the ESD wristband, and remove the drives from their packaging.
- b. Release the levers on the top drive drawer, and slide the drawer out using the levers.
- c. Raise the drive handle to vertical, and align the buttons on the drive with the notches on the drawer.



- d. Pressing gently on the top of the drive, rotate the drive handle down until the drive snaps into place.
- e. After installing the first 12 drives, slide the drawer back in by pushing on the center and closing both levers gently.
- f. Repeat these steps for the other four drawers.

6. Attach the front bezel.

### Cabling the appliance (SG5700)

You must connect the two controllers to each other, connect the management ports on each controller, and connect the 10/25-GbE ports on the E5700SG controller to the Grid Network and optional Client Network for StorageGRID.

#### What you'll need

- You have unpacked the following items, which are included with the appliance:
  - Two power cords.


- Two optical cables for the FC interconnect ports on the controllers.
- Eight SFP+ transceivers, which support either 10-GbE or 16-Gbps FC. The transceivers can be used with the two interconnect ports on both controllers and with the four 10/25-GbE network ports on the E5700SG controller, assuming you want the network ports to use a 10-GbE link speed.
- You have obtained the following items, which are not included with the appliance:
  - One to four optical cables for the 10/25-GbE ports you plan to use.
  - One to four SFP28 transceivers, if you plan to use 25-GbE link speed.
  - Ethernet cables for connecting the management ports.



**Risk of exposure to laser radiation** — Do not disassemble or remove any part of an SFP transceiver. You might be exposed to laser radiation.

### About this task

The figure shows the two controllers in the SG5760, with the E2800 controller on the top and the E5700SG controller on the bottom. In the SG5712, the E2800 controller is to the left of the E5700SG controller when viewed from the back.



	Port	Type of port	Function
1	Two interconnect ports on each controller	16Gb/s FC optical SFP+	Connect the two controllers to each other.
2	Management port 1 on the E2800 controller	1-GbE (RJ-45)	Connects to the network where you access SANtricity System Manager. You can use the Admin Network for StorageGRID or an independent management network.
2	Management port 2 on the E2800 controller	1-GbE (RJ-45)	Reserved for technical support.

	Port	Type of port	Function
3	Management port 1 on the E5700SG controller	1-GbE (RJ-45)	Connects the E5700SG controller to the Admin Network for StorageGRID.
3	Management port 2 on the E5700SG controller	1-GbE (RJ-45)	<ul style="list-style-type: none"> <li>• Can be bonded with management port 1 if you want a redundant connection to the Admin Network.</li> <li>• Can be left unwired and available for temporary local access (IP 169.254.0.1).</li> <li>• During installation, can be used to connect the E5700SG controller to a service laptop if DHCP-assigned IP addresses are not available.</li> </ul>
4	10/25-GbE ports 1-4 on the E5700SG controller	10-GbE or 25-GbE  <b>Note:</b> The SFP+ transceivers included with the appliance support 10-GbE link speeds. If you want to use 25-GbE link speeds for the four network ports, you must provide SFP28 transceivers.	Connect to the Grid Network and the Client Network for StorageGRID. See “10/25-GbE port connections for the E5700SG controller.”

### Steps

1. Connect the E2800 controller to the E5700SG controller, using two optical cables and four of the eight SFP+ transceivers.

Connect this port...	To this port...
Interconnect port 1 on the E2800 controller	Interconnect port 1 on the E5700SG controller
Interconnect port 2 on the E2800 controller	Interconnect port 2 on the E5700SG controller

2. Connect management port 1 (P1) on the E2800 controller (the RJ-45 port on the left) to the management network for SANtricity System Manager, using an Ethernet cable.

Do not use management port 2 (P2) on the E2800 controller (the RJ-45 port on the right). This port is reserved for technical support.

3. If you plan to use the Admin Network for StorageGRID, connect management port 1 on the E5700SG controller (the RJ-45 port on the left) to the Admin Network, using an Ethernet cable.

If you plan to use active-backup network bond mode for the Admin Network, connect management port 2 on the E5700SG controller (the RJ-45 port on the right) to the Admin Network, using an Ethernet cable.

4. Connect the 10/25-GbE ports on the E5700SG controller to the appropriate network switches, using optical cables and SFP+ or SFP28 transceivers.



All ports must use the same link speed. Install SFP+ transceivers if you plan to use 10-GbE link speeds. Install SFP28 transceivers if you plan to use 25-GbE link speeds.

- If you plan to use Fixed port bond mode (default), connect the ports to the StorageGRID Grid and Client Networks, as shown in the table.

Port	Connects to...
Port 1	Client Network (optional)
Port 2	Grid Network
Port 3	Client Network (optional)
Port 4	Grid Network

- If you plan to use the Aggregate port bond mode, connect one or more of the network ports to one or more switches. You should connect at least two of the four ports to avoid having a single point of failure. If you use more than one switch for a single LACP bond, the switches must support MLAG or equivalent.

## Related information

[Accessing the StorageGRID Appliance Installer](#)

[Port bond modes for E5700SG controller ports](#)

## Connecting power cords and applying power (SG5700)

When you apply power to the appliance, both controllers boot up.

### What you'll need

Both appliance power switches must be off before connecting power.



**Risk of electrical shock** — Before connecting the power cords, make sure that the two power switches on the appliance are off.

### Steps

1. Confirm that the two power switches on the appliance are off.

2. Connect the two power cords to the appliance.
3. Connect the two power cords to different power distribution units (PDUs) in the cabinet or rack.
4. Turn on the two power switches on the appliance.
  - Do not turn off the power switches during the power-on process.
  - The fans are very loud when they first start up. The loud noise during start-up is normal.
5. After the controllers have booted up, check their seven-segment displays.

### Viewing SG5700 boot-up status codes

The seven-segment displays on each controller show status and error codes as the appliance powers up.

#### About this task


The E2800 controller and the E5700SG controller display different statuses and error codes.

To understand what these codes mean, see the following resources:

Controller	Reference
E2800 controller	<p><i>E5700 and E2800 System Monitoring Guide</i></p> <p><b>Note:</b> The codes listed for the E-Series E5700 controller do not apply to the E5700SG controller in the appliance.</p>
E5700SG controller	“Status indicators on the E5700SG controller”

#### Steps

1. During boot-up, monitor progress by viewing the codes shown on the seven-segment displays.
  - The seven-segment display on the E2800 controller shows the repeating sequence **OS**, **Sd**, **blank** to indicate that it is performing start-of-day processing.
  - The seven-segment display on the E5700SG controller shows a sequence of codes, ending with **AA** and **FF**.
2. After the controllers have booted up, confirm the seven-segment displays show the following:



Controller	Seven-segment display
E2800 controller	Shows 99, which is the default ID for an E-Series controller shelf.
E5700SG controller	<p>Shows <b>HO</b>, followed by a repeating sequence of two numbers.</p> <div style="border: 1px solid gray; padding: 5px; margin: 10px 0;"> <p>HO -- IP address for Admin Network -- IP address for Grid Network HO</p> </div> <p>In the sequence, the first set of numbers is the DHCP-assigned IP address for the controller's management port 1. This address is used to connect the controller to the Admin Network for StorageGRID. The second set of numbers is the DHCP-assigned IP address used to connect the appliance to the Grid Network for StorageGRID.</p> <p><b>Note:</b> If an IP address could not be assigned using DHCP, 0.0.0.0 is displayed.</p>

3. If the seven-segment displays show other values, see “Troubleshooting the hardware installation,” and confirm you completed the installation steps correctly. If you are unable to resolve the problem, contact technical support.


#### Related information

[Status indicators on the E5700SG controller](#)

Status indicators on the E5700SG controller

The seven-segment display and the LEDs on the E5700SG controller show status and error codes while the appliance powers up and while the hardware is initializing. You can use these displays to determine status and troubleshoot errors.

After the StorageGRID Appliance Installer has started, you should periodically review the status indicators on the E5700SG controller.



	Display	Description
1	Attention LED	Amber: The controller is faulty and requires operator attention, or the installation script was not found.  Off: The controller is operating normally.
2	Seven-segment display	Shows a diagnostic code  Seven-segment display sequences enable you to understand errors and the operational state of the appliance.
3	Expansion Port Attention LEDs	Amber: These LEDs are always amber (no link established) because the appliance does not use the expansion ports.
4	Host Port Link Status LEDs	Green: The link is up.  Off: The link is down.
5	Ethernet Link State LEDs	Green: A link is established.  Off: No link is established.

	Display	Description
6	Ethernet Activity LEDs	<p>Green: The link between the management port and the device to which it is connected (such as an Ethernet switch) is up.</p> <p>Off: There is no link between the controller and the connected device.</p> <p>Blinking Green: There is Ethernet activity.</p>

### General boot-up codes

During boot-up or after a hard reset of the appliance, the following occurs:

1. The seven-segment display on the E5700SG controller shows a general sequence of codes that is not specific to the controller. The general sequence ends with the codes AA and FF.
2. Boot-up codes that are specific to the E5700SG controller appear.

### E5700SG controller boot-up codes

During a normal boot-up of the appliance, the seven-segment display on the E5700SG controller shows the following codes in the order listed:

Code	Indicates
HI	The master boot script has started.
PP	The system is checking to see if the FPGA needs to be updated.
HP	The system is checking to see if the 10/25-GbE controller firmware needs to be updated.
RB	The system is rebooting after applying firmware updates.
FP	The hardware subsystem firmware update checks have been completed. Inter-controller communication services are starting.
HE	<p>The system is awaiting connectivity with the E2800 controller and synchronizing with the SANtricity operating system.</p> <p><b>Note:</b> If this boot procedure does not progress past this stage, check the connections between the two controllers.</p>
HC	The system is checking for existing StorageGRID installation data.
HO	The StorageGRID Appliance Installer is running.
HA	StorageGRID is running.

## E5700SG controller error codes

These codes represent error conditions that might be shown on the E5700SG controller as the appliance boots up. Additional two-digit hexadecimal codes are displayed if specific low-level hardware errors occur. If any of these codes persists for more than a second or two, or if you are unable to resolve the error by following one of the prescribed troubleshooting procedures, contact technical support.

Code	Indicates
22	No master boot record found on any boot device.
23	The internal flash disk is not connected.
2A, 2B	Stuck bus, unable to read DIMM SPD data.
40	Invalid DIMMs.
41	Invalid DIMMs.
42	Memory test failed.
51	SPD reading failure.
92 to 96	PCI bus initialization.
A0 to A3	SATA drive initialization.
AB	Alternate boot code.
AE	Booting OS.
EA	DDR4 training failed.
E8	No memory installed.
EU	The installation script was not found.
EP	Installation or communication with the E2800 controller has failed.

### Related information

[Troubleshooting the hardware installation](#)

[NetApp Support](#)

## Configuring the hardware

After applying power to the appliance, you must configure SANtricity System Manager,

which is the software you will use to monitor the hardware. You must also configure the network connections that will be used by StorageGRID.

### Steps

- [Configuring StorageGRID connections](#)
- [Accessing and Configuring SANtricity System Manager](#)
- [Optional: Enabling node encryption](#)
- [Optional: Changing the RAID mode \(SG5760 only\)](#)
- [Optional: Remapping network ports for the appliance](#)

### Configuring StorageGRID connections

Before you can deploy a StorageGRID appliance as a Storage Node in a StorageGRID grid, you must configure the connections between the appliance and the networks you plan to use. You can configure networking by browsing to the StorageGRID Appliance Installer, which is included on the E5700SG controller (the compute controller in the appliance).

### Steps

- [Accessing the StorageGRID Appliance Installer](#)
- [Verifying and upgrading the StorageGRID Appliance Installer version](#)
- [Configuring network links \(SG5700\)](#)
- [Setting the IP configuration](#)
- [Verifying network connections](#)
- [Verifying port-level network connections](#)

### Accessing the StorageGRID Appliance Installer

You must access the StorageGRID Appliance Installer to configure the connections between the appliance and the three StorageGRID networks: the Grid Network, the Admin Network (optional), and the Client Network (optional).

### What you'll need

- You are using a supported web browser.
- The appliance is connected to all of the StorageGRID networks you plan to use.
- You know the IP address, gateway, and subnet for the appliance on these networks.
- You have configured the network switches you plan to use.

### About this task

When you first access the StorageGRID Appliance Installer, you can use the DHCP-assigned IP address for the Admin Network (assuming the appliance is connected to the Admin Network) or the DHCP-assigned IP address for the Grid Network. Using the IP address for the Admin Network is preferred. Otherwise, if you access the StorageGRID Appliance Installer using the DHCP address for the Grid Network, you might lose connection with the StorageGRID Appliance Installer when you change link settings and when you enter a static IP.

## Steps

1. Obtain the DHCP address for the appliance on the Admin Network (if it is connected) or the Grid Network (if the Admin Network is not connected).

You can do either of the following:

- Look at the seven-segment display on the E5700SG controller. If management port 1 and 10/25-GbE ports 2 and 4 on the E5700SG controller are connected to networks with DHCP servers, the controller attempts to obtain dynamically assigned IP addresses when you power on the enclosure. After the controller has completed the power-on process, its seven-segment display shows **HO**, followed by a repeating sequence of two numbers.

```
HO -- IP address for Admin Network -- IP address for Grid Network HO
```

In the sequence:

- The first set of numbers is the DHCP address for the appliance Storage Node on the Admin Network, if it is connected. This IP address is assigned to management port 1 on the E5700SG controller.
- The second set of numbers is the DHCP address for the appliance Storage Node on the Grid Network. This IP address is assigned to 10/25-GbE ports 2 and 4 when you first apply power to the appliance.



If an IP address could not be assigned using DHCP, 0.0.0.0 is displayed.

- Provide the MAC address for management port 1 to your network administrator, so they can look up the DHCP address for this port on the Admin Network. The MAC address is printed on a label on the E5700SG controller, next to the port.
2. If you were able to obtain either of the DHCP addresses:
    - a. Open a web browser on the service laptop.
    - b. Enter this URL for the StorageGRID Appliance Installer:  
**`https://E5700SG_Controller_IP:8443`**
  - c. If you are prompted with a security alert, view and install the certificate using the browser's installation wizard.

The alert will not appear the next time you access this URL.

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected to StorageGRID networks. Error messages might appear that will be resolved in later steps.

## Home

**i** The installation is ready to be started. Review the settings below, and then click Start Installation.

### This Node

Node type

Storage

Node name

MM-2-108-SGA-lab25

Cancel

Save

### Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

172.16.1.178

Connection state

Connection to 172.16.1.178 ready

Cancel

Save

### Installation

Current state

Ready to start installation of MM-2-108-SGA-lab25 into grid with Admin Node 172.16.1.178 running StorageGRID 11.2.0, using StorageGRID software downloaded from the Admin Node.

Start Installation

3. If the E5700SG controller could not acquire an IP address using DHCP:
  - a. Connect the service laptop to management port 2 on the E5700SG controller, using an Ethernet cable.



- b. Open a web browser on the service laptop.

c. Enter this URL for the StorageGRID Appliance Installer:

**https://169.254.0.1:8443**

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected.



If you cannot access the Home page over a link-local connection, configure the service laptop IP address as 169.254.0.2, and try again.

4. Review any messages displayed on the Home page and configure the link configuration and the IP configuration, as required.

## Related information

[Web browser requirements](#)

## Verifying and upgrading the StorageGRID Appliance Installer version

The StorageGRID Appliance Installer version on the appliance must match the software version installed on your StorageGRID system to ensure that all StorageGRID features are supported.

### What you'll need

You have accessed the StorageGRID Appliance Installer.

### About this task

StorageGRID appliances come from the factory preinstalled with the StorageGRID Appliance Installer. If you are adding an appliance to a recently upgraded StorageGRID system, you might need to manually upgrade the StorageGRID Appliance Installer before installing the appliance as a new node.

The StorageGRID Appliance Installer automatically upgrades when you upgrade to a new StorageGRID version. You do not need to upgrade the StorageGRID Appliance Installer on installed appliance nodes. This procedure is only required when you are installing an appliance that contains an earlier version of the StorageGRID Appliance Installer.

### Steps

1. From the StorageGRID Appliance Installer, select **Advanced > Upgrade Firmware**.
2. Compare the Current Firmware version to the software version installed on your StorageGRID system (from the Grid Manager select **Help > About**).

The second digit in the two versions should match. For example, if your StorageGRID system is running version 11.5.x.y, the StorageGRID Appliance Installer version should be 3.5.z.

3. If the appliance has a down-level version of the StorageGRID Appliance Installer, go to the NetApp Downloads page for StorageGRID.

[NetApp Downloads: StorageGRID](#)

Sign in with the username and password for your NetApp account.

4. Download the appropriate version of the **Support file for StorageGRID Appliances** and the corresponding checksum file.

The Support file for StorageGRID Appliances file is a .zip archive that contains the current and previous firmware versions for all StorageGRID appliance models, in subdirectories for each controller type.

After downloading the Support file for StorageGRID Appliances file, extract the .zip archive and see the README file for important information about installing the StorageGRID Appliance Installer.

5. Follow the instructions on the Upgrade Firmware page of the StorageGRID Appliance Installer to perform these steps:
  - a. Upload the appropriate support file (firmware image) for your controller type and the checksum file.
  - b. Upgrade the inactive partition.
  - c. Reboot and swap partitions.
  - d. Upgrade the second partition.

## Related information

[Accessing the StorageGRID Appliance Installer](#)

### Configuring network links (SG5700)

You can configure network links for the ports used to connect the appliance to the Grid Network, the Client Network, and the Admin Network. You can set the link speed as well as the port and network bond modes.

### What you'll need

If you plan to use the 25-GbE link speed for the 10/25-GbE ports:


- You have installed SFP28 transceivers in the ports you plan to use.
- You have connected the ports to switches that can support these features.
- You understand how to configure the switches to use this higher speed.

If you plan to use Aggregate port bond mode, LACP network bond mode, or VLAN tagging for the 10/25-GbE ports:

- You have connected the ports on the appliance to switches that can support VLAN and LACP.
- If multiple switches are participating in the LACP bond, the switches support multi-chassis link aggregation groups (MLAG), or equivalent.
- You understand how to configure the switches to use VLAN, LACP, and MLAG or equivalent.
- You know the unique VLAN tag to use for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.
- If you plan to use Active-Backup mode for the Admin Network, you have connected Ethernet cables to both management ports on the controller.


### About this task

This figure shows how the four 10/25-GbE ports are bonded in Fixed port bond mode (default configuration).



	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

This figure shows how the four 10/25-GbE ports are bonded in Aggregate port bond mode.



	Which ports are bonded
1	All four ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

The table summarizes the options for configuring the four 10/25-GbE ports. The default settings are shown in bold. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

• **Fixed (default) port bond mode**

Network bond mode	Client Network disabled (default)	Client Network enabled
<b>Active-Backup (default)</b>	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an active-backup bond for the Grid Network.</li> <li>• Ports 1 and 3 are not used.</li> <li>• A VLAN tag is optional.</li> </ul>	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an active-backup bond for the Grid Network.</li> <li>• Ports 1 and 3 use an active-backup bond for the Client Network.</li> <li>• VLAN tags can be specified for both networks for the convenience of the network administrator.</li> </ul>


Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad)	<ul style="list-style-type: none"> <li>Ports 2 and 4 use an LACP bond for the Grid Network.</li> <li>Ports 1 and 3 are not used.</li> <li>A VLAN tag is optional.</li> </ul>	<ul style="list-style-type: none"> <li>Ports 2 and 4 use an LACP bond for the Grid Network.</li> <li>Ports 1 and 3 use an LACP bond for the Client Network.</li> <li>VLAN tags can be specified for both networks for the convenience of the network administrator.</li> </ul>

• **Aggregate port bond mode**

Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad) only	<ul style="list-style-type: none"> <li>Ports 1-4 use a single LACP bond for the Grid Network.</li> <li>A single VLAN tag identifies Grid Network packets.</li> </ul>	<ul style="list-style-type: none"> <li>Ports 1-4 use a single LACP bond for the Grid Network and the Client Network.</li> <li>Two VLAN tags allow Grid Network packets to be segregated from Client Network packets.</li> </ul>

See the information about 10/25-GbE port connections for the E5700SG controller for more information about port bond and network bond modes.

This figure shows how the two 1-GbE management ports on the E5700SG controller are bonded in Active-Backup network bond mode for the Admin Network.




**Steps**

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Link Configuration**.

The Network Link Configuration page displays a diagram of your appliance with the network and management ports numbered.

## Network Link Configuration



⚠ You might lose your connection if you make changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL using one of the other IP addresses assigned to the appliance.

The Link Status table lists the link state (up/down) and speed (1/10/25/40/100 Gbps) of the numbered ports.

### Link Status

Link	State	Speed (Gbps)
1	Up	25
2	Up	25
3	Up	25
4	Up	25
5	Up	1
6	Up	1

The first time you access this page:

- **Link Speed** is set to **10GbE**.
- **Port bond mode** is set to **Fixed**.
- **Network bond mode** for the Grid Network is set to **Active-Backup**.
- The **Admin Network** is enabled, and the network bond mode is set to **Independent**.
- The **Client Network** is disabled.

## Link Settings

Link speed

Port bond mode  Fixed  Aggregate

Choose Fixed port bond mode if you want to use ports 2 and 4 for the Grid Network and ports 1 and 3 for the Client Network (if enabled). Choose Aggregate port bond mode if you want all connected ports to share a single LACP bond for both the Grid and Client Networks.

## Grid Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

Enable VLAN (802.1q) tagging

MAC Addresses 50:6b:4b:42:d7:00 50:6b:4b:42:d7:01 50:6b:4b:42:d7:24 50:6b:4b:42:d7:25

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

## Admin Network

Enable network

Network bond mode  Independent  Active-Backup

Connect the Admin Network to port 5. Leave port 6 unconnected. If necessary, you can make a temporary direct Ethernet connection to port 6 and use link-local IP address 169.254.0.1 for access.

MAC Addresses d8:c4:97:2a:e4:95

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

## Client Network

Enable network

Enabling the Client Network causes the default gateway for this node to move to the Client Network. Before enabling the Client Network, ensure that you've added all necessary subnets to the Grid Network Subnet List. Otherwise, the connection to the node might be lost.

2. If you plan to use the 25-GbE link speed for the 10/25 GbE ports, select **25GbE** from the Link speed drop-down list.

The network switches you are using for the Grid Network and the Client Network must also support and be configured for this speed. SFP28 transceivers must be installed in the ports.

3. Enable or disable the StorageGRID networks you plan to use.

The Grid Network is required. You cannot disable this network.

- a. If the appliance is not connected to the Admin Network, unselect the **Enable network** check box for the Admin Network.

#### **Admin Network**

---

Enable network

- b. If the appliance is connected to the Client Network, select the **Enable network** check box for the Client Network.

The Client Network settings for the 10/25-GbE ports are now shown.

4. Refer to the table, and configure the port bond mode and the network bond mode.

The example shows:

- **Aggregate** and **LACP** selected for the Grid and the Client networks. You must specify a unique VLAN tag for each network. You can select values between 0 and 4095.
- **Active-Backup** selected for the Admin Network.

### Link Settings

Link speed

Port bond mode  Fixed  Aggregate

Choose Fixed port bond mode if you want to use ports 2 and 4 for the Grid Network and ports 1 and 3 for the Client Network (if enabled). Choose Aggregate port bond mode if you want all connected ports to share a single LACP bond for both the Grid and Client Networks.

### Grid Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

### Admin Network

Enable network

Network bond mode  Independent  Active-Backup

Connect the Admin Network to ports 5 and 6. If necessary, you can make a temporary direct Ethernet connection by disconnecting ports 5 and 6, then connecting to port 6 and using link-local IP address 169.254.0.1 for access.

### Client Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

5. When you are satisfied with your selections, click **Save**.



You might lose your connection if you made changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:

**`https://E5700SG_Controller_IP:8443`**

### Related information

[Port bond modes for E5700SG controller ports](#)

### Setting the IP configuration

You use the StorageGRID Appliance Installer to configure the IP addresses and routing information used for the appliance Storage Node on the StorageGRID Grid, Admin, and

## Client Networks.

### About this task

You must either assign a static IP for the appliance on each connected network or assign a permanent lease for the address on the DHCP server.

If you want to change the link configuration, see the instructions for changing the link configuration of the E5700SG controller.

### Steps

1. In the StorageGRID Appliance Installer, select **Configure Networking > IP Configuration**.

The IP Configuration page appears.

2. To configure the Grid Network, select either **Static** or **DHCP** in the **Grid Network** section of the page.

## Grid Network

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

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

 All required Grid Network subnets must also be defined in the Grid Network Subnet List on the Primary Admin Node before starting installation.

Subnets (CIDR)  



MTU  

3. If you selected **Static**, follow these steps to configure the Grid Network:

- a. Enter the static IPv4 address, using CIDR notation.
- b. Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

- c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

**`https://services_appliance_IP:8443`**

e. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

f. Click **Save**.

4. If you selected **DHCP**, follow these steps to configure the Grid Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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

d. Click **Save**.

5. To configure the Admin Network, select either **Static** or **DHCP** in the Admin Network section of the page.



To configure the Admin Network, you must enable the Admin Network on the Link Configuration page.

### Admin Network

The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites.

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

Subnets (CIDR)  **+**

MTU

6. If you selected **Static**, follow these steps to configure the Admin Network:

a. Enter the static IPv4 address, using CIDR notation, for Management Port 1 on the appliance.

Management Port 1 is the left of the two 1-GbE RJ45 ports on the right end of the appliance.

b. Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as

9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

**`https://services_appliance:8443`**

e. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route cannot be made to use the Admin Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

f. Click **Save**.

7. If you selected **DHCP**, follow these steps to configure the Admin Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route cannot be made to use the Admin Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

8. To configure the Client Network, select either **Static** or **DHCP** in the **Client Network** section of the page.



To configure the Client Network, you must enable the Client Network on the Link Configuration page.

## Client Network

The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network enables grid nodes to communicate with any subnet reachable through the Client Network gateway. The Client Network does not become operational until you complete the StorageGRID configuration steps.

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

MTU

9. If you selected **Static**, follow these steps to configure the Client Network:

- Enter the static IPv4 address, using CIDR notation.
- Click **Save**.
- Confirm that the IP address for the Client Network gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and cannot be moved to another interface while the Client Network is enabled.

- If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

- Click **Save**.

10. If you selected **DHCP**, follow these steps to configure the Client Network:

- After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address** and **Gateway** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and cannot be moved to another interface while the Client Network is enabled.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

## Related information

[Changing the link configuration of the E5700SG controller](#)

## Verifying network connections

You should confirm you can access the StorageGRID networks you are using from the appliance. To validate routing through network gateways, you should test connectivity between the StorageGRID Appliance Installer and IP addresses on different subnets. You can also verify the MTU setting.

## Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Ping and MTU Test**.

The Ping and MTU Test page appears.

### Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

#### Ping and MTU Test

Network	<input type="text" value="Grid"/>
Destination IPv4 Address or FQDN	<input type="text"/>
Test MTU	<input type="checkbox"/>
<input type="button" value="Test Connectivity"/>	

2. From the **Network** drop-down box, select the network you want to test: Grid, Admin, or Client.

3. Enter the IPv4 address or fully qualified domain name (FQDN) for a host on that network.

For example, you might want to ping the gateway on the network or the primary Admin Node.

4. Optionally, select the **Test MTU** check box to verify the MTU setting for the entire path through the network to the destination.

For example, you can test the path between the appliance node and a node at a different site.

5. Click **Test Connectivity**.

If the network connection is valid, the "Ping test passed" message appears, with the ping command output listed.

### Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

#### Ping and MTU Test

Network	<input type="text" value="Grid"/>
Destination IPv4 Address or FQDN	<input type="text" value="10.96.104.223"/>
Test MTU	<input checked="" type="checkbox"/>
<input type="button" value="Test Connectivity"/>	

Ping test passed

#### Ping command output

```
PING 10.96.104.223 (10.96.104.223) 1472(1500) bytes of data.  
1480 bytes from 10.96.104.223: icmp_seq=1 ttl=64 time=0.318 ms  
  
--- 10.96.104.223 ping statistics ---  
1 packets transmitted, 1 received, 0% packet loss, time 0ms  
rtt min/avg/max/mdev = 0.318/0.318/0.318/0.000 ms  
  
Found MTU 1500 for 10.96.104.223 via br0
```

### Related information

[Configuring network links \(SG5700\)](#)

[Changing the MTU setting](#)

### Verifying port-level network connections

To ensure that access between the StorageGRID Appliance Installer and other nodes is

not obstructed by firewalls, confirm that the StorageGRID Appliance Installer can connect to a specific TCP port or set of ports at the specified IP address or range of addresses.

### About this task

Using the list of ports provided in the StorageGRID Appliance Installer, you can test the connectivity between the appliance and the other nodes in your Grid Network.

Additionally, you can test connectivity on the Admin and Client Networks and on UDP ports, such as those used for external NFS or DNS servers. For a list of these ports, see the port reference in the StorageGRID networking guidelines.



The Grid Network ports listed in the port connectivity table are valid only for StorageGRID version 11.5.0. To verify which ports are correct for each node type, you should always consult the networking guidelines for your version of StorageGRID.

### Steps

1. From the StorageGRID Appliance Installer, click **Configure Networking > Port Connectivity Test (nmap)**.

The Port Connectivity Test page appears.

The port connectivity table lists node types that require TCP connectivity on the Grid Network. For each node type, the table lists the Grid Network ports that should be accessible to your appliance.

The following node types require TCP connectivity on the Grid Network.

Node Type	Grid Network Ports
Admin Node	22,80,443,1504,1505,1506,1508,7443,9999
Storage Node without ADC	22,1139,1502,1506,1511,7001,9042,9999,18002,18017,18019,18082,18083,18200
Storage Node with ADC	22,1139,1501,1502,1506,1511,7001,9042,9999,18000,18001,18002,18003,18017,18019,18082,18083,18200,19000
API Gateway	22,1506,1507,9999
Archive Node	22,1506,1509,9999,11139

You can test the connectivity between the appliance ports listed in the table and the other nodes in your Grid Network.

2. From the **Network** drop-down, select the network you want to test: **Grid**, **Admin**, or **Client**.
3. Specify a range of IPv4 addresses for the hosts on that network.

For example, you might want to probe the gateway on the network or the primary Admin Node.

Specify a range using a hyphen, as shown in the example.

4. Enter a TCP port number, a list of ports separated by commas, or a range of ports.

The following node types require TCP connectivity on the Grid Network.

Node Type	Grid Network Ports
Admin Node	22,80,443,1504,1505,1506,1508,7443,9999
Storage Node without ADC	22,1139,1502,1506,1511,7001,9042,9999,18002,18017,18019,18082,18083,18200
Storage Node with ADC	22,1139,1501,1502,1506,1511,7001,9042,9999,18000,18001,18002,18003,18017,18019,18082,18083,18200,19000
API Gateway	22,1506,1507,9999
Archive Node	22,1506,1509,9999,11139

### Port Connectivity Test

Network

IPv4 Address Ranges

Port Ranges

Protocol  TCP  UDP

### 5. Click **Test Connectivity**.

- If the selected port-level network connections are valid, the “Port connectivity test passed” message appears in a green banner. The nmap command output is listed below the banner.

Port connectivity test passed

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Fri Nov 13 18:32:03 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,2022 10.224.6.160-161
Nmap scan report for 10.224.6.160
Host is up (0.00072s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp  open  down

Nmap scan report for 10.224.6.161
Host is up (0.00060s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp  open  down

# Nmap done at Fri Nov 13 18:32:04 2020 -- 2 IP addresses (2 hosts up) scanned in 0.55 seconds
```

- If a port-level network connection is made to the remote host, but the host is not listening on one or more of the selected ports, the “Port connectivity test failed” message appears in a yellow banner. The nmap command output is listed below the banner.

Any remote port the host is not listening to has a state of “closed.” For example, you might see this yellow banner when the node you are trying to connect to is in a pre-installed state and the StorageGRID NMS service is not yet running on that node.

 Port connectivity test failed  
Connection not established. Services might not be listening on target ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:07:02 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,80,443,1504,1505,1506,1508,7443,9999
Nmap scan report for 172.16.4.71
Host is up (0.00020s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
80/tcp    open  http
443/tcp    open  https
1504/tcp   closed evb-elm
1505/tcp   open  funkproxy
1506/tcp   open  utcd
1508/tcp   open  diagmond
7443/tcp   open  oracleas-https
9999/tcp   open  abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:07:03 2020 -- 1 IP address (1 host up) scanned in 0.59 seconds
```

- If a port-level network connection cannot be made for one or more selected ports, the “Port connectivity test failed” message appears in a red banner. The nmap command output is listed below the banner.

The red banner indicates that a TCP connection attempt to a port on the remote host was made, but nothing was returned to the sender. When no response is returned, the port has a state of “filtered” and is likely blocked by a firewall.



Ports with “closed” are also listed.

 Port connectivity test failed  
Connection failed to one or more ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:11:01 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,79,80,443,1504,1505,1506,1508,7443,9999 172.16.4.71
Nmap scan report for 172.16.4.71
Host is up (0.00029s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
79/tcp    filtered finger
80/tcp    open  http
443/tcp    open  https
1504/tcp   closed evb-elm
1505/tcp   open  funkproxy
1506/tcp   open  utcd
1508/tcp   open  diagmond
7443/tcp   open  oracleas-https
9999/tcp   open  abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:11:02 2020 -- 1 IP address (1 host up) scanned in 1.60 seconds
```

## Related information

[Network guidelines](#)

## Accessing and Configuring SANtricity System Manager

You can use SANtricity System Manager to monitor the status of the storage controllers, storage disks, and other hardware components in the storage controller shelf. You can

also configure a proxy for E-Series AutoSupport that enables you to send AutoSupport messages from the appliance without the use of the management port.

### Setting up and Accessing SANtricity System Manager

You might need to access SANtricity System Manager on the storage controller to monitor the hardware in the storage controller shelf or to configure E-Series AutoSupport.

#### What you'll need

- You are using a supported web browser.
- To access SANtricity System Manager through Grid Manager, you must have installed StorageGRID, and you must have the Storage Appliance Administrator permission or Root Access permission.
- To access SANtricity System Manager using the StorageGRID Appliance Installer, you must have the SANtricity System Manager administrator username and password.
- To access SANtricity System Manager directly using a web browser, you must have the SANtricity System Manager administrator username and password.



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager or the StorageGRID Appliance Installer. You can check your firmware version by using the StorageGRID Appliance Installer and selecting **Help > About**.



Accessing SANtricity System Manager from the Grid Manager or from the Appliance Installer is generally meant only for monitoring your hardware and configuring E-Series AutoSupport. Many features and operations within SANtricity System Manager such as upgrading firmware do not apply to monitoring your StorageGRID appliance. To avoid issues, always follow the hardware installation and maintenance instructions for your appliance.

#### About this task

There are three ways to access SANtricity System Manager, depending upon what stage of the installation and configuration process you are in:

- If the appliance has not yet been deployed as a node in your StorageGRID system, you should use the Advanced tab in the StorageGRID Appliance Installer.



Once the node is deployed, you can no longer use the StorageGRID Appliance Installer to access SANtricity System Manager.

- If the appliance has been deployed as a node in your StorageGRID system, use the SANtricity System Manager tab on the Nodes page in Grid Manager.
- If you cannot use the StorageGRID Appliance Installer or Grid Manager, you can access SANtricity System Manager directly using a web browser connected to the management port.

This procedure includes steps for your initial access to SANtricity System Manager. If you have already set up SANtricity System Manager, go to the [Configure hardware alerts](#) step.



Using either the Grid Manager or the StorageGRID Appliance Installer enables you to access SANtricity System Manager without having to configure or connect the management port of the appliance.

You use SANtricity System Manager to monitor the following:

- Performance data such as storage array level performance, I/O latency, CPU utilization, and throughput
- Hardware component status
- Support functions including viewing diagnostic data

You can use SANtricity System Manager to configure the following settings:

- Email alerts, SNMP alerts, or syslog alerts for the components in the storage controller shelf
- E-Series AutoSupport settings for the components in the storage controller shelf.

For additional details on E-Series AutoSupport, see the E-Series documentation center.

[NetApp E-Series Systems Documentation Site](#)

- Drive Security keys, which are needed to unlock secured drives (this step is required if the Drive Security feature is enabled)
- Administrator password for accessing SANtricity System Manager

## Steps

1. Do one of the following:

- Use the StorageGRID Appliance Installer and select **Advanced > SANtricity System Manager**
- Use the Grid Manager and select **Nodes > appliance Storage Node > SANtricity System Manager**



If these options are not available or the login page does not appear, you must use the IP address of the storage controller. Access SANtricity System Manager by browsing to the storage controller IP:

**`https://Storage_Controller_IP`**

The login page for SANtricity System Manager appears.

2. Set or enter the administrator password.



SANtricity System Manager uses a single administrator password that is shared among all users.

The Set Up wizard appears.

## 1 Welcome

## 2 Verify Hardware

## 3 Verify Hosts

## 4 Select Applications

## 5 Define Workloads

## 6 Acc

Welcome to the SANtricity® System Manager! With System Manager, you can...

- Configure your storage array and set up alerts.
- Monitor and troubleshoot any problems when they occur.
- Keep track of how your system is performing in real time.

Cancel


Next &gt;

3. Select **Cancel** to close the wizard.



Do not complete the Set Up wizard for a StorageGRID appliance.

The SANtricity System Manager home page appears.



4. Configure hardware alerts.

- a. Select **Help** to access the online help for SANtricity System Manager.
  - b. Use the **Settings > Alerts** section of the online help to learn about alerts.
  - c. Follow the “How To” instructions to set up email alerts, SNMP alerts, or syslog alerts.
5. Manage AutoSupport for the components in the storage controller shelf.
- a. Select **Help** to access the online help for SANtricity System Manager.
  - b. Use the **Support > Support Center** section of the online help to learn about the AutoSupport feature.
  - c. Follow the “How To” instructions to manage AutoSupport.

For specific instructions on configuring a StorageGrid proxy for sending E-Series AutoSupport messages without using the management port, go to the instructions for administering StorageGRID and search for "proxy settings for E-Series AutoSupport."

### Administer StorageGRID

6. If the Drive Security feature is enabled for the appliance, create and manage the security key.
- a. Select **Help** to access the online help for SANtricity System Manager.
  - b. Use the **Settings > System > Security key management** section of the online help to learn about Drive Security.
  - c. Follow the “How To” instructions to create and manage the security key.
7. Optionally, change the administrator password.
- a. Select **Help** to access the online help for SANtricity System Manager.
  - b. Use the **Home > Storage array administration** section of the online help to learn about the administrator password.
  - c. Follow the "How To" instructions to change the password.

### Reviewing the hardware status in SANtricity System Manager

You can use SANtricity System Manager to monitor and manage the individual hardware components in the storage controller shelf and to review hardware diagnostic and environmental information, such as component temperatures, as well as issues related to the drives.

#### What you'll need

- You are using a supported web browser.
- To access SANtricity System Manager through Grid Manager, you must have the Storage Appliance Administrator permission or Root Access permission.
- To access SANtricity System Manager using the StorageGRID Appliance Installer, you must have the SANtricity System Manager administrator username and password.
- To access SANtricity System Manager directly using a web browser, you must have the SANtricity System Manager administrator username and password.



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager or the StorageGRID Appliance Installer.



Accessing SANtricity System Manager from the Grid Manager or from the Appliance Installer is generally meant only for monitoring your hardware and configuring E-Series AutoSupport. Many features and operations within SANtricity System Manager such as upgrading firmware do not apply to monitoring your StorageGRID appliance. To avoid issues, always follow the hardware installation and maintenance instructions for your appliance.


## Steps

1. Access SANtricity System Manager.

### Setting up and Accessing SANtricity System Manager

2. Enter the administrator username and password if required.
3. Click **Cancel** to close the Set Up wizard and to display the SANtricity System Manager home page.

The SANtricity System Manager home page appears. In SANtricity System Manager, the controller shelf is referred to as a storage array.



4. Review the information displayed for appliance hardware and confirm that all hardware components have a status of Optimal.
  - a. Click the **Hardware** tab.
  - b. Click **Show back of shelf**.

## HARDWARE

[Learn More >](#)

Legend ▾

 Show status icon details ?

Controller Shelf 99 ▾

[Show front of shelf](#)

From the back of the shelf, you can view both storage controllers, the battery in each storage controller, the two power canisters, the two fan canisters, and expansion shelves (if any). You can also view component temperatures.

- c. To see the settings for each storage controller, select the controller, and select **View settings** from the context menu.
- d. To see the settings for other components in the back of the shelf, select the component you want to view.
- e. Click **Show front of shelf**, and select the component you want to view.

From the front of the shelf, you can view the drives and the drive drawers for the storage controller shelf or the expansion shelves (if any).

If the status of any component is Needs Attention, follow the steps in the Recovery Guru to resolve the issue or contact technical support.

### Setting the IP addresses for the storage controllers using the StorageGRID Appliance Installer

Management port 1 on each storage controller connects the appliance to the management network for SANtricity System Manager. If you cannot access to the SANtricity System Manager from the StorageGRID Appliance Installer, you must set a static IP address for each storage controller to ensure that you do not lose your management connection to the hardware and the controller firmware in the controller shelf.

#### What you'll need

- You are using any management client that can connect to the StorageGRID Admin Network, or you have a service laptop.
- The client or service laptop has a supported web browser.

## About this task

DHCP-assigned addresses can change at any time. Assign static IP addresses to the controllers to ensure consistent accessibility.



Follow this procedure only if you do not have access to SANtricity System Manager from the StorageGRID Appliance Installer (**Advanced > SANtricity System Manager**) or Grid Manager (**Nodes > SANtricity System Manager**).

## Steps

1. From the client, enter the URL for the StorageGRID Appliance Installer:  
**`https://Appliance_Controller_IP:8443`**

For *Appliance\_Controller\_IP*, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > Storage Controller Network Configuration**.

The Storage Controller Network Configuration page appears.

3. Depending on your network configuration, select **Enabled** for IPv4, IPv6, or both.
4. Make a note of the IPv4 address that is automatically displayed.

DHCP is the default method for assigning an IP address to the storage controller management port.



It might take a few minutes for the DHCP values to appear.

IPv4 Address Assignment	<input type="radio"/> Static	<input checked="" type="radio"/> DHCP
IPv4 Address (CIDR)	<input type="text" value="10.224.5.166/21"/>	
Default Gateway	<input type="text" value="10.224.0.1"/>	

5. Optionally, set a static IP address for the storage controller management port.



You should either assign a static IP for the management port or assign a permanent lease for the address on the DHCP server.

- a. Select **Static**.
- b. Enter the IPv4 address, using CIDR notation.
- c. Enter the default gateway.

IPv4 Address Assignment     Static     DHCP

IPv4 Address (CIDR)	10.224.2.200/21
Default Gateway	10.224.0.1

d. Click **Save**.

It might take a few minutes for your changes to be applied.

When you connect to SANtricity System Manager, you will use the new static IP address as the URL:  
**https://Storage\_Controller\_IP**

### Optional: Enabling node encryption

If you enable node encryption, the disks in your appliance can be protected by secure key management server (KMS) encryption against physical loss or removal from the site. You must select and enable node encryption during appliance installation and cannot unselect node encryption once the KMS encryption process starts.

#### What you'll need

Review the information about KMS in the instructions for administering StorageGRID.

#### About this task

An appliance that has node encryption enabled connects to the external key management server (KMS) that is configured for the StorageGRID site. Each KMS (or KMS cluster) manages the encryption keys for all appliance nodes at the site. These keys encrypt and decrypt the data on each disk in an appliance that has node encryption enabled.

A KMS can be set up in Grid Manager before or after the appliance is installed in StorageGRID. See the information about KMS and appliance configuration in the instructions for administering StorageGRID for additional details.

- If a KMS is set up before installing the appliance, KMS-controlled encryption begins when you enable node encryption on the appliance and add it to a StorageGRID site where KMS is configured.
- If a KMS is not set up before you install the appliance, KMS-controlled encryption is performed on each appliance that has node encryption enabled as soon as a KMS is configured and available for the site that contains the appliance node.



Any data that exists before an appliance that has node encryption enabled connects to the configured KMS is encrypted with a temporary key that is not secure. The appliance is not protected from removal or theft until the key is set to a value provided by the KMS.

Without the KMS key needed to decrypt the disk, data on the appliance cannot be retrieved and is effectively lost. This is the case whenever the decryption key cannot be retrieved from the KMS. The key becomes inaccessible if a customer clears the KMS configuration, a KMS key expires, connection to the KMS is lost, or the appliance is removed from the StorageGRID system where its KMS keys are installed.

#### Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

**https://Controller\_IP:8443**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.



After the appliance has been encrypted with a KMS key, the appliance disks cannot be decrypted without using the same KMS key.

2. Select **Configure Hardware > Node Encryption**.

NetApp® StorageGRID® Appliance Installer Help ▾

Home Configure Networking ▾ Configure Hardware ▾ Monitor Installation Advanced ▾

### Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

#### Encryption Status

⚠ You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

**Save**

Key Management Server Details

3. Select **Enable node encryption**.

You can unselect **Enable node encryption** without risk of data loss until you select **Save** and the appliance node accesses the KMS encryption keys in your StorageGRID system and begins disk encryption. You are not able to disable node encryption after the appliance is installed.



After you add an appliance that has node encryption enabled to a StorageGRID site that has a KMS, you cannot stop using KMS encryption for the node.

4. Select **Save**.
5. Deploy the appliance as a node in your StorageGRID system.

KMS-controlled encryption begins when the appliance accesses the KMS keys configured for your StorageGRID site. The installer displays progress messages during the KMS encryption process, which might take a few minutes depending on the number of disk volumes in the appliance.



Appliances are initially configured with a random non-KMS encryption key assigned to each disk volume. The disks are encrypted using this temporary encryption key, that is not secure, until the appliance that has node encryption enabled accesses the KMS keys configured for your StorageGRID site.

### After you finish

You can view node-encryption status, KMS details, and the certificates in use when the appliance node is in

maintenance mode.

## Related information

[Administer StorageGRID](#)

[Monitoring node encryption in maintenance mode](#)

### Optional: Changing the RAID mode (SG5760 only)

If you have an SG5760 with 60 drives, you can change to a different RAID mode to accommodate your storage and recovery requirements. You can only change the mode before deploying the StorageGRID appliance Storage Node.

#### What you'll need

- You have an SG5760. If you have an SG5712, you must use DDP mode.
- You are using any client that can connect to StorageGRID.
- The client has a supported web browser.

#### About this task

Before deploying the SG5760 appliance as a Storage Node, you can choose one of the following volume configuration options:

- **DDP**: This mode uses two parity drives for every eight data drives. This is the default and recommended mode for all appliances. When compared to RAID6, DDP delivers better system performance, reduced rebuild times after drive failures, and ease of management. DDP also provides drawer loss protection in 60-drive appliances.
- **DDP16**: This mode uses two parity drives for every 16 data drives, which results in higher storage efficiency compared to DDP. When compared to RAID6, DDP16 delivers better system performance, reduced rebuild times after drive failures, ease of management, and comparable storage efficiency. To use DDP16 mode, your configuration must contain at least 20 drives. DDP16 does not provide drawer loss protection.
- **RAID6**: This mode uses two parity drives for every 16 or more data drives. To use RAID 6 mode, your configuration must contain at least 20 drives. Although RAID6 can increase storage efficiency of the appliance when compared to DDP, it is not recommended for most StorageGRID environments.



If any volumes have already been configured or if StorageGRID was previously installed, changing the RAID mode causes the volumes to be removed and replaced. Any data on those volumes will be lost.

#### Steps

1. Using the service laptop, open a web browser and access the StorageGRID Appliance Installer:

**`https://E5700SG_Controller_IP:8443`**

Where `E5700SG_Controller_IP` is any of the IP addresses for the E5700SG controller.

2. Select **Advanced > RAID Mode**.
3. On the **Configure RAID Mode** page, select the desired RAID mode from the Mode drop-down list.
4. Click **Save**.

## Related information

### Optional: Remapping network ports for the appliance

You might need to remap the internal ports on the appliance Storage Node to different external ports. For example, you might need to remap ports because of a firewall issue.

#### What you'll need

- You have previously accessed the StorageGRID Appliance Installer.
- You have not configured and do not plan to configure load balancer endpoints.



If you remap any ports, you cannot use the same ports to configure load balancer endpoints. If you want to configure load balancer endpoints and have already remapped ports, follow the steps in the recovery and maintenance instructions for removing port remaps.

#### Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Remap Ports**.

The Remap Port page appears.

2. From the **Network** drop-down box, select the network for the port you want to remap: Grid, Admin, or Client.
3. From the **Protocol** drop-down box, select the IP protocol: TCP or UDP.
4. From the **Remap Direction** drop-down box, select which traffic direction you want to remap for this port: Inbound, Outbound, or Bi-directional.
5. For **Original Port**, enter the number of the port you want to remap.
6. For **Mapped-To Port**, enter the number of the port you want to use instead.
7. Click **Add Rule**.

The new port mapping is added to the table, and the remapping takes effect immediately.

#### Remap Ports

If required, you can remap the internal ports on the appliance Storage Node to different external ports. For example, you might need to remap ports because of a firewall issue.

The screenshot shows the 'Remap Ports' configuration interface. At the top, there are buttons for 'Remove Selected Rule' and 'Add Rule'. Below these are several fields: 'Network' (set to 'Grid'), 'Protocol' (set to 'TCP'), 'Remap Direction' (set to 'Inbound'), 'Original Port' (set to '1'), and 'Mapped-To Port' (set to '1'). Below the form is a table with the following data:

	Network	Protocol	Remap Direction	Original Port	Mapped-To Port
<input type="radio"/>	Grid	TCP	Bi-directional	1800	1801

8. To remove a port mapping, select the radio button for the rule you want to remove, and click **Remove Selected Rule**.

## Deploying an appliance Storage Node

After installing and configuring the storage appliance, you can deploy it as a Storage Node in a StorageGRID system. When you deploy an appliance as a Storage Node, you use the StorageGRID Appliance Installer included on the appliance.

### What you'll need

- If you are cloning an appliance node, continue following the process in recovery and maintenance.

#### Maintain & recover

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.
- 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.
- The primary Admin Node for the StorageGRID system has been deployed.
- 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 a service laptop with a supported web browser.

### About this task

Each storage appliance functions as a single Storage Node. Any appliance can connect to the Grid Network, the Admin Network, and the Client Network

To deploy an appliance Storage Node in a StorageGRID system, you access the StorageGRID Appliance Installer and perform the following steps:

- You specify or confirm the IP address of the primary Admin Node and the name of the Storage Node.
- You start the deployment and wait as volumes are configured and the software is installed.
- When the installation pauses partway through the appliance installation tasks, you resume the installation by signing into the Grid Manager, approving all grid nodes, and completing the StorageGRID installation and deployment processes.



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.

- If you are performing an expansion or recovery operation, follow the appropriate instructions:
  - To add an appliance Storage Node to an existing StorageGRID system, see the instructions for expanding a StorageGRID system.
  - To deploy an appliance Storage Node as part of a recovery operation, see instructions for recovery and maintenance.

### 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.

## 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

Connection state

Connection to 172.16.4.210 ready



### Node name

Node name




### Installation

Current state

Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

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	<ol style="list-style-type: none"> <li>Unselect the <b>Enable Admin Node discovery</b> check box.</li> <li>Enter the IP address manually.</li> <li>Click <b>Save</b>.</li> <li>Wait for the connection state for the new IP address to become ready.</li> </ol>
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"> <li>Select the <b>Enable Admin Node discovery</b> check box.</li> <li>Wait for the list of discovered IP addresses to be displayed.</li> <li>Select the primary Admin Node for the grid where this appliance Storage Node will be deployed.</li> <li>Click <b>Save</b>.</li> <li>Wait for the connection state for the new IP address to become ready.</li> </ol>

- In the **Node name** field, enter the name you want to use for this appliance node, and click **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.

- 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 installation and maintenance instructions for your appliance.



If you are deploying the Storage Node appliance as a node cloning target, stop the deployment process here and continue the node cloning procedure in [Maintain & recover](#).

- 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**.

- If your grid includes multiple appliance Storage Nodes, repeat these 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. This script applies only to Storage Nodes.

## Related information

[Expand your grid](#)

[Maintain & recover](#)

## Monitoring the storage 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**.

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 are not re-run. If you are re-running an installation, any tasks that do not need to be re-run are shown with a green status bar and a status of "Skipped."

2. Review the progress of the first two installation stages.

#### 1. Configure storage

During this stage, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity software 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. Go to the next step.

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...

```

- Go to the Grid Manager, approve the pending storage node, and complete the StorageGRID installation process.

When you click **Install** from the Grid Manager, Stage 3 completes and stage 4, **Finalize Installation**, begins. When stage 4 completes, the controller is rebooted.

## Automating appliance installation and configuration

You can automate the installation and configuration of your appliances and configuration of the whole StorageGRID system.

### About this task

Automating installation and configuration can be useful for deploying multiple StorageGRID instances or one large, complex StorageGRID instance.

To automate installation and configuration, use one or more of the following options:

- Create a JSON file that specifies the configuration settings for your appliances. Upload the JSON file using the StorageGRID Appliance Installer.



You can use the same file to configure more than one appliance.

- Use the `StorageGRIDconfigure-sga.py` Python script to automate the configuration of your appliances.
- Use additional Python scripts to configure other components of the whole StorageGRID system (the "grid").



You can use StorageGRID automation Python scripts directly, or you can use them as examples of how to use the StorageGRID Installation REST API in grid deployment and configuration tools you develop yourself. See the information about downloading and extracting the StorageGRID installation files in the Recovery and Maintenance instructions.

## Automating appliance configuration using the StorageGRID Appliance Installer

You can automate the configuration of an appliance by using a JSON file that contains the configuration information. You upload the file using the StorageGRID Appliance Installer.

### What you'll need

- Your appliance must be on the latest firmware compatible with StorageGRID 11.5 or higher.
- You must be connected to the StorageGRID Appliance Installer on the appliance you are configuring using a supported browser.

### About this task

You can automate appliance configuration tasks such as configuring the following:

- Grid Network, Admin Network, and Client Network IP addresses
- BMC interface
- Network links
  - Port bond mode
  - Network bond mode
  - Link speed

Configuring your appliance using an uploaded JSON file is often more efficient than performing the configuration manually using multiple pages in the StorageGRID Appliance Installer, especially if you have to configure many nodes. You must apply the configuration file for each node one at a time.



Experienced users who want to automate both the installation and configuration of their appliances can use the `configure-sga.py` script.

[Automating installation and configuration of appliance nodes using the `configure-sga.py` script](#)

## Steps

1. Generate the JSON file using one of the following methods:

- The ConfigBuilder application

[ConfigBuilder.netapp.com](http://ConfigBuilder.netapp.com)

- The `configure-sga.py` appliance configuration script. You can download the script from StorageGRID Appliance Installer (**Help > Appliance Configuration Script**). See the instructions on automating the configuration using the `configure-sga.py` script.

[Automating installation and configuration of appliance nodes using the configure-sga.py script](#)

The node names in the JSON file must follow these requirements:

- Must be a valid hostname containing at least 1 and no more than 32 characters
- Can use letters, numbers, and hyphens are allowed
- Cannot start or end with a hyphen or contain only numbers



Ensure that the node names (the top-level names) in the JSON file are unique, or you will not be able to configure more than one node using the JSON file.

2. Select **Advanced > Update Appliance Configuration**.

The Update Appliance Configuration page appears.

### Update Appliance Configuration

Use a JSON file to update this appliance's configuration. You can generate the JSON file from the [ConfigBuilder](#) application or from the [appliance configuration script](#).

**⚠** You might lose your connection if the applied configuration from the JSON file includes "link\_config" and/or "networks" sections. If you are not reconnected within 1 minute, re-enter the URL using one of the other IP addresses assigned to the appliance.

#### Upload JSON

JSON configuration	<input type="button" value="Browse"/>
Node name	<input type="button" value="-- Upload a file"/>
<input type="button" value="Apply JSON configuration"/>	

3. Select the JSON file with the configuration you want to upload.

- a. Select **Browse**.
- b. Locate and select the file.
- c. Select **Open**.

The file is uploaded and validated. When the validation process is complete, the file name is shown

next to a green check mark.



You might lose connection to the appliance if the configuration from the JSON file includes sections for "link\_config", "networks", or both. If you are not reconnected within 1 minute, re-enter the appliance URL using one of the other IP addresses assigned to the appliance.

#### Upload JSON

JSON configuration  ✓ appliances.orig.json

Node name -- Select a node ▼

The **Node name** drop down is populated with the top-level node names defined in the JSON file.



If the file is not valid, the file name is shown in red and an error message is displayed in a yellow banner. The invalid file is not applied to the appliance. You can use ConfigBuilder to ensure you have a valid JSON file.

4. Select a node from the list in the **Node name** drop down.

The **Apply JSON configuration** button is enabled.

#### Upload JSON

JSON configuration  ✓ appliances.orig.json

Node name Lab-80-1000 ▼

5. Select **Apply JSON configuration**.

The configuration is applied to the selected node.

### Automating installation and configuration of appliance nodes using the `configure-sga.py` script

You can use the `configure-sga.py` script to automate many of the installation and configuration tasks for StorageGRID appliance nodes, including installing and configuring a primary Admin Node. This script can be useful if you have a large number of appliances to configure. You can also use the script to generate a JSON file that contains appliance configuration information.

## About this task

- 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 primary Admin Node using the StorageGRID Appliance Installer.
- If you are installing the primary Admin Node, you know its IP address.
- If you are installing and configuring other nodes, the primary Admin Node has been deployed, and you know its IP address.
- For all nodes other than the primary Admin Node, 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 downloaded the `configure-sga.py` file. The file is included in the installation archive, or you can access it by clicking **Help > Appliance Installation Script** in the StorageGRID Appliance Installer.



This procedure is for advanced users with some experience using command-line interfaces. Alternatively, you can also use the StorageGRID Appliance Installer to automate the configuration.

[Automating appliance configuration using the StorageGRID Appliance Installer](#)

## Steps

1. Log in to the Linux machine you are using to run the Python script.
2. For general help with the script syntax and to see a list of the available parameters, enter the following:

```
configure-sga.py --help
```

The `configure-sga.py` script uses five subcommands:

- `advanced` for advanced StorageGRID appliance interactions, including BMC configuration and creating a JSON file containing the current configuration of the appliance
- `configure` for configuring the RAID mode, node name, and networking parameters
- `install` for starting a StorageGRID installation
- `monitor` for monitoring a StorageGRID installation
- `reboot` for rebooting the appliance

If you enter a subcommand (`advanced`, `configure`, `install`, `monitor`, or `reboot`) argument followed by the `--help` option you will get a different help text providing more detail on the options available within that subcommand:

```
configure-sga.py subcommand --help
```

3. To confirm the current configuration of the appliance node, enter the following where `SGA-install-ip` is any one of the IP addresses for the appliance node:

```
configure-sga.py configure SGA-INSTALL-IP
```

The results show current IP information for the appliance, including the IP address of the primary Admin Node and information about the Admin, Grid, and Client Networks.

```
Connecting to +https://10.224.2.30:8443+ (Checking version and
```

connectivity.)

2021/02/25 16:25:11: Performing GET on /api/versions... Received 200

2021/02/25 16:25:11: Performing GET on /api/v2/system-info... Received 200

2021/02/25 16:25:11: Performing GET on /api/v2/admin-connection... Received 200

2021/02/25 16:25:11: Performing GET on /api/v2/link-config... Received 200

2021/02/25 16:25:11: Performing GET on /api/v2/networks... Received 200

2021/02/25 16:25:11: Performing GET on /api/v2/system-config... Received 200

#### StorageGRID Appliance

Name: LAB-SGA-2-30

Node type: storage

#### StorageGRID primary Admin Node

IP: 172.16.1.170

State: unknown

Message: Initializing...

Version: Unknown

#### Network Link Configuration

##### Link Status

Link	State	Speed (Gbps)
----	-----	-----
1	Up	10
2	Up	10
3	Up	10
4	Up	10
5	Up	1
6	Down	N/A

##### Link Settings

Port bond mode: FIXED

Link speed: 10GBE

Grid Network: ENABLED

Bonding mode: active-backup

VLAN: novlan

MAC Addresses: 00:a0:98:59:8e:8a 00:a0:98:59:8e:82

Admin Network: ENABLED

Bonding mode: no-bond

MAC Addresses: 00:80:e5:29:70:f4

```

Client Network:      ENABLED
    Bonding mode:    active-backup
    VLAN:            novlan
    MAC Addresses:   00:a0:98:59:8e:89  00:a0:98:59:8e:81

Grid Network
CIDR:      172.16.2.30/21 (Static)
MAC:       00:A0:98:59:8E:8A
Gateway:   172.16.0.1
Subnets:  172.17.0.0/21
           172.18.0.0/21
           192.168.0.0/21
MTU:       1500

Admin Network
CIDR:      10.224.2.30/21 (Static)
MAC:       00:80:E5:29:70:F4
Gateway:   10.224.0.1
Subnets:  10.0.0.0/8
           172.19.0.0/16
           172.21.0.0/16
MTU:       1500

Client Network
CIDR:      47.47.2.30/21 (Static)
MAC:       00:A0:98:59:8E:89
Gateway:   47.47.0.1
MTU:       2000

#####
##### If you are satisfied with this configuration, #####
##### execute the script with the "install" sub-command. #####
#####

```

4. If you need to change any of the values in the current configuration, use the `configure` subcommand to update them. For example, if you want to change the IP address that the appliance uses for connection to the primary Admin Node to `172.16.2.99`, enter the following:

```
configure-sga.py configure --admin-ip 172.16.2.99 SGA-INSTALL-IP
```

5. If you want to back up the appliance configuration to a JSON file, use the `advanced` and `backup-file` subcommands. For example, if you want to back up the configuration of an appliance with IP address `SGA-INSTALL-IP` to a file named `appliance-SG1000.json`, enter the following:

```
configure-sga.py advanced --backup-file appliance-SG1000.json SGA-INSTALL-IP
```

The JSON file containing the configuration information is written to the same directory you executed the script from.



Check that the top-level node name in the generated JSON file matches the appliance name. Do not make any changes to this file unless you are an experienced user and have a thorough understanding of StorageGRID APIs.

- When you are satisfied with the appliance configuration, use the `install` and `monitor` subcommands to install the appliance:

```
configure-sga.py install --monitor SGA-INSTALL-IP
```

- If you want to reboot the appliance, enter the following:

```
configure-sga.py reboot SGA-INSTALL-IP
```

## Automating the configuration of StorageGRID

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

### What you'll need

- 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>	Sample 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 sample 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

- 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`.

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

For example:

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

## After you finish

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, you need to extract 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.
```

## Overview of installation REST APIs

StorageGRID provides two REST APIs for performing installation tasks: the StorageGRID Installation API and the StorageGRID Appliance Installer API.

Both APIs use 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 (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs 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 in the event that 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.

### StorageGRID Appliance Installer API

The StorageGRID Appliance Installer API can be accessed over HTTPS from `Controller_IP:8443`.

To access the API documentation, go to the StorageGRID Appliance Installer on the appliance and select **Help > API Docs** from the menu bar.

The StorageGRID Appliance Installer API includes the following sections:

- **clone** — Operations to configure and control node cloning.
- **encryption** — Operations to manage encryption and view encryption status.
- **hardware configuration** — Operations to configure system settings on attached hardware.
- **installation** — Operations for starting the appliance installation and for monitoring installation status.
- **networking** — Operations related to the Grid, Admin, and Client Network configuration for a StorageGRID appliance and appliance port settings.
- **setup** — Operations to help with initial appliance installation setup including requests to get information about the system and update the primary Admin Node IP.
- **support** — Operations for rebooting the controller and getting logs.
- **upgrade** — Operations related to upgrading appliance firmware.
- **uploadsg** — Operations for uploading StorageGRID installation files.

## Troubleshooting the hardware installation

If you encounter issues during the installation, you might find it helpful to review troubleshooting information related to hardware setup and connectivity issues.

### Related information

[Hardware setup appears to hang](#)

## Hardware setup appears to hang

The StorageGRID Appliance Installer might not be available if hardware faults or cabling errors prevent the E5700SG controller from completing its boot-up processing.

### Steps

1. Watch the codes on the seven-segment displays.

While the hardware is initializing during power up, the two seven-segment displays show a sequence of codes. When the hardware boots successfully, the seven-segment displays show different codes for each controller.

2. Review the codes on the seven-segment display for the E5700SG controller.



The installation and provisioning take time. Some installation phases do not report updates to the StorageGRID Appliance Installer for several minutes.

If an error occurs, the seven-segment display flashes a sequence, such as HE.

3. To understand what these codes mean, see the following resources:

Controller	Reference
E5700SG controller	<ul style="list-style-type: none"><li>• “Status indicators on the E5700SG controller”</li><li>• “HE error: Error synchronizing with SANtricity OS Software”</li></ul>
E2800 controller	<p><i>E5700 and E2800 System Monitoring Guide</i></p> <p><b>Note:</b> The codes described for the E-Series E5700 controller do not apply to the E5700SG controller in the appliance.</p>

4. If this does not resolve the issue, contact technical support.

### Related information

[Status indicators on the E5700SG controller](#)

[HE error: Error synchronizing with SANtricity OS Software](#)

[NetApp E-Series Systems Documentation Site](#)

### HE error: Error synchronizing with SANtricity OS Software

The seven-segment display on the compute controller shows an HE error code if the StorageGRID Appliance Installer cannot synchronize with SANtricity OS Software.

### About this task

If an HE error code is displayed, perform this corrective action.

## Steps

1. Check the two interconnect cables between the two controllers, and confirm that the cables and SFP+ transceivers are securely connected.
2. As required, replace one or both of the cables or SFP+ transceivers, and try again.
3. If this does not resolve the issue, contact technical support.

## Troubleshooting connection issues

If you encounter connection issues during the StorageGRID appliance installation, you should perform the corrective action steps listed.

### Unable to connect to the appliance

If you cannot connect to the appliance, there might be a network issue, or the hardware installation might not have been completed successfully.

## Steps

1. If you are unable to connect to SANtricity System Manager:
  - a. Try to ping the appliance using the IP address for the E2800 controller on the management network for SANtricity System Manager:  
**ping E2800\_Controller\_IP**
  - b. If you receive no response from the ping, confirm you are using the correct IP address.  
  
Use the IP address for management port 1 on the E2800 controller.
  - c. If the IP address is correct, check appliance cabling and the network setup.  
  
If that does not resolve the issue, contact technical support.
  - d. If the ping was successful, open a web browser.
  - e. Enter the URL for SANtricity System Manager:  
**https://E2800\_Controller\_IP**  
  
The log in page for SANtricity System Manager appears.
2. If you are unable to connect to the E5700SG controller:
  - a. Try to ping the appliance using the IP address for the E5700SG controller:  
**ping E5700SG\_Controller\_IP**
  - b. If you receive no response from the ping, confirm you are using the correct IP address.  
  
You can use the IP address of the appliance on the Grid Network, the Admin Network, or the Client Network.
  - c. If the IP address is correct, check appliance cabling, SFP transceivers, and the network setup.  
  
If that does not resolve the issue, contact technical support.
  - d. If the ping was successful, open a web browser.
  - e. Enter the URL for the StorageGRID Appliance Installer:

`https://E5700SG_Controller_IP:8443`

The Home page appears.

## Rebooting the controller while the StorageGRID Appliance Installer is running


You might need to reboot the compute controller while the StorageGRID Appliance Installer is running. For example, you might need to reboot the controller if the installation fails.

### About this task

This procedure only applies when the compute controller is running the StorageGRID Appliance Installer. Once the installation is completed, this step no longer works because the StorageGRID Appliance Installer is no longer available.

### Steps

1. From the StorageGRID Appliance Installer, click **Advanced** > **Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
  - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



The SG6000-CN controller is rebooted.

## Maintaining the SG5700 appliance

You might need to upgrade the SANtricity OS Software on the E2800 controller, change the Ethernet link configuration of the E5700SG controller, replace the E2800 controller or the E5700SG controller, or replace specific components. The procedures in this section assume that the appliance has already been deployed as a Storage Node in a StorageGRID system.

### Steps

- [Placing an appliance into maintenance mode](#)
- [Upgrading SANtricity OS on the storage controller](#)
- [Upgrading drive firmware using SANtricity System Manager](#)
- [Replacing the E2800 controller](#)
- [Replacing the E5700SG controller](#)
- [Replacing other hardware components](#)
- [Changing the link configuration of the E5700SG controller](#)
- [Changing the MTU setting](#)
- [Checking the DNS server configuration](#)
- [Monitoring node encryption in maintenance mode](#)

## Placing an appliance into maintenance mode

You must place the appliance into maintenance mode before performing specific maintenance procedures.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission. For details, see the instructions for administering StorageGRID.

### About this task

Placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.



The password and host key for a StorageGRID appliance in maintenance mode remain the same as they were when the appliance was in service.

### Steps

1. From the Grid Manager, select **Nodes**.
2. From the tree view of the Nodes page, select the appliance Storage Node.
3. Select **Tasks**.

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

Places the appliance's compute controller into maintenance mode.

Maintenance Mode

#### 4. Select **Maintenance Mode**.

A confirmation dialog box appears.

### Enter Maintenance Mode on SGA-106-15

You must place the appliance's compute controller into maintenance mode to perform certain maintenance procedures on the appliance.

Attention: All StorageGRID services on this node will be shut down. Wait a few minutes for the node to reboot into maintenance mode.

If you are ready to start, enter the provisioning passphrase and click OK.

Provisioning Passphrase

Cancel

OK

#### 5. Enter the provisioning passphrase, and select **OK**.

A progress bar and a series of messages, including "Request Sent," "Stopping StorageGRID," and "Rebooting," indicate that the appliance is completing the steps for entering maintenance mode.

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

**Attention:** Your request has been sent, but the appliance might take 10-15 minutes to enter maintenance mode. Do not perform maintenance procedures until this tab indicates maintenance mode is ready, or data could become corrupted.



Request Sent

When the appliance is in maintenance mode, a confirmation message lists the URLs you can use to access the StorageGRID Appliance Installer.

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

This node is currently in maintenance mode. Navigate to one of the URLs listed below and perform any necessary maintenance procedures.

- <https://172.16.2.106:8443>
- <https://10.224.2.106:8443>
- <https://47.47.2.106:8443>
- <https://169.254.0.1:8443>

When you are done with any required maintenance procedures, you must exit maintenance mode by clicking Reboot Controller from the StorageGRID Appliance Installer.

6. To access the StorageGRID Appliance Installer, browse to any of the URLs displayed.

If possible, use the URL containing the IP address of the appliance's Admin Network port.




Accessing <https://169.254.0.1:8443> requires a direct connection to the local management port.


7. From the StorageGRID Appliance Installer, confirm that the appliance is in maintenance mode.

⚠ This node is in maintenance mode. Perform any required maintenance procedures. If you want to exit maintenance mode manually to resume normal operation, go to **Advanced > Reboot Controller** to **reboot** the controller.

8. Perform any required maintenance tasks.
9. After completing maintenance tasks, exit maintenance mode and resume normal node operation. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Upgrading SANtricity OS on the storage controller

To ensure optimal functioning of the storage controller, you must upgrade to the latest maintenance release of the SANtricity OS that is qualified for your StorageGRID appliance. Consult the NetApp Interoperability Matrix Tool (IMT) to determine which version you should be using. If you need assistance, contact technical support.

- If the storage controller is using SANtricity OS 08.42.20.00 (11.42) or newer, use the Grid Manager to perform the upgrade.

## [Upgrading SANtricity OS on the storage controllers using the Grid Manager](#)

- If the storage controller is using a SANtricity OS version older than 08.42.20.00 (11.42), use maintenance mode to perform the upgrade.

## [Upgrading SANtricity OS on the E2800 controller using maintenance mode](#)

### **Related information**

[NetApp Interoperability Matrix Tool](#)

[NetApp Downloads: SANtricity OS](#)

[Monitor & troubleshoot](#)

### **Upgrading SANtricity OS on the storage controllers using the Grid Manager**

For storage controllers currently using SANtricity OS 08.42.20.00 (11.42) or newer, you must use the Grid Manager to apply an upgrade.

### **What you'll need**

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- You must have the Maintenance permission.
- You must be signed in to the Grid Manager using a supported browser.
- You must have the provisioning passphrase.
- You must have access to the NetApp downloads page for SANtricity OS.

### **About this task**

You cannot perform other software updates (StorageGRID software upgrade or a hotfix) until you have completed the SANtricity OS upgrade process. If you attempt to start a hotfix or a StorageGRID software upgrade before the SANtricity OS upgrade process has finished, you are redirected to the SANtricity OS upgrade page.

The procedure will not be complete until the SANtricity OS upgrade has been successfully applied to all applicable nodes. It might take more than 30 minutes to load the SANtricity OS on each node and up to 90 minutes to reboot each StorageGRID storage appliance.



The following steps are only applicable when you are using the Grid Manager to perform the upgrade. The storage controllers in the SG5700 series appliance cannot be upgraded using the Grid Manager when the controllers are using SANtricity OS older than 08.42.20.00 (11.42).



This procedure will automatically upgrade the NVSRAM to the most recent version associated with the SANtricity OS upgrade. You do not need to apply a separate NVSRAM upgrade file.

### **Steps**

1. From a service laptop, download the new SANtricity OS Software file from the NetApp support site.

Be sure to choose the SANtricity OS version for the E2800 storage controllers.

[NetApp Downloads: SANtricity OS](#)

2. Sign in to the Grid Manager using a supported browser.
3. Select **Maintenance**. Then, in the System section of the menu, select **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.

- To perform a major version upgrade of StorageGRID, see the [instructions for upgrading StorageGRID](#), and then select **StorageGRID Upgrade**.
- To apply a hotfix to all nodes in your system, see "Hotfix procedure" in the [recovery and maintenance instructions](#), and then select **StorageGRID Hotfix**.
- To upgrade SANtricity OS software on a storage controller, see "Upgrading SANtricity OS Software on the storage controllers" in the installation and maintenance instructions for your storage appliance, and then select **SANtricity OS**:

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)



4. Click **SANtricity OS**.

The SANtricity OS page appears.

#### SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

#### SANtricity OS Upgrade File

SANtricity OS Upgrade File



Browse

#### Passphrase

Provisioning Passphrase



Start

5. Select the SANtricity OS upgrade file you downloaded from the NetApp support site.
  - a. Click **Browse**.
  - b. Locate and select the file.
  - c. Click **Open**.

The file is uploaded and validated. When the validation process is done, the file name is shown in the Details field.



Do not change the file name since it is part of the verification process.

### SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

### SANtricity OS Upgrade File

SANtricity OS Upgrade File



✓ RC\_XXXXXXXXXX\_XXXXXX.dlp

Details

RC\_XXXXXXXXXX\_XXXXXX.dlp

### Passphrase

Provisioning Passphrase



6. Enter the provisioning passphrase.

The **Start** button is enabled.

## SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

### SANtricity OS Upgrade File

SANtricity OS Upgrade File

Browse

✓ RC\_20240311\_143\_145\_146\_1701.dlp

Details

RC\_20240311\_143\_145\_146\_1701.dlp

### Passphrase

Provisioning Passphrase

\*\*\*\*\*

Start

#### 7. Click **Start**.

A warning box appears stating that your browser's connection might be lost temporarily as services on nodes that are upgraded are restarted.

### Warning

**Nodes can disconnect and services might be affected**

The node will be automatically rebooted at the end of upgrade and services will be affected. Are you sure you want to start the SANtricity OS upgrade?

Cancel

OK

#### 8. Click **OK** to stage the SANtricity OS upgrade file to the primary Admin Node.

When the SANtricity OS upgrade starts:

- a. The health check is run. This process checks that no nodes have the status of Needs Attention.



If any errors are reported, resolve them and click **Start** again.

- b. The SANtricity OS Upgrade Progress table appears. This table shows all Storage Nodes in your grid and the current stage of the upgrade for each node.



The table shows all Storage Nodes, including software-based Storage Nodes. You must approve the upgrade for all Storage Nodes, even though a SANtricity OS upgrade has no effect on software-based Storage Nodes. The upgrade message returned for software-based Storage Nodes is “SANtricity OS upgrade is not applicable to this node.”

### SANtricity OS Upgrade Progress

Approve All Remove All

▲ Storage Nodes - 0 out of 4 completed
Approve All Remove All

Site	Name	Progress	Stage	Details	Action
RTP Lab 1	DT-10-224-1-181-S1		Waiting for you to approve		<span>Approve</span>
RTP Lab 1	DT-10-224-1-182-S2		Waiting for you to approve		<span>Approve</span>
RTP Lab 1	DT-10-224-1-183-S3		Waiting for you to approve		<span>Approve</span>
RTP Lab 1	NetApp-SGA-Lab2-002-024		Waiting for you to approve		<span>Approve</span>

◀ ▶

- Optionally, sort the list of nodes 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.

You can scroll through the list of nodes by using the left and right arrows at the bottom right corner of the section.

- Approve the grid nodes you are ready to add to the upgrade queue. Approved nodes of the same type are upgraded one at a time.



Do not approve the SANtricity OS upgrade for an appliance storage node unless you are sure the node is ready to be stopped and rebooted. When the SANtricity OS upgrade is approved on a node, the services on that node are stopped. Later, when the node is upgraded, the appliance node is rebooted. These operations might cause service interruptions for clients that are communicating with the node.

- Click either of the **Approve All** buttons to add all Storage Nodes to the SANtricity OS upgrade queue.



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).

- Click one or more **Approve** buttons to add one or more nodes to the SANtricity OS upgrade queue.



You can delay applying a SANtricity OS upgrade to a node, but the SANtricity OS upgrade process will not be complete until you approve the SANtricity OS upgrade on all the listed Storage Nodes.

After you click **Approve**, the upgrade process determines if the node can be upgraded. If a node can

be upgraded, it is added to the upgrade queue. +

For some nodes, the selected upgrade file is intentionally not applied and you can complete the upgrade process without upgrading these specific nodes. For nodes intentionally not upgraded, the process will show stage of Complete with one of the following messages in the Details column:

- Storage Node was already upgraded.
- SANtricity OS upgrade is not applicable to this node.
- SANtricity OS file is not compatible with this node.

The message “SANtricity OS upgrade is not applicable to this node” indicates that the node does not have a storage controller that can be managed by the StorageGRID system. This message will appear for non-appliance Storage Nodes. You can complete the SANtricity OS upgrade process without upgrading the node displaying this message.

The message “SANtricity OS file is not compatible with this node” indicates that the node requires a SANtricity OS file different than the one the process is attempting to install. After you complete the current SANtricity OS upgrade, download the SANtricity OS appropriate for the node and repeat the upgrade process.

11. If you need to remove a node or all nodes from the SANtricity OS upgrade queue, click **Remove** or **Remove All**.

As shown in the example, when the stage progresses beyond Queued, the **Remove** button is hidden and you can no longer remove the node from the SANtricity OS upgrade process.

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

12. Wait while the SANtricity OS upgrade is applied to each approved grid node.



If any node shows a stage of Error while the SANtricity OS upgrade is being applied, the upgrade has failed for that node. The appliance might need to be placed in maintenance mode to recover from the failure. Contact technical support before continuing.

If the firmware on the node is too old to be upgraded with the Grid Manager, the node shows a stage of Error with the details: “You must use maintenance mode to upgrade SANtricity OS on this node. See the installation and maintenance instructions for your appliance. After the upgrade, you can use this utility for

future upgrades.” To resolve the error, do the following:

- a. Use maintenance mode to upgrade SANtricity OS on the node that shows a stage of Error.
- b. Use the Grid Manager to re-start and complete the SANtricity OS upgrade.

When the SANtricity OS upgrade is complete on all approved nodes, the SANtricity OS Upgrade Progress table closes and a green banner shows the date and time the SANtricity OS upgrade was completed.

SANtricity OS upgrade completed at 2020-04-07 13:26:02 EDT.

SANtricity OS Upgrade File

SANtricity OS Upgrade File

Passphrase

Provisioning Passphrase

13. Repeat this upgrade procedure for any nodes with a stage of Complete that require a different SANtricity OS upgrade file.



For any nodes with a status of Needs Attention, use maintenance mode to perform the upgrade.

## Related information

[Upgrading SANtricity OS on the E2800 controller using maintenance mode](#)

### Upgrading SANtricity OS on the E2800 controller using maintenance mode

For storage controllers currently using SANtricity OS older than 08.42.20.00 (11.42), you must use the maintenance mode procedure to apply an upgrade.

### What you'll need

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- You must place the E5700SG controller into maintenance mode, which interrupts the connection to the E2800 controller. Putting a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

[Placing an appliance into maintenance mode](#)

### About this task

Do not upgrade the SANtricity OS or NVSRAM in the E-Series controller on more than one StorageGRID appliance at a time.



Upgrading more than one StorageGRID appliance at a time might cause data unavailability, depending on your deployment model and ILM policies.

## Steps

1. From a service laptop, access SANtricity System Manager and sign in.
2. Download the new SANtricity OS Software file and NVSRAM file to the management client.



The NVSRAM is specific to the StorageGRID appliance. Do not use the standard NVSRAM download.

3. Follow the instructions in the *E2800 and E5700 SANtricity Software and Firmware Upgrade Guide* or the SANtricity System Manager online help to upgrade the E2800 controller's firmware and NVSRAM.




Activate the upgrade files immediately. Do not defer activation.

4. Once the upgrade operation has completed, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Upgrading SANtricity OS on the storage controllers using the Grid Manager](#)

## Upgrading drive firmware using SANtricity System Manager

You upgrade your drive firmware to make sure you have all the latest features and bug fixes.

### What you'll need

- The storage appliance has an Optimal status.
- All drives have an Optimal status.
- You have the latest version of SANtricity System Manager installed that is compatible with your StorageGRID version.
- You have placed the StorageGRID appliance in maintenance mode.

### Placing an appliance into maintenance mode



Maintenance mode interrupts the connection to the storage controller, stopping all I/O activity and placing all drives offline.



Do not upgrade the drive firmware on more than one StorageGRID appliance at a time. Doing so might cause data unavailability, depending on your deployment model and ILM policies.

### Steps

1. Access SANtricity System Manager using one of these methods:
  - Use the StorageGRID Appliance Installer and select **Advanced > SANtricity System Manager**
  - Use the Grid Manager and select **Nodes > appliance Storage Node > SANtricity System Manager**



If these options are not available or the SANtricity System Manager login page does not appear, access SANtricity System Manager by browsing to the storage controller IP:  
**`https://Storage_Controller_IP`**

2. Enter the SANtricity System Manager administrator username and password, if required.
3. Verify the drive firmware version currently installed in the storage appliance:
  - a. From SANtricity System Manager, select **Support > Upgrade Center**.
  - b. Under Drive Firmware upgrade, select **Begin Upgrade**.

The Upgrade Drive Firmware displays the drive firmware files currently installed.

- c. Note the current drive firmware revisions and drive identifiers in the Current Drive Firmware column.

## Upgrade Drive Firmware

1 Select Upgrade Files
2 Select Drives

Review your current drive firmware and select upgrade files below...

[What do I need to know before upgrading drive firmware?](#)

Current Drive Firmware	Associated Drives
MS02, KPM51VUG800G	<a href="#">View drives</a>

Total rows: 1 | [↻](#)

Select up to four drive firmware files: Browse...

In this example:

- The drive firmware revision is **MS02**.
- The drive identifier is **KPM51VUG800G**.

Select **View drives** in the Associated Drives column to display where these drives are installed in your storage appliance.

- d. Close the Upgrade Drive Firmware window.
4. Download and prepare the available drive firmware upgrade:
  - a. Under Drive Firmware upgrade, select **NetApp Support**.
  - b. On the NetApp Support web site, select the **Downloads** tab, and then select **E-Series Disk Drive Firmware**.

The E-Series Disk Firmware page displays.

- c. Search for each **Drive Identifier** installed in your storage appliance and verify that each drive identifier has the latest firmware revision.
  - If the firmware revision is not a link, this drive identifier has the latest firmware revision.
  - If one or more drive part numbers are listed for a drive identifier, a firmware upgrade is available for these drives. You can select any link to download the firmware file.

PRODUCTS ▾ SYSTEMS ▾ DOCS & KNOWLEDGEBASE ▾ COMMUNITY ▾ DOWNLOADS ▾ TOOLS ▾ CASES ▾ PARTS ▾

Downloads > Firmware > E-Series Disk Firmware

## E-Series Disk Firmware

Download all current E-Series Disk Firmware

Drive Part Number ▾	Descriptions ▾	Drive Identifier ▾	Firmware Rev. (Download)	Notes and Config Info	Release Date ▾
Drive Part Number	Descriptions	KPM51VUG800G	Firmware Rev. (Download)		
E-X4041C	SSD, 800GB, SAS, PI	KPM51VUG800G	MS03	MS02 Fixes <a href="#">Bug 1194908</a> MS03 Fixes <a href="#">Bug 1334862</a>	04-Sep-2020

d. If a later firmware revision is listed, select the link in the Firmware Rev. (Download) column to download a .zip archive containing the firmware file.

e. Extract (unzip) the drive firmware archive files you downloaded from the Support site.

5. Install the drive firmware upgrade:

a. From SANtricity System Manager, under Drive Firmware upgrade, select **Begin Upgrade**.

b. Select **Browse**, and select the new drive firmware files that you downloaded from the Support site.

Drive firmware files have a filename similar to +

D\_HUC101212CSS600\_30602291\_MS01\_2800\_0002.dlp

You can select up to four drive firmware files, one at a time. If more than one drive firmware file is compatible with the same drive, you get a file conflict error. Decide which drive firmware file you want to use for the upgrade and remove the other one.

c. Select **Next**.

**Select Drives** lists the drives that you can upgrade with the selected firmware files.

Only drives that are compatible appear.

The selected firmware for the drive appears in **Proposed Firmware**. If you must change this firmware, select **Back**.

d. Select **Offline (parallel)** upgrade.

You can use the offline upgrade method because the appliance is in maintenance mode, where I/O activity is stopped for all drives and all volumes.

e. In the first column of the table, select the drive or drives you want to upgrade.

The best practice is to upgrade all drives of the same model to the same firmware revision.

f. Select **Start**, and confirm that you want to perform the upgrade.

If you need to stop the upgrade, select **Stop**. Any firmware downloads currently in progress complete. Any firmware downloads that have not started are canceled.



Stopping the drive firmware upgrade might result in data loss or unavailable drives.

g. (Optional) To see a list of what was upgraded, select **Save Log**.

The log file is saved in the downloads folder for your browser with the name `latest-upgrade-log-timestamp.txt`.

If any of the following errors occur during the upgrade procedure, take the appropriate recommended action.

- **Failed assigned drives**

One reason for the failure might be that the drive does not have the appropriate signature. Make sure that the affected drive is an authorized drive. Contact technical support for more information.

When replacing a drive, make sure that the replacement drive has a capacity equal to or greater than the failed drive you are replacing.

You can replace the failed drive while the storage array is receiving I/O.

- **Check storage array**

- Make sure that an IP address has been assigned to each controller.
- Make sure that all cables connected to the controller are not damaged.
- Make sure that all cables are tightly connected.

- **Integrated hot spare drives**

This error condition must be corrected before you can upgrade the firmware.

- **Incomplete volume groups**

If one or more volume groups or disk pools are incomplete, you must correct this error condition before you can upgrade the firmware.

- **Exclusive operations (other than background media/parity scan) currently running on any volume groups**

If one or more exclusive operations are in progress, the operations must complete before the firmware can be upgraded. Use System Manager to monitor the progress of the operations.

- **Missing volumes**

You must correct the missing volume condition before the firmware can be upgraded.

- **Either controller in a state other than Optimal**

One of the storage array controllers needs attention. This condition must be corrected before the firmware can be upgraded.

- **Mismatched Storage Partition information between Controller Object Graphs**

An error occurred while validating the data on the controllers. Contact technical support to resolve this issue.

- **SPM Verify Database Controller check fails**

A storage partitions mapping database error occurred on a controller. Contact technical support to resolve this issue.

- **Configuration Database Validation (If supported by the storage array's controller version)**

A configuration database error occurred on a controller. Contact technical support to resolve this issue.

- **MEL Related Checks**

Contact technical support to resolve this issue.

- **More than 10 DDE Informational or Critical MEL events were reported in the last 7 days**

Contact technical support to resolve this issue.

- **More than 2 Page 2C Critical MEL Events were reported in the last 7 days**

Contact technical support to resolve this issue.

- **More than 2 Degraded Drive Channel Critical MEL events were reported in the last 7 days**


Contact technical support to resolve this issue.

- **More than 4 critical MEL entries in the last 7 days**


Contact technical support to resolve this issue.

6. Once the upgrade operation has completed, reboot the appliance. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Upgrading SANtricity OS on the storage controller](#)

## Replacing the E2800 controller

You might need to replace the E2800 controller if it is not functioning optimally or if it has failed.

### About this task

- You have a replacement controller with the same part number as the controller you are replacing.
- You have downloaded the instructions for replacing the simplex configuration of a failed E2800 controller canister.



Refer to the E-Series instructions only when directed or if you need more details to perform a specific step. Do not rely on the E-Series instructions to replace a controller in the StorageGRID appliance, because the procedures are not the same.

- You have labels to identify each cable that is connected to the controller.
- If all drives are secured, you have reviewed the steps in the simplex E2800 controller replacement procedure, which include downloading and installing E-Series SANtricity Storage Manager from the NetApp Support Site and then using the Enterprise Management Window (EMW) to unlock the secured drives after you have replaced the controller.



You will not be able to use the appliance until you unlock the drives with the saved key.

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

### About this task

You can determine if you have a failed controller canister in two ways:

- The Recovery Guru in SANtricity System Manager directs you to replace the controller.
- The amber Attention LED on the controller is on, indicating that the controller has a fault.

The appliance Storage Node will not be accessible when you replace the controller. If the E2800 controller is functioning sufficiently, you can place the E5700SG controller into maintenance mode.

## Placing an appliance into maintenance mode

When you replace a controller, you must remove the battery from the original controller and install it in the replacement controller.



The E2800 controller in the appliance does not include a host interface card (HIC).

### Steps

1. Follow the instructions in the E2800 controller replacement procedure to prepare to remove the controller.

You use SANtricity System Manager to perform these steps.

- a. Make a note of which version of SANtricity OS software is currently installed on the controller.
- b. Make a note of which version of NVSRAM is currently installed.
- c. If the Drive Security feature is enabled, be sure a saved key exists and that you know the pass phrase required to install it.



**Possible loss of data access** -- If all drives in the appliance are security enabled, the new controller will not be able to access the appliance until you unlock the secured drives using the Enterprise Management Window in SANtricity Storage Manager.

- d. Back up the configuration database.

If a problem occurs when you remove a controller, you can use the saved file to restore your configuration.

- e. Collect support data for the appliance.



Collecting support data before and after replacing a component ensures you can send a full set of logs to technical support in case the replacement does not resolve the problem.

2. If the StorageGRID appliance is running in a StorageGRID system, place the E5700SG controller into maintenance mode.

## Placing an appliance into maintenance mode

3. If the E2800 controller is functioning sufficiently to allow for a controlled shutdown, confirm that all operations have completed.
  - a. From the home page of SANtricity System Manager, select **View Operations in Progress**.
  - b. Confirm that all operations have completed.
4. Remove the controller from the appliance:
  - a. Put on an ESD wristband or take other antistatic precautions.
  - b. Label the cables and then disconnect the cables and SFPs.



To prevent degraded performance, do not twist, fold, pinch, or step on the cables.

- c. Release the controller from the appliance by squeezing the latch on the cam handle until it releases, and then open the cam handle to the right.

d. Using two hands and the cam handle, slide the controller out of the appliance.



Always use two hands to support the weight of the controller.

e. Place the controller on a flat, static-free surface with the removable cover facing up.

f. Remove the cover by pressing down on the button and sliding the cover off.

5. Remove the battery from the failed controller, and install it into the replacement controller:

a. Confirm that the green LED inside the controller (between the battery and the DIMMs) is off.


If this green LED is on, the controller is still using battery power. You must wait for this LED to go off before removing any components.



Item	Description
	Internal Cache Active LED
	Battery

b. Locate the blue release latch for the battery.

c. Unlatch the battery by pushing the release latch down and away from the controller.



Item	Description
	Battery release latch
	Battery

- d. Lift up on the battery, and slide it out of the controller.
- e. Remove the cover from the replacement controller.
- f. Orient the replacement controller so that the slot for the battery faces toward you.
- g. Insert the battery into the controller at a slight downward angle.

You must insert the metal flange at the front of the battery into the slot on the bottom of the controller, and slide the top of the battery beneath the small alignment pin on the left side of the controller.

- h. Move the battery latch up to secure the battery.

When the latch clicks into place, the bottom of the latch hooks into a metal slot on the chassis.

- i. Turn the controller over to confirm that the battery is installed correctly.



**Possible hardware damage**— The metal flange at the front of the battery must be completely inserted into the slot on the controller (as shown in the first figure). If the battery is not installed correctly (as shown in the second figure), the metal flange might contact the controller board, causing damage.

- **Correct**— The battery's metal flange is completely inserted in the slot on the controller:



- **Incorrect**— The battery's metal flange is not inserted into the slot on the controller:




- j. Replace the controller cover.
6. Install the replacement controller into the appliance.
    - a. Turn the controller over, so that the removable cover faces down.
    - b. With the cam handle in the open position, slide the controller all the way into the appliance.
    - c. Move the cam handle to the left to lock the controller in place.
    - d. Replace the cables and SFPs.
    - e. Wait for the E2800 controller to reboot. Verify that the seven-segment display shows a state of 99.
    - f. Determine how you will assign an IP address to the replacement controller.




The steps for assigning an IP address to the replacement controller depend on whether you connected management port 1 to a network with a DHCP server and on whether all drives are secured.

- If management port 1 is connected to a network with a DHCP server, the new controller will obtain its IP address from the DHCP server. This value might be different than the original controller's IP address.
  - If all drives are secured, you must use the Enterprise Management Window (EMW) in SANtricity Storage Manager to unlock the secured drives. You cannot access the new controller until you unlock the drives with the saved key. See the E-Series instructions for replacing a simplex E2800 controller.
7. If the appliance uses secured drives, follow the instructions in the E2800 controller replacement procedure to import the drive security key.
  8. Return the appliance to normal operating mode. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.




During the reboot, the following screen appears:



The appliance reboots and rejoins the grid. This process can take up to 20 minutes.

9. Confirm that the reboot is complete and that the node has rejoined the grid. In the Grid Manager, verify that

the **Nodes** tab displays a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



10. From SANtricity System Manager, confirm that the new controller is Optimal, and collect support data.

### Related information

[NetApp E-Series Systems Documentation Site](#)

### Replacing the E5700SG controller

You might need to replace the E5700SG controller if it is not functioning optimally or if it has failed.

#### What you'll need

- You have a replacement controller with the same part number as the controller you are replacing.
- You have downloaded the E-Series instructions for replacing a failed E5700 controller.



Use the E-Series instructions for reference only if you need more details to perform a specific step. Do not rely on the E-Series instructions to replace a controller in the StorageGRID appliance, because the procedures are not the same. For example, the E-Series instructions for the E5700 controller describe how to remove the battery and the host interface card (HIC) from a failed controller and install them in a replacement controller. These steps do not apply to the E5700SG controller.

- You have labels to identify each cable that is connected to the controller.
- The appliance has been placed maintenance mode.

[Placing an appliance into maintenance mode](#)

#### About this task

The appliance Storage Node will not be accessible when you replace the controller. If the E5700SG controller is functioning sufficiently, you can perform a controlled shutdown at the start of this procedure.



If you are replacing the controller before installing StorageGRID software, you might not be able to access the StorageGRID Appliance Installer immediately after completing this procedure. While you can access the StorageGRID Appliance Installer from other hosts on the same subnet as the appliance, you cannot access it from hosts on other subnets. This condition should resolve itself within 15 minutes (when any ARP cache entries for the original controller time out), or you can clear the condition immediately by purging any old ARP cache entries manually from the local router or gateway.

## Steps

1. When the appliance has been placed in maintenance mode, shut down the E5700SG controller.

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. Shut down the E5700SG controller:

**shutdown -h now**

c. Wait for any data in cache memory to be written to the drives.

The green Cache Active LED on the back of the E2800 controller is on when cached data needs to be written to the drives. You must wait for this LED to turn off.

2. Turn off the power.

- a. From the home page of SANtricity System Manager, select **View Operations in Progress**.
- b. Confirm that all operations have completed.
- c. Turn off both power switches on the appliance.
- d. Wait for all LEDs to turn off.

3. If the StorageGRID networks attached to the controller use DHCP servers:

- a. Note the MAC addresses for the ports on the replacement controller (located on labels on the controller).
- b. Ask your network administrator to update the IP address settings for the original controller to reflect the MAC addresses for the replacement controller.



You must ensure that the IP addresses for the original controller have been updated before you apply power to the replacement controller. Otherwise, the controller will obtain new DHCP IP addresses when it boots up and might not be able to reconnect to StorageGRID. This step applies to all StorageGRID networks that are attached to the controller.

4. Remove the controller from the appliance:

- a. Put on an ESD wristband or take other antistatic precautions.
- b. Label the cables and then disconnect the cables and SFPs.



To prevent degraded performance, do not twist, fold, pinch, or step on the cables.

- c. Release the controller from the appliance by squeezing the latch on the cam handle until it releases, and then open the cam handle to the right.
- d. Using two hands and the cam handle, slide the controller out of the appliance.



Always use two hands to support the weight of the controller.

5. Install the replacement controller into the appliance.
  - a. Turn the controller over, so that the removable cover faces down.
  - b. With the cam handle in the open position, slide the controller all the way into the appliance.
  - c. Move the cam handle to the left to lock the controller in place.
  - d. Replace the cables and SFPs.
6. Power on the appliance, and monitor the controller LEDs and seven-segment displays.

After the controllers have successfully booted up, the seven-segment displays should show the following:

- E2800 controller:

The final state is 99.

- E5700SG controller:

The final state is HA.

7. Confirm that the appliance Storage Node appears in the Grid Manager and that no alarms appear.

### Related information

[NetApp E-Series Systems Documentation Site](#)

### Replacing other hardware components

You might need to replace a controller battery, drive, fan, or power supply, in the StorageGRID appliance.

#### What you'll need

- You have the E-Series hardware replacement procedure.
- The appliance has been placed in maintenance mode if the component replacement procedure requires that you shut down the appliance.

[Placing an appliance into maintenance mode](#)

#### About this task

To replace the battery in the E2800 controller, see the instructions in these instructions for replacing the E2800 controller. Those instructions describe how to remove the controller from the appliance, remove the battery from the controller, install the battery, and replace the controller.

To replace a drive, power-fan canister, fan canister, power canister, or drive drawer in the appliance, access the E-Series procedures for maintaining E2800 hardware.

## SG5712 component replacement instructions

FRU	See E-Series instructions for
Drive	Replacing a drive in E2800 12-drive or 24-drive shelves
Power-fan canister	Replacing a power-fan canister in E2800 shelves

## SG5760 component replacement instructions

FRU	See E-Series instructions for
Drive	Replacing a drive in E2860 shelves
Power canister	Replacing a power canister in E2860 shelves
Fan canister	Replacing a fan canister in E2860 shelves
Drive drawer	Replacing a drive drawer in E2860 shelves

### Related information

[Replacing the E2800 controller](#)

[NetApp E-Series Systems Documentation Site](#)

## Changing the link configuration of the E5700SG controller

You can change the Ethernet link configuration of the E5700SG controller. You can change the port bond mode, the network bond mode, and the link speed.

### What you'll need

You must place the E5700SG controller into maintenance mode. Putting a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

[Placing an appliance into maintenance mode](#)


### About this task

Options for changing the Ethernet link configuration of the E5700SG controller include:

- Changing **Port bond mode** from Fixed to Aggregate, or from Aggregate to Fixed
- Changing **Network bond mode** from Active-Backup to LACP, or from LACP to Active-Backup
- Enabling or disabling VLAN tagging, or changing the value of a VLAN tag
- Changing the link speed from 10-GbE to 25-GbE, or from 25-GbE to 10-GbE

### Steps

1. Select **Configure Networking > Link Configuration** from the menu.



2. Make the desired changes to the link configuration.

For more information on the options, see “Configuring network links.”

3. When you are satisfied with your selections, click **Save**.



You might lose your connection if you made changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:

**`https://E5700SG_Controller_IP:8443`**

If you made changes to the VLAN settings, the subnet for the appliance might have changed. If you need to change the IP addresses for the appliance, follow the instructions for configuring IP addresses.

### Setting the IP configuration


4. From the StorageGRID Appliance Installer, select **Configure Networking > Ping Test**.

5. Use the Ping Test tool to check connectivity to IP addresses on any networks that might have been affected by the link configuration changes you made in the [Change link configuration](#) step.


In addition to any other tests you choose to perform, confirm that you can ping the grid IP address of the primary Admin Node, and the grid IP address of at least one other Storage Node. If necessary, correct any link configuration issues.

6. Once you are satisfied that your link configuration changes are working, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Configuring network links \(SG5700\)](#)

## Changing the MTU setting

You can change the MTU setting that you assigned when you configured IP addresses for the appliance node.

### What you'll need

The appliance has been placed maintenance mode.

[Placing an appliance into maintenance mode](#)

### Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > IP Configuration**.
2. Make the desired changes to the MTU settings for the Grid Network, Admin Network, and Client Network.

## Grid Network

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

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

 All required Grid Network subnets must also be defined in the Grid Network Subnet List on the Primary Admin Node before starting installation.

Subnets (CIDR)  



MTU  



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.



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


- When you are satisfied with the settings, select **Save**.
- Reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this

option if you are done working in maintenance mode and are ready to return the node to normal operation.

- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Administer StorageGRID](#)

## Checking the DNS server configuration

You can check and temporarily change the domain name system (DNS) servers that are currently in use by this appliance node.

### What you'll need

The appliance has been placed in maintenance mode.

[Placing an appliance into maintenance mode](#)

### About this task

You might need to change the DNS server settings if an encrypted appliance cannot connect to the key management server (KMS) or KMS cluster because the hostname for the KMS was specified as a domain name instead of an IP address. Any changes that you make to the DNS settings for the appliance are temporary and are lost when you exit maintenance mode. To make these changes permanent, specify the DNS servers in Grid Manager (**Maintenance > Network > DNS Servers**).

- Temporary changes to the DNS configuration are necessary only for node-encrypted appliances where the KMS server is defined using a fully qualified domain name, instead of an IP address, for the hostname.
- When a node-encrypted appliance connects to a KMS using a domain name, it must connect to one of the DNS servers defined for the grid. One of these DNS servers then translates the domain name into an IP address.
- If the node cannot reach a DNS server for the grid, or if you changed the grid-wide DNS settings when a node-encrypted appliance node was offline, the node is unable to connect to the KMS. Encrypted data on the appliance cannot be decrypted until the DNS issue is resolved.

To resolve a DNS issue preventing KMS connection, specify the IP address of one or more DNS servers in the StorageGRID Appliance Installer. These temporary DNS settings allow the appliance to connect to the KMS and decrypt data on the node.

For example, if the DNS server for the grid changes while an encrypted node was offline, the node will not be able to reach the KMS when it comes back online, since it is still using the previous DNS values. Entering the new DNS server IP address in the StorageGRID Appliance Installer allows a temporary KMS connection to decrypt the node data.

### Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > DNS Configuration**.
2. Verify that the DNS servers specified are correct.

#### DNS Servers

 Configuration changes made on this page will not be passed to the StorageGRID software after appliance installation.

#### Servers

Server 1	<input type="text" value="10.224.223.135"/>	
Server 2	<input type="text" value="10.224.223.136"/>	 
<input type="button" value="Cancel"/>		<input type="button" value="Save"/>

3. If required, change the DNS servers.




Changes made to the DNS settings are temporary and are lost when you exit maintenance mode.

4. When you are satisfied with the temporary DNS settings, select **Save**.

The node uses the DNS server settings specified on this page to reconnect to the KMS, allowing data on the node to be decrypted.


5. After node data is decrypted, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



When the node reboots and rejoins the grid, it uses the system-wide DNS servers listed in the Grid Manager. After rejoining the grid, the appliance will no longer use the temporary DNS servers specified in the StorageGRID Appliance Installer while the appliance was in maintenance mode.

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



### Monitoring node encryption in maintenance mode

If you enabled node encryption for the appliance during installation, you can monitor the node-encryption status of each appliance node, including the node-encryption state and key management server (KMS) details.

## What you'll need

- Node encryption must have been enabled for the appliance during installation. You cannot enable node encryption after the appliance is installed.
- The appliance has been placed into maintenance mode.

### [Placing an appliance into maintenance mode](#)

## Steps

1. From the StorageGRID Appliance Installer, select **Configure Hardware > Node Encryption**.

### Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

### Encryption Status

 You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

Save

### Key Management Server Details

View the status and configuration details for the KMS that manages the encryption key for this appliance. You must use the Grid Manager to make configuration changes.

KMS display name	thales
External key UID	41b0306abcce451facfe01b1b4870ae1c1ec6bd5e3849d790223766baf35c57
Hostnames	10.96.99.164 10.96.99.165
Port	5696

Server certificate >

Client certificate >

### Clear KMS Key

 Do not clear the KMS key if you need to access or preserve any data on this appliance.

If you want to reinstall this appliance node (for example, in another grid), you must clear the KMS key. When the KMS key is cleared, all data on this appliance is deleted.

Clear KMS Key and Delete Data

The Node Encryption page includes these three sections:

- Encryption Status shows whether node encryption is enabled or disabled for the appliance.
- Key Management Server Details shows information about the KMS being used to encrypt the appliance. You can expand the server and client certificate sections to view certificate details and status.

- To address issues with the certificates themselves, such as renewing expired certificates, see the information about KMS in the instructions for administering StorageGRID.
- If there are unexpected problems connecting to KMS hosts, verify that the domain name system (DNS) servers are correct and that appliance networking is correctly configured.

### Checking the DNS server configuration

- If you are unable to resolve your certificate issues, contact technical support.
- Clear KMS Key disables node encryption for the appliance, removes the association between the appliance and the key management server that was configured for the StorageGRID site, and deletes all data from the appliance. You must clear the KMS key before you can install the appliance into another StorageGRID system.


### Clearing the key management server configuration




Clearing the KMS configuration deletes data from the appliance, rendering it permanently inaccessible. This data is not recoverable.

2. When you are done checking node-encryption status, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Administer StorageGRID](#)

### Clearing the key management server configuration

Clearing the key management server (KMS) configuration disables node encryption on your appliance. After clearing the KMS configuration, the data on your appliance is permanently deleted and is no longer accessible. This data is not recoverable.

### What you'll need

If you need to preserve data on the appliance, you must perform a node decommission procedure before you clear the KMS configuration.



When KMS is cleared, data on the appliance will be permanently deleted and no longer accessible. This data is not recoverable.

Decommission the node to move any data it contains to other nodes in StorageGRID. See the recovery and maintenance instructions for grid node decommissioning.

### About this task

Clearing the appliance KMS configuration disables node encryption, removing the association between the appliance node and the KMS configuration for the StorageGRID site. Data on the appliance is then deleted and the appliance is left in a pre-install state. This process cannot be reversed.

You must clear the KMS configuration:

- Before you can install the appliance into another StorageGRID system, that does not use a KMS or that uses a different KMS.



Do not clear the KMS configuration if you plan to reinstall an appliance node in a StorageGRID system that uses the same KMS key.

- Before you can recover and reinstall a node where the KMS configuration was lost and the KMS key is not recoverable.
- Before returning any appliance that was previously in use at your site.
- After decommissioning a appliance that had node encryption enabled.



Decommission the appliance before clearing KMS to move its data to other nodes in your StorageGRID system. Clearing KMS before decommissioning the appliance will result in data loss and might render the appliance inoperable.

## Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > Node Encryption**.

### Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

### Encryption Status

You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

Save

### Key Management Server Details

View the status and configuration details for the KMS that manages the encryption key for this appliance. You must use the Grid Manager to make configuration changes.

KMS display name	thales
External key UID	41b0306abcce451facfe01b1b4870ae1c1ec6bd5e3849d790223766baf35c57
Hostnames	10.96.99.164 10.96.99.165
Port	5696

Server certificate >

Client certificate >

### Clear KMS Key

Do not clear the KMS key if you need to access or preserve any data on this appliance.


If you want to reinstall this appliance node (for example, in another grid), you must clear the KMS key. When the KMS key is cleared, all data on this appliance is deleted.

Clear KMS Key and Delete Data



If the KMS configuration is cleared, data on the appliance will be permanently deleted. This data is not recoverable.

- At the bottom of the window, select **Clear KMS Key and Delete Data**.
- If you are sure that you want to clear the KMS configuration, type **clear** and select **Clear KMS Key and Delete Data**.



The KMS encryption key and all data are deleted from the node, and the appliance reboots. This can take up to 20 minutes.

- Open a browser, and enter one of the IP addresses for the appliance's compute controller.  
**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

- Select **Configure Hardware > Node Encryption**.
- Verify that node encryption is disabled and that the key and certificate information in **Key Management Server Details** and the **Clear KMS Key and Delete Data** control are removed from the window.

Node encryption cannot be reenabled on the appliance until it is reinstalled in a grid.

### After you finish

After the appliance reboots and you have verified that KMS has been cleared and that the appliance is in a pre-install state, you can physically remove the appliance from your StorageGRID system. See the recovery and maintenance instructions for information about preparing an appliance for reinstallation.

### Related information

[Administer StorageGRID](#)

## SG5600 storage appliances

Learn how to install and maintain StorageGRID SG5612 and SG5660 appliances.

- [StorageGRID appliance overview](#)
- [Installation and deployment overview](#)
- [Preparing for installation](#)
- [Installing the hardware](#)
- [Configuring the hardware](#)
- [Deploying an appliance Storage Node](#)
- [Monitoring the storage appliance installation](#)
- [Automating appliance installation and configuration](#)
- [Overview of installation REST APIs](#)
- [Troubleshooting the hardware installation](#)
- [Maintaining the SG5600 appliance](#)

### StorageGRID appliance overview

The StorageGRID SG5600 appliance is an integrated storage and computing platform that operates as a Storage Node in a StorageGRID grid.

The StorageGRID SG5600 appliance includes the following components:

Component	Description
E5600SG controller	<p>Compute serverThe E5600SG controller runs the Linux operating system and the StorageGRID software.</p> <p>This controller connects to the following:</p> <ul style="list-style-type: none"><li>• The Admin, Grid, and Client Networks for the StorageGRID system</li><li>• The E2700 controller, using dual SAS paths (active/active) with the E5600SG controller operating as the initiator</li></ul>

Component	Description
E2700 controller	<p>Storage controllerThe E2700 controller operates as a standard E-Series storage array in simplex mode, and runs the SANtricity operating system (controller firmware).</p> <p>This controller connects to the following:</p> <ul style="list-style-type: none"> <li>• The management network where SANtricity Storage Manager is installed</li> <li>• The E5600SG controller, using dual SAS paths (active/active) with the E2700 controller operating as the target</li> </ul>

The SG5600 appliance also includes the following components, depending on the model:

Component	Model SG5612	Model SG5660
Drives	12 NL-SAS drives	60 NL-SAS drives
Enclosure	DE1600 enclosure, a two rack-unit (2U) chassis that houses the drives and the controllers	DE6600 enclosure, a four rack-unit (4U) chassis that houses the drives and the controllers
Power supplies and fans	Two power-fan canisters	Two power supplies and two fans



The E5600SG controller is highly customized for use in the StorageGRID appliance. All other components operate as described in E-Series documentation, except as indicated in these instructions.

The maximum raw storage available on each StorageGRID appliance Storage Node is fixed, based on the appliance model and configuration. You cannot expand the available storage by adding a shelf with additional drives.

### StorageGRID appliance features

The StorageGRID SG5600 appliance provides an integrated storage solution for creating a new StorageGRID system or for expanding the capacity of an existing system.

The StorageGRID appliance provides the following features:

- Combines the StorageGRID Storage Node computing and storage elements into a single, efficient, integrated solution
- Simplifies the installation and configuration of a Storage Node, automating most of the process required
- Provides a high-density storage solution with two enclosure options: one that is 2U and one that is 4U
- Uses 10-GbE IP interfaces directly to the Storage Node, without the need for intermediate storage interfaces such as FC or iSCSI

- Can be used in a hybrid grid environment that uses StorageGRID appliances and virtual (software-based) Storage Nodes
- Includes preconfigured storage and comes preloaded with the StorageGRID Appliance Installer (on the E5600SG controller) for field-ready software deployment and integration

### Hardware diagrams

The SG5612 and SG5660 models of the StorageGRID appliance both include an E2700 controller and an E5600SG controller. You should review the diagrams to learn the differences between the models and the controllers.


**Model SG5612 2U: Rear view of the E2700 controller and E5600SG controller**



	Description
1	E2700 controller
2	E5600SG controller


**Model SG5660 4U: Rear view of the E2700 controller and E5600SG controller**

The E2700 controller is above the E5600SG controller.



	Description
1	E2700 controller
2	E5600SG controller

**Rear view of the E2700 controller**




	Description
1	Management port 1 (Connect to the network where SANtricity Storage Manager is installed.)
2	Management port 2 (Use during installation to connect to a laptop.)
3	SAS interconnect port 1
4	SAS interconnect port 2
5	Serial connection port
6	Seven-segment display



The two SAS ports labeled Drive Expansion (green) on the rear of the E2700 controller are not used. The StorageGRID appliance does not support expansion drive shelves.

**Rear view of the E5600SG controller**



	Description
1	Management port 1 (Connect to the Admin network for StorageGRID.)
2	Management port 2 Options: <ul style="list-style-type: none"> <li>• Bond with management port 1 for a redundant connection to the Admin Network for StorageGRID.</li> <li>• Leave unwired and available for temporary local access (IP 169.254.0.1).</li> <li>• During installation, use for IP configuration if DHCP-assigned IP addresses are not available.</li> </ul>
3	SAS interconnect port 1
4	SAS interconnect port 2
5	Fault and Active LEDs for 10-GbE network port 1
6	Fault and Active LEDs for 10-GbE network port 2
7	Fault and Active LEDs for 10-GbE network port 3
8	Fault and Active LEDs for 10-GbE network port 4
9	Needs Attention LED
10	Seven-segment display
11	10-GbE network port 1
12	10-GbE network port 2
13	10-GbE network port 3

	Description
14	10-GbE network port 4



The host interface card (HIC) on the StorageGRID appliance E5600SG controller supports only 10-Gb Ethernet connections. It cannot be used for iSCSI connections.

## Installation and deployment overview

You can install one or more StorageGRID appliances when you first deploy StorageGRID, or you can add appliance Storage Nodes later as part of an expansion. You might also need to install an appliance Storage Node as part of a recovery operation.

Adding a StorageGRID storage appliance to a StorageGRID system includes four primary steps:

1. Preparing for installation:
  - Preparing the installation site
  - Unpacking the boxes and checking the contents
  - Obtaining additional equipment and tools
  - Gathering IP addresses and network information
  - Optional: Configuring an external key management server (KMS) if you plan to encrypt all appliance data. See details about external key management in the instructions for administering StorageGRID.
2. Installing the hardware:
  - Registering the hardware
  - Installing the appliance into a cabinet or rack
  - Installing the drives (SG5660 only)
  - Cabling the appliance
  - Connecting the power cords and applying power
  - Viewing boot-up status codes
3. Configuring the hardware:
  - Accessing SANtricity Storage Manager, setting a static IP address for management port 1 on the E2700 controller, and configuring SANtricity Storage Manager settings
  - Accessing StorageGRID Appliance Installer and configuring the link and network IP settings required to connect to StorageGRID networks
  - Optional: Enabling node encryption if you plan to use an external KMS to encrypt appliance data.
  - Optional: Changing the RAID mode.
4. Deploying the appliance as a Storage Node:

Task	Refer to
Deploying an appliance Storage Node in a new StorageGRID system	<a href="#">Deploying an appliance Storage Node</a>

Task	Refer to
Adding an appliance Storage Node to an existing StorageGRID system	Instructions for expanding a StorageGRID system
Deploying an appliance Storage Node as part of a Storage Node recovery operation	Instructions for recovery and maintenance

### Related information

[Preparing for installation](#)

[Installing the hardware](#)

[Configuring the hardware](#)

[Expand your grid](#)

[Maintain & recover](#)

[Administer StorageGRID](#)

## Preparing for installation

Preparing to install a StorageGRID appliance entails preparing the site and obtaining all required hardware, cables, and tools. You should also gather IP addresses and network information.

### Steps

- [Preparing the site \(SG5600\)](#)
- [Unpacking the boxes \(SG5600\)](#)
- [Obtaining additional equipment and tools \(SG5600\)](#)
- [Service laptop requirements](#)
- [Web browser requirements](#)
- [Reviewing appliance network connections](#)
- [Gathering installation information \(SG5600\)](#)

### Preparing the site (SG5600)

Before installing the appliance, you must make sure that the site and the cabinet or rack you plan to use meet the specifications for a StorageGRID appliance.

### Steps

1. Confirm that the site meets the requirements for temperature, humidity, altitude range, airflow, heat dissipation, wiring, power, and grounding. See the NetApp Hardware Universe for more information.
2. Obtain a 19-inch (48.3-cm) cabinet or rack to fit shelves of this size (without cables):

Appliance model	Height	Width	Depth	Maximum weight
SG5612 (12 drives)	3.40 in. (8.64 cm)	19.0 in. (48.26 cm)	21.75 in. (55.25 cm)	59.5 lb (27 kg)
SG5660 (60 drives)	7.00 in. (17.78 cm)	17.75 in. (45.08 cm)	32.50 in. (82.55 cm)	236.2 lb. (107.1 kg)

3. Install any required network switches. See the NetApp Interoperability Matrix Tool for compatibility information.

#### Related information

[NetApp Hardware Universe](#)

[NetApp Interoperability](#)

#### Unpacking the boxes (SG5600)

Before installing the StorageGRID appliance, unpack all boxes and compare the contents to the items on the packing slip.

- **SG5660 enclosure, a 4U chassis with 60 drives**




- **SG5612 enclosure, a 2U chassis with 12 drives**



- **4U bezel or 2U endcaps**



- **NL-SAS drives**



Drives are preinstalled in the 2U SG5612, but not in the 4U SG5660 for shipment safety.

- **E5600SG controller**




- **E2700 controller**



- **Mounting rails and screws**



- **Enclosure handles (4U enclosures only)**



### Cables and connectors

The shipment for the StorageGRID appliance includes the following cables and connectors:

- **Power cords for your country**



The appliance ships with two AC power cords for connecting to an external power source, such as a wall plug. Your cabinet might have special power cords that you use instead of the power cords that ship with the appliance.

- **SAS interconnect cables**



Two 0.5-meter SAS interconnect cables with mini-SAS-HD and mini-SAS connectors.

The square connector plugs into the E2700 controller, and the rectangular connector plugs into the E5600SG controller.

### Obtaining additional equipment and tools (SG5600)

Before installing the SG5600 appliance, confirm you have all of the additional equipment and tools that you need.

- **Screwdrivers**



Phillips No. 2 screwdriver

Medium flat-blade screwdrivers

- **ESD wrist strap**



- **Ethernet cables**



- **Ethernet switch**



- **Service laptop**



## Service laptop requirements

Before you install the StorageGRID appliance hardware, you should check to see if the service laptop has the minimum required resources.

The service laptop, which is needed for the hardware installation, must meet the following requirements:

- Microsoft Windows operating system
- Network port
- Supported web browser
- NetApp SANtricity Storage Manager version 11.40 or later
- SSH client (for example, PuTTY)

## Related information

[Web browser requirements](#)

[NetApp Documentation: SANtricity Storage Manager](#)

## Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## Reviewing appliance network connections

Before installing the StorageGRID appliance, you should understand which networks can be connected to the appliance and how the ports on each controller are used.

### StorageGRID appliance networks

When you deploy a StorageGRID appliance as a Storage Node, you can connect it to the following networks:

- **Grid Network for StorageGRID:** The Grid Network is used for all internal StorageGRID traffic. It provides connectivity between all nodes in the grid, across all sites and subnets. The Grid Network is required.


- **Admin Network for StorageGRID:** The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites. The Admin Network is optional.
- **Client Network for StorageGRID:** The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network provides client protocol access to the grid, so the Grid Network can be isolated and secured. The Client Network is optional.
- **Management network for SANtricity Storage Manager:** The E2700 controller connects to the management network where SANtricity Storage Manager is installed, allowing you to monitor and manage the hardware components in the appliance. This management network can be the same as the Admin Network for StorageGRID, or it can be an independent management network.



For detailed information about StorageGRID networks, see the *Grid Primer*.

### StorageGRID appliance connections

When you install a StorageGRID appliance, you must connect the two controllers to each other and to the required networks. The figure shows the two controllers in the SG5660, with the E2700 controller on the top and the E5600SG controller on the bottom. In the SG5612, the E2700 controller is to the left of the E5600SG controller.



Item	Port	Type of port	Function
1	Management port 1 on the E2700 controller	1-Gb (RJ-45) Ethernet	Connects the E2700 controller to the network where SANtricity Storage Manager is installed.
2	Management port 2 on the E2700 controller	1-Gb (RJ-45) Ethernet	Connects the E2700 controller to a service laptop during installation.

Item	Port	Type of port	Function
3	Two SAS interconnect ports on each controller, labelled Ch 1 and Ch 2	E2700 controller: mini-SAS-HD  E5600SG controller: mini-SAS	Connect the two controllers to each other.
4	Management port 1 on the E5600SG controller	1-Gb (RJ-45) Ethernet	Connects the E5600SG controller to the Admin Network for StorageGRID.
5	Management port 2 on the E5600SG controller	1-Gb (RJ-45) Ethernet	<ul style="list-style-type: none"> <li>• Can be bonded with management port 1 if you want a redundant connection to the Admin Network.</li> <li>• Can be left unwired and available for temporary local access (IP 169.254.0.1).</li> <li>• Can be used to connect the E5600SG controller to a service laptop during installation, if a DHCP-assigned IP address is not available.</li> </ul>
6	Four network ports on the E5600SG controller	10-GbE (optical)	Connect to the Grid Network and the Client Network for StorageGRID. See “10-GbE port connections for the E5600SG controller.”

### Related information

[Port bond modes for the E5600SG controller ports](#)

[Gathering installation information \(SG5600\)](#)

[Cabling the appliance \(SG5600\)](#)

[Network guidelines](#)

[Install VMware](#)

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

[Install Ubuntu or Debian](#)

## Port bond modes for the E5600SG controller ports

When configuring network links for the E5600SG controller ports, you can use port bonding for the 10-GbE ports that connect to the Grid Network and optional Client Network, and the 1-GbE management ports that connect to the optional Admin Network. Port bonding helps protect your data by providing redundant paths between StorageGRID networks and the appliance.

### Related information


[Configuring network links \(SG5600\)](#)

## Network bond modes for the 10-GbE ports

The 10-GbE networking ports on the E5600SG controller support Fixed port bond mode or Aggregate port bond mode for the Grid Network and Client Network connections.

### Fixed port bond mode

Fixed mode is the default configuration for the 10-GbE networking ports.



	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

When using Fixed port bond mode, the ports can be bonded using active-backup mode or Link Aggregation Control Protocol mode (LACP 802.3ad).


- In active-backup mode (default), only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Port 4 provides a backup path for port 2 (Grid Network), and port 3 provides a backup path for port 1 (Client Network).
- In LACP mode, each pair of ports forms a logical channel between the controller and the network, allowing for higher throughput. If one port fails, the other port continues to provide the channel. Throughput is reduced, but connectivity is not impacted.



If you do not need redundant connections, you can use only one port for each network. However, be aware that an alarm will be raised in the Grid Manager after StorageGRID is installed, indicating that a cable is unplugged. You can safely acknowledge this alarm to clear it.

## Aggregate port bond mode

Aggregate port bond mode significantly increases the throughput for each StorageGRID network and provides additional failover paths.



	Which ports are bonded
1	All connected ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

If you plan to use Aggregate port bond mode:

- You must use LACP network bond mode.
- You must specify a unique VLAN tag for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.
- The ports must be connected to switches that can support VLAN and LACP. If multiple switches are participating in the LACP bond, the switches must support multi-chassis link aggregation groups (MLAG), or equivalent.
- You must understand how to configure the switches to use VLAN, LACP, and MLAG, or equivalent.

If you do not want to use all four 10-GbE ports, you can use one, two, or three ports. Using more than one port maximizes the chance that some network connectivity will remain available if one of the 10-GbE ports fails.



If you choose to use fewer than four ports, be aware that one or more alarms will be raised in the Grid Manager after StorageGRID is installed, indicating that cables are unplugged. You can safely acknowledge the alarms to clear them.

## Network bond modes for the 1-GbE management ports

For the two 1-GbE management ports on the E5600SG controller, you can choose Independent network bond mode or Active-Backup network bond mode to connect to the optional Admin Network.

In Independent mode, only management port 1 is connected to the Admin Network. This mode does not provide a redundant path. Management port 2 is left unwired and available for temporary local connections (use IP address 169.254.0.1)

In Active-Backup mode, both management ports 1 and 2 are connected to the Admin Network. Only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Bonding these two physical ports into one logical management port provides a redundant path to the Admin Network.



If you need to make a temporary local connection to the E5600SG controller when the 1-GbE management ports are configured for Active-Backup mode, remove the cables from both management ports, plug your temporary cable into management port 2, and access the appliance using IP address 169.254.0.1.



### Gathering installation information (SG5600)

As you install and configure the StorageGRID appliance, you must make decisions and gather information about Ethernet switch ports, IP addresses, and port and network bond modes.

#### About this task

You can use the following tables to record information for each network you connect to the appliance. These values are required to install and configure the hardware.

#### Information needed to connect the E2700 controller to SANtricity Storage Manager

You must connect the E2700 controller to the management network you will use for SANtricity Storage Manager.

Information needed	Your value
Ethernet switch port you will connect to management port 1	
MAC address for management port 1 (printed on a label near port P1)	
DHCP-assigned IP address for management port 1, if available after power on  <b>Note:</b> If the network you will connect to the E2700 controller includes a DHCP server, the network administrator can use the MAC address to determine the IP address that was assigned by the DHCP server.	
Speed and duplex mode  <b>Note:</b> You must make sure the Ethernet switch for the SANtricity Storage Manager management network is set to autonegotiate.	Must be: <ul style="list-style-type: none"> <li>• Autonegotiate (default)</li> </ul>

Information needed	Your value
IP address format	Choose one: <ul style="list-style-type: none"> <li>• IPv4</li> <li>• IPv6</li> </ul>
Static IP address you plan to use for the appliance on the management network	For IPv4: <ul style="list-style-type: none"> <li>• IPv4 address:</li> <li>• Subnet mask:</li> <li>• Gateway:</li> </ul> For IPv6: <ul style="list-style-type: none"> <li>• IPv6 address:</li> <li>• Routable IP address:</li> <li>• E2700 controller router IP address:</li> </ul>

**Information needed to connect the E5600SG controller to the Admin Network**

The Admin Network for StorageGRID is an optional network, used for system administration and maintenance. The appliance connects to the Admin Network using the 1-GbE management ports on the E5600SG controller.

Information needed	Your value
Admin Network enabled	Choose one: <ul style="list-style-type: none"> <li>• No</li> <li>• Yes (default)</li> </ul>
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>• Independent</li> <li>• Active-Backup</li> </ul>
Switch port for management port 1 (P1)	
Switch port for management port 2 (P2; Active-Backup network bond mode only)	
MAC address for management port 1 (printed on a label near port P1)	

Information needed	Your value
DHCP-assigned IP address for management port 1, if available after power on  <b>Note:</b> If the Admin Network includes a DHCP server, the E5600SG controller displays the DHCP-assigned IP address on its seven-segment display after it boots up. You can also determine the DHCP-assigned IP address by using the MAC address to look up the assigned IP.	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Static IP address you plan to use for the appliance Storage Node on the Admin Network  <b>Note:</b> If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Admin Network subnets (CIDR)	

**Information needed to connect and configure the 10-GbE ports on the E5600SG controller**

The four 10-GbE ports on the E5600SG controller connect to the StorageGRID Grid Network and Client Network.



See "10-GbE port connections for the E5600SG controller" for more information about the options for these ports.

Information needed	Your value
Port bond mode	Choose one: <ul style="list-style-type: none"> <li>• Fixed (default)</li> <li>• Aggregate</li> </ul>
Switch port for port 1 (Client Network for Fixed mode)	
Switch port for port 2 (Grid Network for Fixed mode)	
Switch port for port 3 (Client Network for Fixed mode)	
Switch port for port 4 (Grid Network for Fixed mode)	

**Information needed to connect the E5600SG controller to the Grid Network**

The Grid Network for StorageGRID is a required network, used for all internal StorageGRID traffic. The appliance connects to the Grid Network using the 10-GbE ports on the E5600SG controller.



See "10-GbE port connections for the E5600SG controller" for more information about the options for these ports.

Information needed	Your value
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>• Active-Backup (default)</li> <li>• LACP (802.3ad)</li> </ul>
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>
VLAN tag(if VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Grid Network, if available after power on  <b>Note:</b> If the Grid Network includes a DHCP server, the E5600SG controller displays the DHCP-assigned IP address for the Grid Network on its seven-segment display after it boots up.	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Static IP address you plan to use for the appliance Storage Node on the Grid Network  <b>Note:</b> If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Grid Network subnets (CIDR)  <b>Note:</b> If the Client Network is not enabled, the default route on the controller will use the gateway specified here.	

#### Information needed to connect the E5600SG controller to the Client Network

The Client Network for StorageGRID is an optional network, used to provides client protocol access to the grid. The appliance connects to the Client Network using the 10-GbE ports on the E5600SG controller.



See "10-GbE port connections for the E5600SG controller" for more information about the options for these ports.

Information needed	Your value
Client Network enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>• Active-Backup (default)</li> <li>• LACP (802.3ad)</li> </ul>
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>
VLAN tag(if VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Client Network, if available after power on	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Static IP address you plan to use for the appliance Storage Node on the Client Network  <b>Note:</b> If the Client Network is enabled, the default route on the controller will use the gateway specified here.	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>

### Related information

[Reviewing appliance network connections](#)

[Configuring the hardware](#)

[Port bond modes for the E5600SG controller ports](#)

## Installing the hardware

Hardware installation includes several major tasks, including installing hardware components, cabling those components, and configuring ports.

### Steps

- [Registering the hardware](#)
- [Installing the appliance in a cabinet or rack \(SG5600\)](#)
- [Cabling the appliance \(SG5600\)](#)
- [Connecting the AC power cords \(SG5600\)](#)
- [Turning power on \(SG5600\)](#)

- [Viewing boot-up status and reviewing error codes on the SG5600 controllers](#)

## Registering the hardware

Registering the appliance hardware provides support benefits.

### Steps

1. Locate the chassis serial number.

You can find the number on the packing slip, in your confirmation email, or on the appliance after you unpack it.



2. Go to the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com).
3. Determine whether you need to register the hardware:

If you are a...	Follow these steps...
Existing NetApp customer	<ol style="list-style-type: none"> <li>a. Sign in with your username and password.</li> <li>b. Select <b>Products &gt; My Products</b>.</li> <li>c. Confirm that the new serial number is listed.</li> <li>d. If it is not, follow the instructions for new NetApp customers.</li> </ol>
New NetApp customer	<ol style="list-style-type: none"> <li>a. Click <b>Register Now</b>, and create an account.</li> <li>b. Select <b>Products &gt; Register Products</b>.</li> <li>c. Enter the product serial number and requested details.</li> </ol> <p>After your registration is approved, you can download any required software. The approval process might take up to 24 hours.</p>

## Installing the appliance in a cabinet or rack (SG5600)

You must install rails in your cabinet or rack and then slide the appliance onto the rails. If you have an SG5660, you must also install the drives after installing the appliance.

### What you'll need

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the E-Series installation instructions for the hardware.



Install hardware from the bottom of the rack or cabinet or rack up to prevent the equipment from tipping over.



The SG5612 weighs approximately 60 lb (27 kg) when fully loaded with drives. Two people or a mechanized lift are required to safely move the SG5612.



The SG5660 weighs approximately 132 lb (60 kg) with no drives installed. Four people or a mechanized lift are required to safely move an empty SG5660.



To avoid damaging the hardware, never move an SG5660 if drives are installed. You must remove all drives before moving the appliance.

### About this task

Complete the following tasks to install the SG5660 appliance in a cabinet or rack.

#### • Install the mounting rails

Install the mounting rails in the cabinet or rack.

See the E-Series installation instructions for the E2700 or the E5600.

#### • Install the appliance in the cabinet or rack

Slide the appliance into the cabinet or rack, and secure it.




If you are lifting the SG5660 by hand, attach the four handles to the sides of the chassis. You remove these handles as you slide the appliance onto the rails.

#### • Install the drives

If you have an SG5660, install 12 drives in each of the 5 drive drawers.

You must install all 60 drives to ensure correct operation.

- a. Put on the ESD wristband, and remove the drives from their packaging.
- b. Release the levers on the top drive drawer, and slide the drawer out using the levers.
- c. Raise the drive handle to vertical, and align the buttons on the drive with the notches on the drawer.



- d. Pressing gently on the top of the drive, rotate the drive handle down until the drive snaps into place.

e. After installing the first 12 drives, slide the drawer back in by pushing on the center and closing both levers gently.

f. Repeat these steps for the other four drawers.

- **Attach the front bezel**

**SG5612:** Attach the left and right end caps to the front.

**SG5660:** Attach the bezel to the front.

#### **Related information**

[E2700 Controller-Drive Tray and Related Drive Trays Installation Guide](#)

[E5600 Controller-Drive Tray and Related Drive Trays Installation Guide](#)

#### **Cabling the appliance (SG5600)**

You must connect the two controllers to each other with SAS interconnect cables, connect the management ports to the appropriate management network, and connect the 10 GbE ports on the E5600SG controller to the Grid Network and optional Client Network for StorageGRID.

#### **What you'll need**


- You have Ethernet cables for connecting the management ports.
- You have optical cables for connecting the four 10-GbE ports (these are not provided with the appliance).



**Risk of exposure to laser radiation** — Do not disassemble or remove any part of an SFP transceiver. You might be exposed to laser radiation.

#### **About this task**

When connecting the cables, refer to the following diagram, which shows the E2700 controller on the top and the E5600SG controller on the bottom. The diagram shows the SG5660 model; the controllers in the SG5612 model are side by side instead of stacked.



Item	Port	Type of port	Function
1	Management port 1 on the E2700 controller	1-Gb (RJ-45) Ethernet	Connects the E2700 controller to the network where SANtricity Storage Manager is installed.
2	Management port 2 on the E2700 controller	1-Gb (RJ-45) Ethernet	Connects the E2700 controller to a service laptop during installation.
3	Two SAS interconnect ports on each controller, labelled Ch 1 and Ch 2	E2700 controller: mini-SAS-HD E5600SG controller: mini-SAS	Connect the two controllers to each other.
4	Management port 1 on the E5600SG controller	1-Gb (RJ-45) Ethernet	Connects the E5600SG controller to the Admin Network for StorageGRID.

Item	Port	Type of port	Function
5	Management port 2 on the E5600SG controller	1-Gb (RJ-45) Ethernet	<ul style="list-style-type: none"> <li>• Can be bonded with management port 1 if you want a redundant connection to the Admin Network.</li> <li>• Can be left unwired and available for temporary local access (IP 169.254.0.1).</li> <li>• Can be used to connect the E5600SG controller to a service laptop during installation if DHCP-assigned IP addresses are not available.</li> </ul>
6	Four network ports on the E5600SG controller	10-GbE (optical)	Connect the E5600SG controller to the Grid Network and to the Client Network (if used) for StorageGRID. The ports can be bonded together to provide redundant paths to the controller.

### Steps

1. Connect the E2700 controller to the E5600SG controller, using the two SAS interconnect cables.

Connect this port...	To this port...
SAS interconnect port 1 (labeled Ch 1) on the E2700 controller	SAS interconnect port 1 (labeled Ch 1) on the E5600SG controller
SAS interconnect port 2 (labeled Ch 2) on the E2700 controller	SAS interconnect port 2 (labeled Ch 2) on the E5600SG controller

Use the square connector (mini-SAS HD) for the E2700 controller, and use the rectangular connector (mini-SAS) for the E5600SG controller.



Make sure the pull tabs on the SAS connectors are at the bottom, and carefully insert each connector until it clicks into place. Do not push on the connector if there is any resistance. Verify the position of the pull tab before continuing.

2. Connect the E2700 controller to the management network where SANtricity Storage Manager software is installed, using an Ethernet cable.

Connect this port...	To this port...
Port 1 on the E2700 controller (the RJ-45 port on the left)	Switch port on the management network used for SANtricity Storage Manager
Port 2 on the E2700 controller	Service laptop, if not using DHCP

3. If you plan to use the Admin Network for StorageGRID, connect the E5600SG controller, using an Ethernet cable.

Connect this port...	To this port...
Port 1 on the E5600SG controller (the RJ-45 port on the left)	Switch port on the Admin Network for StorageGRID
Port 2 on the E5600SG controller	Service laptop, if not using DHCP

4. Connect the 10-GbE ports on the E5600SG controller to the appropriate network switches, using optical cables and SFP+ transceivers.
- If you plan to use Fixed port bond mode (default), connect the ports to the StorageGRID Grid and Client Networks, as shown in the table.

Port	Connects to...
Port 1	Client Network (optional)
Port 2	Grid Network
Port 3	Client Network (optional)
Port 4	Grid Network

- If you plan to use the Aggregate port bond mode, connect one or more of the network ports to one or more switches. You should connect at least two of the four ports to avoid having a single point of failure. If you use more than one switch for a single LACP bond, the switches must support MLAG or equivalent.

### Related information

[Port bond modes for the E5600SG controller ports](#)

[Accessing the StorageGRID Appliance Installer](#)

### Connecting the AC power cords (SG5600)

You must connect the AC power cords to the external power source and to the AC power connector on each controller. After you have connected the power cords, you can turn the power on.

### What you'll need

Both appliance power switches must be off before connecting power.



**Risk of electrical shock**— Before connecting the power cords, make sure that the two power switches on the appliance are off.

### About this task

- You should use separate power sources for each power supply.

Connecting to independent power sources maintains power redundancy.

- You can use the power cords shipped with the controller with typical outlets used in the destination country, such as wall receptacles of an uninterrupted power supply (UPS).

However, these power cords are not intended for use in most EIA-compliant cabinets.

### Steps

1. Turn off the power switches in the enclosure or chassis.
2. Turn off the power switches on the controllers.
3. Connect the primary power cords from the cabinet to the external power sources.
4. Connect the power cords to the AC power connector on each controller.

### Turning power on (SG5600)

Powering on the enclosure provides power to both controllers.

### Steps

1. Turn on the two power supply switches at the rear of the enclosure.

While the power is being applied, the LEDs on the controllers go on and off intermittently.

The power-on process can take up to ten minutes to complete. The controllers reboot several times during the initial startup sequence, which causes the fans to ramp up and down and the LEDs to flash.

2. Check the Power LED and the Host Link Active LEDs on each controller to verify that the power was turned on.
3. Wait for all drives to show a persistent green LED, indicating that they have come online.
4. Check for green LEDs on the front and rear of the enclosure.

If you see any amber LEDs, make a note of their locations.

5. Look at the seven-segment display for the E5600SG controller.

This display shows **HO**, followed by a repeating sequence of two digits.

```
HO -- IP address for Admin Network -- IP address for Grid Network HO
```

In the sequence, the first set of numbers is the DHCP-assigned IP address for the controller's management port 1. This address is used to connect the controller to the Admin Network for StorageGRID. The second set of numbers is the DHCP-assigned IP address used to connect the appliance to the Grid Network for

StorageGRID.



If an IP address could not be assigned using DHCP, 0.0.0.0 is displayed.

## Viewing boot-up status and reviewing error codes on the SG5600 controllers

The seven-segment display on each controller shows status and error codes when the appliance powers up, while the hardware is initializing, and when the hardware fails and must back out of the initialization. If you are monitoring the progress or troubleshooting, you should watch the sequence of the codes as they appear.

### About this task

The status and error codes for the E5600SG controller are not the same as those for the E2700 controller.

### Steps

1. During boot-up, view the codes shown on the seven-segment displays to monitor progress.
2. To review error codes for the E5600SG controller, see the seven-segment display status and error code information.
3. To review error codes for the E2700 controller, see the E2700 controller documentation on the Support Site.

### Related information

[E5600SG controller seven-segment display codes](#)

[NetApp Documentation: E2700 Series](#)

### E5600SG controller seven-segment display codes

The seven-segment display on the E5600SG controller shows status and error codes while the appliance powers up and while the hardware is initializing. You can use these codes to determine status and troubleshoot errors.

When reviewing status and error codes on the E5600SG controller, you should look at the following types of codes:

- **General boot-up codes**

Represent the standard boot-up events.

- **Normal boot-up codes**


Represent the normal boot-up events that occur in the appliance.

- **Error codes**

Indicate issues during the boot-up events.

StorageGRID controls only the following LEDs on the E5600SG controller and only after the StorageGRID Appliance Installer has started:

- Service Action Allowed LED
- Service Action Required LED
- Seven-segment display



The decimal points on the seven-segment display are not used by the StorageGRID appliance:

- The upper decimal point adjacent to the least significant digit is the platform diagnostic LED. This is turned on during reset and initial hardware configuration. Otherwise, it is turned off.
- The lower decimal point adjacent to the most significant digit is turned off.

To diagnose other issues, you might want to look at these resources:

- To see all other hardware and environmental diagnostic information, see the E-Series operating system hardware diagnostics.

This includes looking for hardware issues such as power, temperature, and disk drives. The appliance relies on the E-Series operating system to monitor all platform environmental statuses.

- To determine firmware and driver issues, look at the link lights on the SAS and network ports.

For details, see the E-Series E5600 documentation.

### General boot-up codes

During boot-up or after a hard reset of the hardware, the Service Action Allowed and the Service Action Required LEDs come on while the hardware is initializing. The seven-segment display shows a sequence of codes that are the same for E-Series hardware and not specific to the E5600SG controller.

During boot-up, the Field Programmable Gate Array (FPGA) controls the functions and initialization on the hardware.

Code	Indication
19	FPGA initialization.
68	FPGA initialization.
...	FPGA initialization. This is a quick succession of codes.
AA	Platform BIOS booting.
FF	Bios boot-up complete. This is an intermediate state before E5600SG controller initializes and manages LEDs to indicate status.

After the AA and FF codes appear, either the normal boot-up codes appear or error codes appear. Additionally, the Service Action Allowed and the Service Action Required LEDs are turned off.

### Normal boot-up codes

These codes represent the normal boot-up events that occur in the appliance, in chronological order.

Code	Indication
HI	The master boot script has started.
PP	The platform FPGA firmware is checking for updates.
HP	The host interface card (HIC) is checking for updates.
RB	After firmware updates, the system is rebooting, if necessary.
FP	The firmware update checks have been completed. Starting the process (utmagent) to communicate with and manage the E2700 controller. This process facilitates appliance provisioning.
HE	The system is synchronizing with the E-Series operating system.
HC	The StorageGRID installation is being checked.
HO	Installation management and active interfacing are occurring.
HA	The Linux operating system and StorageGRID are running.

## E5600SG controller error codes

These codes represent error conditions that might be shown on the E5600SG controller as the appliance boots up. Additional two-digit hexadecimal codes are displayed if specific low-level hardware errors occur. If any of these codes persists for more than a second or two, or if you are unable to resolve the error by following one of the prescribed troubleshooting procedures, contact technical support.

Code	Indication
22	No master boot record found on any boot device.
23	No SATA drive installed.
2A, 2B	Stuck bus, unable to read DIMM SPD data.
40	Invalid DIMMs.
41	Invalid DIMMs.
42	Memory test failed.
51	SPD reading failure.
92 to 96	PCI bus initialization.
A0 to A3	SATA drive initialization.
AB	Alternate boot code.
AE	Booting OS.
EA	DDR3 training failed.
E8	No memory installed.
EU	The installation script was not found.
EP	"ManageSGA" code indicates that pregrid communication with the E2700 controller failed.

### Related information

[Troubleshooting the hardware installation](#)

[NetApp Support](#)

## Configuring the hardware

After applying power to the appliance, you must configure SANtricity Storage Manager, which is the software you will use to monitor the hardware. You must also configure the network connections that will be used by StorageGRID.

### Steps

- [Configuring StorageGRID connections](#)
- [Configuring SANtricity Storage Manager](#)
- [Optional: Enabling node encryption](#)
- [Optional: Changing to RAID6 mode \(SG5660 only\)](#)
- [Optional: Remapping network ports for the appliance](#)

## Configuring StorageGRID connections

Before you can deploy a StorageGRID appliance as a Storage Node in a StorageGRID grid, you must configure the connections between the appliance and the networks you plan to use. You can configure networking by browsing to the StorageGRID Appliance Installer, which is included on the E5600SG controller (the compute controller in the appliance).

### Steps

- [Accessing the StorageGRID Appliance Installer](#)
- [Verifying and upgrading the StorageGRID Appliance Installer version](#)
- [Configuring network links \(SG5600\)](#)
- [Setting the IP configuration](#)
- [Verifying network connections](#)
- [Verifying port-level network connections](#)

### Accessing the StorageGRID Appliance Installer

You must access the StorageGRID Appliance Installer to configure the connections between the appliance and the three StorageGRID networks: the Grid Network, the Admin Network (optional), and the Client Network (optional).

### What you'll need

- You are using a supported web browser.
- The appliance is connected to all of the StorageGRID networks you plan to use.
- You know the IP address, gateway, and subnet for the appliance on these networks.
- You have configured the network switches you plan to use.

### About this task

When you first access the StorageGRID Appliance Installer, you can use the DHCP-assigned IP address for the Admin Network (assuming the appliance is connected to the Admin Network) or the DHCP-assigned IP address for the Grid Network. Using the IP address for the Admin Network is preferred. Otherwise, if you access the StorageGRID Appliance Installer using the DHCP address for the Grid Network, you might lose

connection with the StorageGRID Appliance Installer when you change link settings and when you enter a static IP.

## Steps

1. Obtain the DHCP address for the appliance on the Admin Network (if it is connected) or the Grid Network (if the Admin Network is not connected).

You can do either of the following:

- Provide the MAC address for management port 1 to your network administrator, so they can look up the DHCP address for this port on the Admin Network. The MAC address is printed on a label on the E5600SG controller, next to the port.
- Look at the seven-segment display on the E5600SG controller. If management port 1 and 10-GbE ports 2 and 4 on the E5600SG controller are connected to networks with DHCP servers, the controller attempts to obtain dynamically assigned IP addresses when you power on the enclosure. After the controller has completed the power-on process, its seven-segment display shows **HO**, followed by a repeating sequence of two numbers.

```
HO -- IP address for Admin Network -- IP address for Grid Network HO
```

In the sequence:

- The first set of numbers is the DHCP address for the appliance Storage Node on the Admin Network, if it is connected. This IP address is assigned to management port 1 on the E5600SG controller.
- The second set of numbers is the DHCP address for the appliance Storage Node on the Grid Network. This IP address is assigned to 10-GbE ports 2 and 4 when you first apply power to the appliance.



If an IP address could not be assigned using DHCP, 0.0.0.0 is displayed.

2. If you were able to obtain either of the DHCP addresses:

- a. Open a web browser on the service laptop.
- b. Enter this URL for the StorageGRID Appliance Installer:

**`https://E5600SG_Controller_IP:8443`**

For *E5600SG\_Controller\_IP*, use the DHCP address for the controller (use the IP address for the Admin Network if you have it).

- c. If you are prompted with a security alert, view and install the certificate using the browser's installation wizard.

The alert will not appear the next time you access this URL.

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected to StorageGRID networks. Error messages might appear that will be resolved in later steps.

## Home

**i** The installation is ready to be started. Review the settings below, and then click Start Installation.

### This Node

Node type

Node name

Cancel

Save

### Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

Connection state Connection to 172.16.1.178 ready

Cancel

Save

### Installation

Current state Ready to start installation of MM-2-108-SGA-lab25 into grid with Admin Node 172.16.1.178 running StorageGRID 11.2.0, using StorageGRID software downloaded from the Admin Node.

Start Installation

3. If the E5600SG controller could not acquire an IP address using DHCP:
  - a. Connect the service laptop to management port 2 on the E5600SG controller, using an Ethernet cable.



- b. Open a web browser on the service laptop.

c. Enter this URL for the StorageGRID Appliance Installer:

**https://169.254.0.1:8443**

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected.



If you cannot access the Home page over a link-local connection, configure the service laptop IP address as 169.254.0.2, and try again.

4. Review any messages displayed on the Home page and configure the link configuration and the IP configuration, as required.

### Related information

[Web browser requirements](#)

### Verifying and upgrading the StorageGRID Appliance Installer version

The StorageGRID Appliance Installer version on the appliance must match the software version installed on your StorageGRID system to ensure that all StorageGRID features are supported.

### What you'll need

You have accessed the StorageGRID Appliance Installer.

StorageGRID appliances come from the factory preinstalled with the StorageGRID Appliance Installer. If you are adding an appliance to a recently upgraded StorageGRID system, you might need to manually upgrade the StorageGRID Appliance Installer before installing the appliance as a new node.

The StorageGRID Appliance Installer automatically upgrades when you upgrade to a new StorageGRID version. You do not need to upgrade the StorageGRID Appliance Installer on installed appliance nodes. This procedure is only required when you are installing an appliance that contains an earlier version of the StorageGRID Appliance Installer.

### Steps

1. From the StorageGRID Appliance Installer, select **Advanced > Upgrade Firmware**.
2. Compare the Current Firmware version to the software version installed on your StorageGRID system (from the Grid Manager select **Help > About**).

The second digit in the two versions should match. For example, if your StorageGRID system is running version 11.5.x.y, the StorageGRID Appliance Installer version should be 3.5.z.

3. If the appliance has a down-level version of the StorageGRID Appliance Installer, go to the NetApp Downloads page for StorageGRID.

[NetApp Downloads: StorageGRID](#)

Sign in with the username and password for your NetApp account.

4. Download the appropriate version of the **Support file for StorageGRID Appliances** and the corresponding checksum file.

The Support file for StorageGRID Appliances file is a .zip archive that contains the current and previous

firmware versions for all StorageGRID appliance models, in subdirectories for each controller type.

After downloading the Support file for StorageGRID Appliances file, extract the .zip archive and see the README file for important information about installing the StorageGRID Appliance Installer.

5. Follow the instructions on the Upgrade Firmware page of the StorageGRID Appliance Installer to perform these steps:
  - a. Upload the appropriate support file (firmware image) for your controller type and the checksum file.
  - b. Upgrade the inactive partition.
  - c. Reboot and swap partitions.
  - d. Upgrade the second partition.

## Related information

[Accessing the StorageGRID Appliance Installer](#)

### Configuring network links (SG5600)

You can configure network links for the ports used to connect the appliance to the Grid Network, the Client Network, and the Admin Network. You can set the link speed as well as the port and network bond modes.


### What you'll need

If you plan to use Aggregate port bond mode, LACP network bond mode, or VLAN tagging:

- You have connected the 10-GbE ports on the appliance to switches that can support VLAN and LACP.
- If multiple switches are participating in the LACP bond, the switches support multi-chassis link aggregation groups (MLAG), or equivalent.
- You understand how to configure the switches to use VLAN, LACP, and MLAG or equivalent.
- You know the unique VLAN tag to use for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.

### About this task


This figure shows how the four 10-GbE ports are bonded in fixed port bond mode (default configuration).



	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.

	<b>Which ports are bonded</b>
G	Ports 2 and 4 are bonded together for the Grid Network.

This figure shows how the four 10-GbE ports are bonded in aggregate port bond mode.



	<b>Which ports are bonded</b>
1	All four ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

The table summarizes the options for configuring the four 10-GbE ports. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.

• **Fixed (default) port bond mode**

Network bond mode	Client Network disabled (default)	Client Network enabled
<b>Active-Backup (default)</b>	<ul style="list-style-type: none"> <li>Ports 2 and 4 use an active-backup bond for the Grid Network.</li> <li>Ports 1 and 3 are not used.</li> <li>A VLAN tag is optional.</li> </ul>	<ul style="list-style-type: none"> <li>Ports 2 and 4 use an active-backup bond for the Grid Network.</li> <li>Ports 1 and 3 use an active-backup bond for the Client Network.</li> <li>VLAN tags can be specified for both networks for the convenience of the network administrator.</li> </ul>
LACP (802.3ad)	<ul style="list-style-type: none"> <li>Ports 2 and 4 use an LACP bond for the Grid Network.</li> <li>Ports 1 and 3 are not used.</li> <li>A VLAN tag is optional.</li> </ul>	<ul style="list-style-type: none"> <li>Ports 2 and 4 use an LACP bond for the Grid Network.</li> <li>Ports 1 and 3 use an LACP bond for the Client Network.</li> <li>VLAN tags can be specified for both networks for the convenience of the network administrator.</li> </ul>

• **Aggregate port bond mode**

Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad) only	<ul style="list-style-type: none"> <li>Ports 1-4 use a single LACP bond for the Grid Network.</li> <li>A single VLAN tag identifies Grid Network packets.</li> </ul>	<ul style="list-style-type: none"> <li>Ports 1-4 use a single LACP bond for the Grid Network and the Client Network.</li> <li>Two VLAN tags allow Grid Network packets to be segregated from Client Network packets.</li> </ul>

See “10-GbE port connections for the E5600SG controller” for more information about port bond and network bond modes.

This figure shows how the two 1-GbE management ports on the E5600SG controller are bonded in Active-Backup network bond mode for the Admin Network.




### Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Link Configuration**.

The Network Link Configuration page displays a diagram of your appliance with the network and management ports numbered.

### Network Link Configuration



**⚠** You might lose your connection if you make changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL using one of the other IP addresses assigned to the appliance.

The Link Status table lists the link state (up/down) and speed (1/10/25/40/100 Gbps) of the numbered ports.

## Link Status

---

Link	State	Speed (Gbps)
1	Down	N/A
2	Up	10
3	Up	10
4	Down	N/A
5	Up	1
6	Up	1

The first time you access this page:

- **Link Speed** is set to **10GbE**. This is the only link speed available for the E5600SG controller.
- **Port bond mode** is set to **Fixed**.
- **Network bond mode** for the Grid Network is set to **Active-Backup**.
- The **Admin Network** is enabled, and the network bond mode is set to **Independent**.
- The **Client Network** is disabled.

## Link Settings

Link speed

Port bond mode  Fixed  Aggregate

Choose Fixed port bond mode if you want to use ports 2 and 4 for the Grid Network and ports 1 and 3 for the Client Network (if enabled). Choose Aggregate port bond mode if you want all connected ports to share a single LACP bond for both the Grid and Client Networks.

## Grid Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

Enable VLAN (802.1q) tagging

MAC Addresses 50:6b:4b:42:d7:00 50:6b:4b:42:d7:01 50:6b:4b:42:d7:24 50:6b:4b:42:d7:25

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

## Admin Network

Enable network

Network bond mode  Independent  Active-Backup

Connect the Admin Network to port 5. Leave port 6 unconnected. If necessary, you can make a temporary direct Ethernet connection to port 6 and use link-local IP address 169.254.0.1 for access.

MAC Addresses d8:c4:97:2a:e4:95

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

## Client Network

Enable network

Enabling the Client Network causes the default gateway for this node to move to the Client Network. Before enabling the Client Network, ensure that you've added all necessary subnets to the Grid Network Subnet List. Otherwise, the connection to the node might be lost.

2. Enable or disable the StorageGRID networks you plan to use.

The Grid Network is required. You cannot disable this network.

- a. If the appliance is not connected to the Admin Network, unselect the **Enable network** check box for the Admin Network.

## Admin Network

---

Enable network

- b. If the appliance is connected to the Client Network, select the **Enable network** check box for the Client Network.

The Client Network settings for the 10-GbE ports are now shown.

3. Refer to the table, and configure the port bond mode and the network bond mode.

The example shows:

- **Aggregate** and **LACP** selected for the Grid and the Client networks. You must specify a unique VLAN tag for each network. You can select values between 0 and 4095.
- **Active-Backup** selected for the Admin Network.

### Link Settings

Link speed

Port bond mode  Fixed  Aggregate

Choose Fixed port bond mode if you want to use ports 2 and 4 for the Grid Network and ports 1 and 3 for the Client Network (if enabled). Choose Aggregate port bond mode if you want all connected ports to share a single LACP bond for both the Grid and Client Networks.

### Grid Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

### Admin Network

Enable network

Network bond mode  Independent  Active-Backup

Connect the Admin Network to ports 5 and 6. If necessary, you can make a temporary direct Ethernet connection by disconnecting ports 5 and 6, then connecting to port 6 and using link-local IP address 169.254.0.1 for access.

### Client Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

4. When you are satisfied with your selections, click **Save**.



You might lose your connection if you made changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:

**`https://E5600SG_Controller_IP:8443`**

### Related information

[Port bond modes for the E5600SG controller ports](#)

### Setting the IP configuration

You use the StorageGRID Appliance Installer to configure the IP addresses and routing information used for the appliance Storage Node on the StorageGRID Grid, Admin, and

## Client Networks.

### About this task

You must either assign a static IP for the appliance on each connected network or assign a permanent lease for the address on the DHCP server.

If you want to change the link configuration, see the instructions for changing the link configuration of the E5600SG controller.

### Steps

1. In the StorageGRID Appliance Installer, select **Configure Networking > IP Configuration**.

The IP Configuration page appears.

2. To configure the Grid Network, select either **Static** or **DHCP** in the **Grid Network** section of the page.

## Grid Network

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

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

 All required Grid Network subnets must also be defined in the Grid Network Subnet List on the Primary Admin Node before starting installation.

Subnets (CIDR)  



MTU  

3. If you selected **Static**, follow these steps to configure the Grid Network:

- Enter the static IPv4 address, using CIDR notation.
- Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

- If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

**`https://services_appliance_IP:8443`**

e. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

f. Click **Save**.

4. If you selected **DHCP**, follow these steps to configure the Grid Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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

d. Click **Save**.

5. To configure the Admin Network, select either **Static** or **DHCP** in the Admin Network section of the page.



To configure the Admin Network, you must enable the Admin Network on the Link Configuration page.

### Admin Network

The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites.

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

Subnets (CIDR)  +

MTU

6. If you selected **Static**, follow these steps to configure the Admin Network:

a. Enter the static IPv4 address, using CIDR notation, for Management Port 1 on the appliance.

Management Port 1 is the left of the two 1-GbE RJ45 ports on the right end of the appliance.

b. Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as

9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

**`https://services_appliance:8443`**

e. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route cannot be made to use the Admin Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

f. Click **Save**.

7. If you selected **DHCP**, follow these steps to configure the Admin Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route cannot be made to use the Admin Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

8. To configure the Client Network, select either **Static** or **DHCP** in the **Client Network** section of the page.



To configure the Client Network, you must enable the Client Network on the Link Configuration page.

## Client Network

The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network enables grid nodes to communicate with any subnet reachable through the Client Network gateway. The Client Network does not become operational until you complete the StorageGRID configuration steps.

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

MTU

9. If you selected **Static**, follow these steps to configure the Client Network:

- Enter the static IPv4 address, using CIDR notation.
- Click **Save**.
- Confirm that the IP address for the Client Network gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and cannot be moved to another interface while the Client Network is enabled.

- If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

- Click **Save**.

10. If you selected **DHCP**, follow these steps to configure the Client Network:

- After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address** and **Gateway** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and cannot be moved to another interface while the Client Network is enabled.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

## Related information

[Changing the link configuration of the E5600SG controller](#)

## Verifying network connections

You should confirm you can access the StorageGRID networks you are using from the appliance. To validate routing through network gateways, you should test connectivity between the StorageGRID Appliance Installer and IP addresses on different subnets. You can also verify the MTU setting.

## Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Ping and MTU Test**.

The Ping and MTU Test page appears.

### Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

#### Ping and MTU Test

Network	<input type="text" value="Grid"/>
Destination IPv4 Address or FQDN	<input type="text"/>
Test MTU	<input type="checkbox"/>
<input type="button" value="Test Connectivity"/>	

2. From the **Network** drop-down box, select the network you want to test: Grid, Admin, or Client.

3. Enter the IPv4 address or fully qualified domain name (FQDN) for a host on that network.

For example, you might want to ping the gateway on the network or the primary Admin Node.

4. Optionally, select the **Test MTU** check box to verify the MTU setting for the entire path through the network to the destination.

For example, you can test the path between the appliance node and a node at a different site.

5. Click **Test Connectivity**.

If the network connection is valid, the "Ping test passed" message appears, with the ping command output listed.

### Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

#### Ping and MTU Test

Network	<input type="text" value="Grid"/>
Destination IPv4 Address or FQDN	<input type="text" value="10.96.104.223"/>
Test MTU	<input checked="" type="checkbox"/>
<input type="button" value="Test Connectivity"/>	

Ping test passed

#### Ping command output

```
PING 10.96.104.223 (10.96.104.223) 1472(1500) bytes of data.  
1480 bytes from 10.96.104.223: icmp_seq=1 ttl=64 time=0.318 ms  
  
--- 10.96.104.223 ping statistics ---  
1 packets transmitted, 1 received, 0% packet loss, time 0ms  
rtt min/avg/max/mdev = 0.318/0.318/0.318/0.000 ms  
  
Found MTU 1500 for 10.96.104.223 via br0
```

### Related information

[Configuring network links \(SG5600\)](#)

[Changing the MTU setting](#)

### Verifying port-level network connections

To ensure that access between the StorageGRID Appliance Installer and other nodes is

not obstructed by firewalls, confirm that the StorageGRID Appliance Installer can connect to a specific TCP port or set of ports at the specified IP address or range of addresses.

### About this task

Using the list of ports provided in the StorageGRID Appliance Installer, you can test the connectivity between the appliance and the other nodes in your Grid Network.

Additionally, you can test connectivity on the Admin and Client Networks and on UDP ports, such as those used for external NFS or DNS servers. For a list of these ports, see the port reference in the StorageGRID networking guidelines.



The Grid Network ports listed in the port connectivity table are valid only for StorageGRID version 11.5.0. To verify which ports are correct for each node type, you should always consult the networking guidelines for your version of StorageGRID.

### Steps

1. From the StorageGRID Appliance Installer, click **Configure Networking > Port Connectivity Test (nmap)**.

The Port Connectivity Test page appears.

The port connectivity table lists node types that require TCP connectivity on the Grid Network. For each node type, the table lists the Grid Network ports that should be accessible to your appliance.

The following node types require TCP connectivity on the Grid Network.

Node Type	Grid Network Ports
Admin Node	22,80,443,1504,1505,1506,1508,7443,9999
Storage Node without ADC	22,1139,1502,1506,1511,7001,9042,9999,18002,18017,18019,18082,18083,18200
Storage Node with ADC	22,1139,1501,1502,1506,1511,7001,9042,9999,18000,18001,18002,18003,18017,18019,18082,18083,18200,19000
API Gateway	22,1506,1507,9999
Archive Node	22,1506,1509,9999,11139

You can test the connectivity between the appliance ports listed in the table and the other nodes in your Grid Network.

2. From the **Network** drop-down, select the network you want to test: **Grid**, **Admin**, or **Client**.
3. Specify a range of IPv4 addresses for the hosts on that network.

For example, you might want to probe the gateway on the network or the primary Admin Node.

Specify a range using a hyphen, as shown in the example.

4. Enter a TCP port number, a list of ports separated by commas, or a range of ports.

The following node types require TCP connectivity on the Grid Network.

Node Type	Grid Network Ports
Admin Node	22,80,443,1504,1505,1506,1508,7443,9999
Storage Node without ADC	22,1139,1502,1506,1511,7001,9042,9999,18002,18017,18019,18082,18083,18200
Storage Node with ADC	22,1139,1501,1502,1506,1511,7001,9042,9999,18000,18001,18002,18003,18017,18019,18082,18083,18200,19000
API Gateway	22,1506,1507,9999
Archive Node	22,1506,1509,9999,11139

### Port Connectivity Test

Network

IPv4 Address Ranges

Port Ranges

Protocol  TCP  UDP

### 5. Click **Test Connectivity**.

- If the selected port-level network connections are valid, the “Port connectivity test passed” message appears in a green banner. The nmap command output is listed below the banner.

Port connectivity test passed

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Fri Nov 13 18:32:03 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,2022 10.224.6.160-161
Nmap scan report for 10.224.6.160
Host is up (0.00072s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp  open  down

Nmap scan report for 10.224.6.161
Host is up (0.00060s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp  open  down

# Nmap done at Fri Nov 13 18:32:04 2020 -- 2 IP addresses (2 hosts up) scanned in 0.55 seconds
```

- If a port-level network connection is made to the remote host, but the host is not listening on one or more of the selected ports, the “Port connectivity test failed” message appears in a yellow banner. The nmap command output is listed below the banner.

Any remote port the host is not listening to has a state of “closed.” For example, you might see this yellow banner when the node you are trying to connect to is in a pre-installed state and the StorageGRID NMS service is not yet running on that node.

 Port connectivity test failed  
Connection not established. Services might not be listening on target ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:07:02 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,80,443,1504,1505,1506,1508,7443,9999
Nmap scan report for 172.16.4.71
Host is up (0.00020s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
80/tcp    open  http
443/tcp    open  https
1504/tcp   closed evb-elm
1505/tcp   open  funkproxy
1506/tcp   open  utcd
1508/tcp   open  diagmond
7443/tcp   open  oracleas-https
9999/tcp   open  abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:07:03 2020 -- 1 IP address (1 host up) scanned in 0.59 seconds
```

- If a port-level network connection cannot be made for one or more selected ports, the “Port connectivity test failed” message appears in a red banner. The nmap command output is listed below the banner.

The red banner indicates that a TCP connection attempt to a port on the remote host was made, but nothing was returned to the sender. When no response is returned, the port has a state of “filtered” and is likely blocked by a firewall.



Ports with “closed” are also listed.

 Port connectivity test failed  
Connection failed to one or more ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:11:01 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,79,80,443,1504,1505,1506,1508,7443,9999 172.16.4.71
Nmap scan report for 172.16.4.71
Host is up (0.00029s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
79/tcp    filtered finger
80/tcp    open  http
443/tcp    open  https
1504/tcp   closed evb-elm
1505/tcp   open  funkproxy
1506/tcp   open  utcd
1508/tcp   open  diagmond
7443/tcp   open  oracleas-https
9999/tcp   open  abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:11:02 2020 -- 1 IP address (1 host up) scanned in 1.60 seconds
```

## Related information

[Network guidelines](#)

## Configuring SANtricity Storage Manager

You can use SANtricity Storage Manager to monitor the status of the storage disks and hardware components in your StorageGRID appliance. To access this software, you must

know the IP address of management port 1 on the E2700 controller (the storage controller in the appliance).

### Steps

- [Setting the IP address for the E2700 controller](#)
- [Adding the appliance to SANtricity Storage Manager](#)
- [Setting up SANtricity Storage Manager](#)

### Setting the IP address for the E2700 controller

Management port 1 on the E2700 controller connects the appliance to the management network for SANtricity Storage Manager. You must set a static IP address for the E2700 controller to ensure that you do not lose your management connection to the hardware and the controller firmware in the StorageGRID appliance.

### What you'll need

You are using a supported web browser.

### About this task

DHCP-assigned addresses could change at any time. Assign a static IP address to the controller to ensure consistent accessibility.

### Steps

1. From the client, enter the URL for the StorageGRID Appliance Installer:

**`https://E5600SG_Controller_IP:8443`**

For `E5600SG_Controller_IP`, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Hardware Configuration > Storage Controller Network Configuration**.

The Storage Controller Network Configuration page appears.

3. Depending on your network configuration, select **Enabled** for IPv4, IPv6, or both.
4. Make a note of the IPv4 address that is automatically displayed.

DHCP is the default method for assigning an IP address to this port.



It might take a few minutes for the DHCP values to appear.

IPv4 Address Assignment

Static  DHCP

IPv4 Address (CIDR)

10.224.5.166/21

Default Gateway

10.224.0.1

5. Optionally, set a static IP address for the E2700 controller management port.



You should either assign a static IP for the management port or assign a permanent lease for the address on the DHCP server.

- a. Select **Static**.
- b. Enter the IPv4 address, using CIDR notation.
- c. Enter the default gateway.

IPv4 Address Assignment     Static     DHCP

IPv4 Address (CIDR)	10.224.2.200/21
Default Gateway	10.224.0.1

- d. Click **Save**.

It might take a few minutes for your changes to be applied.

When you connect to SANtricity Storage Manager, you will use the new static IP address as the URL:  
**https://E2700\_Controller\_IP**

#### Related information

[NetApp Documentation: SANtricity Storage Manager](#)

#### Adding the appliance to SANtricity Storage Manager

You connect the E2700 controller in the appliance to SANtricity Storage Manager and then add the appliance as a storage array.

#### What you'll need

You are using a supported web browser.

#### About this task

For detailed instructions, see the SANtricity Storage Manager documentation.

#### Steps

1. Open a web browser, and enter the IP address as the URL for SANtricity Storage Manager:  
**https://E2700\_Controller\_IP**

The login page for SANtricity Storage Manager appears.

2. On the **Select Addition Method** page, select **Manual**, and click **OK**.
3. Select **Edit > Add Storage Array**.

The Add New Storage Array - Manual page appears.

4. In the **Out-of-band management** box, enter one of the following values:

- **Using DHCP:** The IP address assigned by the DHCP server to management port 1 on the E2700 controller
- **Not using DHCP:** 192.168.128.101



Only one of the appliance's controllers is connected to SANtricity Storage Manager, so you only need to enter one IP address.

5. Click **Add**.

#### Related information

[NetApp Documentation: SANtricity Storage Manager](#)

#### Setting up SANtricity Storage Manager

After accessing SANtricity Storage Manager, you can use it to configure hardware settings. Typically, you configure these settings before deploying the appliance as a Storage Node in a StorageGRID system.

#### Steps

- [Configuring AutoSupport](#)
- [Verifying receipt of AutoSupport](#)
- [Configuring email and SNMP trap alert notifications](#)
- [Setting passwords for SANtricity Storage Manager](#)

## Configuring AutoSupport

The AutoSupport tool collects data in a customer support bundle from the appliance and automatically sends the data to technical support. Configuring AutoSupport assists technical support with remote troubleshooting and problem analysis.

### What you'll need

- The AutoSupport feature must be enabled and activated on the appliance.

The AutoSupport feature is activated and deactivated globally on a storage management station.

- The Storage Manager Event Monitor must be running on at least one machine with access to the appliance and, preferably, on no more than one machine.

### About this task

All of the data is compressed into a single compressed archive file format (.7z) at the location you specify.

AutoSupport provides the following types of messages:

Message types	Description
Event messages	<ul style="list-style-type: none"> <li>• Sent when a support event on the managed appliance occurs</li> <li>• Include system configuration and diagnostic information</li> </ul>
Daily messages	<ul style="list-style-type: none"> <li>• Sent once every day during a user configurable time interval in the local time of the appliance</li> <li>• Include the current system event logs and performance data</li> </ul>
Weekly messages	<ul style="list-style-type: none"> <li>• Sent once every week during a user configurable time interval in the local time of the appliance</li> <li>• Include configuration and system state information</li> </ul>

### Steps

1. From the Enterprise Management Window in SANtricity Storage Manager, select the **Devices** tab, and then select **Discovered Storage Arrays**.
2. Select **Tools > AutoSupport > Configuration**.
3. Use SANtricity Storage Manager online help, if needed, to complete the task.

### Related information

## Verifying receipt of AutoSupport

You should verify that technical support is receiving your AutoSupport messages. You can find the status of AutoSupport for your systems on the Active IQ portal. Verifying receipt of these messages ensures that technical support has your information if you need assistance.

### About this task

AutoSupport can show one of the following statuses:

- **ON**

An ON status indicates that technical support is currently receiving AutoSupport messages from the system.

- **OFF**

An OFF status suggests that you might have disabled AutoSupport because technical support has not received a Weekly Log from the system in the last 15 calendar days or there might have been a change in your environment or configuration (as an example).

- **DECLINE**

A DECLINE status means that you have notified technical support that you will not enable AutoSupport.

After technical support receives a Weekly Log from the system, the AutoSupport status changes to ON.

### Steps

1. Go to the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com), and sign in to the Active IQ portal.
2. If the AutoSupport status is OFF, and you believe that is incorrect, complete the following:
  - a. Check your system configuration to ensure that you have turned AutoSupport on.
  - b. Check your network environment and configuration to ensure that the system can send messages to technical support.

## Configuring email and SNMP trap alert notifications

SANtricity Storage Manager can notify you when the status of the appliance or one of its components changes. This is called an alert notification. You can receive alert notifications by two different methods: email and SNMP traps. You must configure the alert notifications you want to receive.

### Steps

1. From the Enterprise Management Window in SANtricity Storage Manager, select the **Devices** tab, and then select a node.
2. Select **Edit > Configure Alerts**.
3. Select the **Email** tab to configure email alert notifications.
4. Select the **SNMP** tab to configure SNMP trap alert notifications.

5. Use SANtricity Storage Manager online help, if needed, to complete the task.

## Setting passwords for SANtricity Storage Manager

You can set the passwords used for the appliance in SANtricity Storage Manager. Setting passwords maintains system security.

### Steps

1. From the Enterprise Management Window in SANtricity Storage Manager, double-click the controller.
2. From the Array Management Window, select the **Storage Array** menu, and select **Security > Set Password**.
3. Configure the passwords.
4. Use SANtricity Storage Manager online help, if needed, to complete the task.

### Optional: Enabling node encryption

If you enable node encryption, the disks in your appliance can be protected by secure key management server (KMS) encryption against physical loss or removal from the site. You must select and enable node encryption during appliance installation and cannot unselect node encryption once the KMS encryption process starts.

### What you'll need

Review the information about KMS in the instructions for administering StorageGRID.

### About this task

An appliance that has node encryption enabled connects to the external key management server (KMS) that is configured for the StorageGRID site. Each KMS (or KMS cluster) manages the encryption keys for all appliance nodes at the site. These keys encrypt and decrypt the data on each disk in an appliance that has node encryption enabled.

A KMS can be set up in Grid Manager before or after the appliance is installed in StorageGRID. See the information about KMS and appliance configuration in the instructions for administering StorageGRID for additional details.

- If a KMS is set up before installing the appliance, KMS-controlled encryption begins when you enable node encryption on the appliance and add it to a StorageGRID site where KMS is configured.
- If a KMS is not set up before you install the appliance, KMS-controlled encryption is performed on each appliance that has node encryption enabled as soon as a KMS is configured and available for the site that contains the appliance node.



Any data that exists before an appliance that has node encryption enabled connects to the configured KMS is encrypted with a temporary key that is not secure. The appliance is not protected from removal or theft until the key is set to a value provided by the KMS.

Without the KMS key needed to decrypt the disk, data on the appliance cannot be retrieved and is effectively lost. This is the case whenever the decryption key cannot be retrieved from the KMS. The key becomes inaccessible if a customer clears the KMS configuration, a KMS key expires, connection to the KMS is lost, or the appliance is removed from the StorageGRID system where its KMS keys are installed.

### Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

**https://Controller\_IP:8443**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.



After the appliance has been encrypted with a KMS key, the appliance disks cannot be decrypted without using the same KMS key.

2. Select **Configure Hardware > Node Encryption**.

NetApp® StorageGRID® Appliance Installer Help ▾

Home Configure Networking ▾ Configure Hardware ▾ Monitor Installation Advanced ▾

### Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

#### Encryption Status

⚠ You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

**Save**

#### Key Management Server Details

3. Select **Enable node encryption**.

You can unselect **Enable node encryption** without risk of data loss until you select **Save** and the appliance node accesses the KMS encryption keys in your StorageGRID system and begins disk encryption. You are not able to disable node encryption after the appliance is installed.



After you add an appliance that has node encryption enabled to a StorageGRID site that has a KMS, you cannot stop using KMS encryption for the node.

4. Select **Save**.
5. Deploy the appliance as a node in your StorageGRID system.

KMS-controlled encryption begins when the appliance accesses the KMS keys configured for your StorageGRID site. The installer displays progress messages during the KMS encryption process, which might take a few minutes depending on the number of disk volumes in the appliance.



Appliances are initially configured with a random non-KMS encryption key assigned to each disk volume. The disks are encrypted using this temporary encryption key, that is not secure, until the appliance that has node encryption enabled accesses the KMS keys configured for your StorageGRID site.

### After you finish

You can view node-encryption status, KMS details, and the certificates in use when the appliance node is in

maintenance mode.

## Related information

[Administer StorageGRID](#)

[Monitoring node encryption in maintenance mode](#)

### Optional: Changing to RAID6 mode (SG5660 only)

If you have an SG5660 with 60 drives, you can change the volume configuration from its default and recommended setting, Dynamic Disk Pools (DDP), to RAID6. You can only change the mode before deploying the StorageGRID appliance Storage Node.

#### What you'll need

- You have an SG5660. The SG5612 does not support RAID6. If you have an SG5612, you must use DDP mode.



If any volumes have already been configured or if StorageGRID was previously installed, changing the RAID mode causes the volumes to be removed and replaced. Any data on those volumes will be lost.

#### About this task

Before deploying a StorageGRID appliance Storage Node, you can choose from two volume configuration options:

- **Dynamic Disk Pools (DDP)** — This is the default and recommended setting. DDP is an enhanced hardware data protection scheme that delivers better system performance, reduced rebuild times after drive failures, and ease of management.
- **RAID6** — This is a hardware protection scheme that uses parity stripes on each disk, and allows for two disk failures within the RAID set before any data is lost.



Using RAID6 is not recommended for most StorageGRID environments. Although RAID6 can increase storage efficiency to 88% (compared to 80% for DDP), DDP mode provides more efficient recovery from drive failures.

#### Steps

1. Using the service laptop, open a web browser and access the StorageGRID Appliance Installer:  
**`https://E5600SG_Controller_IP:8443`**

Where *E5600SG\_Controller\_IP* is any of the IP addresses for the E5600SG controller.

2. From the menu bar, select **Advanced > RAID Mode**.
3. On the **Configure RAID Mode** page, select **RAID6** from the Mode drop-down list.
4. Click **Save**.

### Optional: Remapping network ports for the appliance

You might need to remap the internal ports on the appliance Storage Node to different external ports. For example, you might need to remap ports because of a firewall issue.

## What you'll need

- You have previously accessed the StorageGRID Appliance Installer.
- You have not configured and do not plan to configure load balancer endpoints.



If you remap any ports, you cannot use the same ports to configure load balancer endpoints. If you want to configure load balancer endpoints and have already remapped ports, follow the steps in the recovery and maintenance instructions for removing port remaps.

## Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Remap Ports**.

The Remap Port page appears.

2. From the **Network** drop-down box, select the network for the port you want to remap: Grid, Admin, or Client.
3. From the **Protocol** drop-down box, select the IP protocol: TCP or UDP.
4. From the **Remap Direction** drop-down box, select which traffic direction you want to remap for this port: Inbound, Outbound, or Bi-directional.
5. For **Original Port**, enter the number of the port you want to remap.
6. For **Mapped-To Port**, enter the number of the port you want to use instead.
7. Click **Add Rule**.

The new port mapping is added to the table, and the remapping takes effect immediately.

## Remap Ports

If required, you can remap the internal ports on the appliance Storage Node to different external ports. For example, you might need to remap ports because of a firewall issue.

Network	Protocol	Remap Direction	Original Port	Mapped-To Port
<input type="radio"/> Grid	TCP	Bi-directional	1800	1801

8. To remove a port mapping, select the radio button for the rule you want to remove, and click **Remove Selected Rule**.

## Related information

[Maintain & recover](#)

## Deploying an appliance Storage Node

After installing and configuring the storage appliance, you can deploy it as a Storage Node in a StorageGRID system. When you deploy an appliance as a Storage Node, you

use the StorageGRID Appliance Installer included on the appliance.

### What you'll need

- If you are cloning an appliance node, continue following the process in recovery and maintenance.

#### Maintain & recover

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.
- 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.
- The primary Admin Node for the StorageGRID system has been deployed.
- 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 a service laptop with a supported web browser.

### About this task

Each storage appliance functions as a single Storage Node. Any appliance can connect to the Grid Network, the Admin Network, and the Client Network

To deploy an appliance Storage Node in a StorageGRID system, you access the StorageGRID Appliance Installer and perform the following steps:

- You specify or confirm the IP address of the primary Admin Node and the name of the Storage Node.
- You start the deployment and wait as volumes are configured and the software is installed.
- When the installation pauses partway through the appliance installation tasks, you resume the installation by signing into the Grid Manager, approving all grid nodes, and completing the StorageGRID installation and deployment processes.



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.

- If you are performing an expansion or recovery operation, follow the appropriate instructions:
  - To add an appliance Storage Node to an existing StorageGRID system, see the instructions for expanding a StorageGRID system.
  - To deploy an appliance Storage Node as part of a recovery operation, see instructions for recovery and maintenance.

### 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.

## 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

Connection state

Connection to 172.16.4.210 ready



### Node name

Node name




### Installation

Current state

Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

- 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.

- If this IP address is not shown or you need to change it, specify the address:

Option	Description
Manual IP entry	<ol style="list-style-type: none"> <li>Unselect the <b>Enable Admin Node discovery</b> check box.</li> <li>Enter the IP address manually.</li> <li>Click <b>Save</b>.</li> <li>Wait for the connection state for the new IP address to become ready.</li> </ol>
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"> <li>Select the <b>Enable Admin Node discovery</b> check box.</li> <li>Wait for the list of discovered IP addresses to be displayed.</li> <li>Select the primary Admin Node for the grid where this appliance Storage Node will be deployed.</li> <li>Click <b>Save</b>.</li> <li>Wait for the connection state for the new IP address to become ready.</li> </ol>

- In the **Node name** field, enter the name you want to use for this appliance node, and click **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.

- 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 installation and maintenance instructions for your appliance.



If you are deploying the Storage Node appliance as a node cloning target, stop the deployment process here and continue the node cloning procedure in recovery and maintenance.

#### Maintain & recover

- 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**.

- If your grid includes multiple appliance Storage Nodes, repeat these 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. This script applies only to Storage Nodes.

## Related information

[Expand your grid](#)

[Maintain & recover](#)

## Monitoring the storage 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**.

The Monitor Installation page shows the installation progress.

Monitor Installation

1. Configure storage			Running
Step	Progress	Status	
Connect to storage controller	<div style="width: 100%; height: 10px; background-color: green;"></div>	Complete	
Clear existing configuration	<div style="width: 100%; height: 10px; background-color: green;"></div>	Complete	
Configure volumes	<div style="width: 30%; height: 10px; background-color: blue;"></div>	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 are not re-run. If you are re-running an installation, any tasks that do not need to be re-run are shown with a green status bar and a status of "Skipped."

2. Review the progress of the first two installation stages.

### 1. Configure storage

During this stage, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity software 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. Go to the next step.



### 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...
```

4. Go to the Grid Manager, approve the pending storage node, and complete the StorageGRID installation process.

When you click **Install** from the Grid Manager, Stage 3 completes and stage 4, **Finalize Installation**, begins. When stage 4 completes, the controller is rebooted.

## Automating appliance installation and configuration

You can automate the installation and configuration of your appliances and configuration of the whole StorageGRID system.

## About this task

Automating installation and configuration can be useful for deploying multiple StorageGRID instances or one large, complex StorageGRID instance.

To automate installation and configuration, use one or more of the following options:

- Create a JSON file that specifies the configuration settings for your appliances. Upload the JSON file using the StorageGRID Appliance Installer.



You can use the same file to configure more than one appliance.

- Use the `StorageGRIDconfigure-sga.py` Python script to automate the configuration of your appliances.
- Use additional Python scripts to configure other components of the whole StorageGRID system (the "grid").



You can use StorageGRID automation Python scripts directly, or you can use them as examples of how to use the StorageGRID Installation REST API in grid deployment and configuration tools you develop yourself. See the information about downloading and extracting the StorageGRID installation files in the Recovery and Maintenance instructions.

## Automating appliance configuration using the StorageGRID Appliance Installer

You can automate the configuration of an appliance by using a JSON file that contains the configuration information. You upload the file using the StorageGRID Appliance Installer.

### What you'll need

- Your appliance must be on the latest firmware compatible with StorageGRID 11.5 or higher.
- You must be connected to the StorageGRID Appliance Installer on the appliance you are configuring using a supported browser.

## About this task

You can automate appliance configuration tasks such as configuring the following:

- Grid Network, Admin Network, and Client Network IP addresses
- BMC interface
- Network links
  - Port bond mode
  - Network bond mode
  - Link speed

Configuring your appliance using an uploaded JSON file is often more efficient than performing the configuration manually using multiple pages in the StorageGRID Appliance Installer, especially if you have to configure many nodes. You must apply the configuration file for each node one at a time.



Experienced users who want to automate both the installation and configuration of their appliances can use the `configure-sga.py` script.

[Automating installation and configuration of appliance nodes using the `configure-sga.py` script](#)

## Steps

1. Generate the JSON file using one of the following methods:

- The ConfigBuilder application

[ConfigBuilder.netapp.com](https://configbuilder.netapp.com)

- The `configure-sga.py` appliance configuration script. You can download the script from StorageGRID Appliance Installer (**Help > Appliance Configuration Script**). See the instructions on automating the configuration using the `configure-sga.py` script.

[Automating installation and configuration of appliance nodes using the configure-sga.py script](#)

The node names in the JSON file must follow these requirements:

- Must be a valid hostname containing at least 1 and no more than 32 characters
- Can use letters, numbers, and hyphens are allowed
- Cannot start or end with a hyphen or contain only numbers



Ensure that the node names (the top-level names) in the JSON file are unique, or you will not be able to configure more than one node using the JSON file.

2. Select **Advanced > Update Appliance Configuration**.

The Update Appliance Configuration page appears.

### Update Appliance Configuration

Use a JSON file to update this appliance's configuration. You can generate the JSON file from the [ConfigBuilder](#) application or from the [appliance configuration script](#).

**⚠** You might lose your connection if the applied configuration from the JSON file includes "link\_config" and/or "networks" sections. If you are not reconnected within 1 minute, re-enter the URL using one of the other IP addresses assigned to the appliance.

### Upload JSON

JSON configuration	<input type="button" value="Browse"/>
Node name	<input type="button" value="-- Upload a file"/>
<input type="button" value="Apply JSON configuration"/>	

3. Select the JSON file with the configuration you want to upload.

- a. Select **Browse**.
- b. Locate and select the file.
- c. Select **Open**.

The file is uploaded and validated. When the validation process is complete, the file name is shown next to a green check mark.



You might lose connection to the appliance if the configuration from the JSON file includes sections for "link\_config", "networks", or both. If you are not reconnected within 1 minute, re-enter the appliance URL using one of the other IP addresses assigned to the appliance.

#### Upload JSON

JSON configuration

Node name

The **Node name** drop down is populated with the top-level node names defined in the JSON file.



If the file is not valid, the file name is shown in red and an error message is displayed in a yellow banner. The invalid file is not applied to the appliance. You can use ConfigBuilder to ensure you have a valid JSON file.

4. Select a node from the list in the **Node name** drop down.

The **Apply JSON configuration** button is enabled.

#### Upload JSON

JSON configuration

Node name

5. Select **Apply JSON configuration**.

The configuration is applied to the selected node.

### Automating installation and configuration of appliance nodes using the `configure-sga.py` script

You can use the `configure-sga.py` script to automate many of the installation and configuration tasks for StorageGRID appliance nodes, including installing and configuring a primary Admin Node. This script can be useful if you have a large number of appliances to configure. You can also use the script to generate a JSON file that contains appliance

configuration information.

### What you'll need

- 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 primary Admin Node using the StorageGRID Appliance Installer.
- If you are installing the primary Admin Node, you know its IP address.
- If you are installing and configuring other nodes, the primary Admin Node has been deployed, and you know its IP address.
- For all nodes other than the primary Admin Node, 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 downloaded the `configure-sga.py` file. The file is included in the installation archive, or you can access it by clicking **Help > Appliance Installation Script** in the StorageGRID Appliance Installer.



This procedure is for advanced users with some experience using command-line interfaces. Alternatively, you can also use the StorageGRID Appliance Installer to automate the configuration.

[Automating appliance configuration using the StorageGRID Appliance Installer](#)

### Steps

1. Log in to the Linux machine you are using to run the Python script.
2. For general help with the script syntax and to see a list of the available parameters, enter the following:

```
configure-sga.py --help
```

The `configure-sga.py` script uses five subcommands:

- `advanced` for advanced StorageGRID appliance interactions, including BMC configuration and creating a JSON file containing the current configuration of the appliance
- `configure` for configuring the RAID mode, node name, and networking parameters
- `install` for starting a StorageGRID installation
- `monitor` for monitoring a StorageGRID installation
- `reboot` for rebooting the appliance

If you enter a subcommand (`advanced`, `configure`, `install`, `monitor`, or `reboot`) argument followed by the `--help` option you will get a different help text providing more detail on the options available within that subcommand:

```
configure-sga.py subcommand --help
```

3. To confirm the current configuration of the appliance node, enter the following where `SGA-install-ip` is any one of the IP addresses for the appliance node:

```
configure-sga.py configure SGA-INSTALL-IP
```

The results show current IP information for the appliance, including the IP address of the primary Admin Node and information about the Admin, Grid, and Client Networks.

```
Connecting to +https://10.224.2.30:8443+ (Checking version and
connectivity.)
2021/02/25 16:25:11: Performing GET on /api/versions... Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/system-info... Received
200
2021/02/25 16:25:11: Performing GET on /api/v2/admin-connection...
Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/link-config... Received
200
2021/02/25 16:25:11: Performing GET on /api/v2/networks... Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/system-config... Received
200
```

#### StorageGRID Appliance

```
Name:          LAB-SGA-2-30
Node type:     storage
```

#### StorageGRID primary Admin Node

```
IP:           172.16.1.170
State:        unknown
Message:      Initializing...
Version:      Unknown
```

#### Network Link Configuration

##### Link Status

Link	State	Speed (Gbps)
----	-----	-----
1	Up	10
2	Up	10
3	Up	10
4	Up	10
5	Up	1
6	Down	N/A

##### Link Settings

```
Port bond mode:    FIXED
Link speed:        10GBE

Grid Network:      ENABLED
  Bonding mode:    active-backup
  VLAN:            novlan
  MAC Addresses:   00:a0:98:59:8e:8a  00:a0:98:59:8e:82

Admin Network:     ENABLED
  Bonding mode:    no-bond
  MAC Addresses:   00:80:e5:29:70:f4
```

```
Client Network:      ENABLED
  Bonding mode:      active-backup
  VLAN:              novlan
  MAC Addresses:     00:a0:98:59:8e:89  00:a0:98:59:8e:81
```

#### Grid Network

```
CIDR:      172.16.2.30/21 (Static)
MAC:       00:A0:98:59:8E:8A
Gateway:   172.16.0.1
Subnets:  172.17.0.0/21
           172.18.0.0/21
           192.168.0.0/21
MTU:      1500
```

#### Admin Network

```
CIDR:      10.224.2.30/21 (Static)
MAC:       00:80:E5:29:70:F4
Gateway:   10.224.0.1
Subnets:  10.0.0.0/8
           172.19.0.0/16
           172.21.0.0/16
MTU:      1500
```

#### Client Network

```
CIDR:      47.47.2.30/21 (Static)
MAC:       00:A0:98:59:8E:89
Gateway:   47.47.0.1
MTU:      2000
```

```
#####
##### If you are satisfied with this configuration, #####
##### execute the script with the "install" sub-command. #####
#####
```

4. If you need to change any of the values in the current configuration, use the `configure` subcommand to update them. For example, if you want to change the IP address that the appliance uses for connection to the primary Admin Node to `172.16.2.99`, enter the following:  
`configure-sga.py configure --admin-ip 172.16.2.99 SGA-INSTALL-IP`
5. If you want to back up the appliance configuration to a JSON file, use the `advanced` and `backup-file` subcommands. For example, if you want to back up the configuration of an appliance with IP address `SGA-INSTALL-IP` to a file named `appliance-SG1000.json`, enter the following:  
`configure-sga.py advanced --backup-file appliance-SG1000.json SGA-INSTALL-IP`

The JSON file containing the configuration information is written to the same directory you executed the script from.



Check that the top-level node name in the generated JSON file matches the appliance name. Do not make any changes to this file unless you are an experienced user and have a thorough understanding of StorageGRID APIs.

- When you are satisfied with the appliance configuration, use the `install` and `monitor` subcommands to install the appliance:

```
configure-sga.py install --monitor SGA-INSTALL-IP
```

- If you want to reboot the appliance, enter the following:

```
configure-sga.py reboot SGA-INSTALL-IP
```

## Automating the configuration of StorageGRID

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

### What you'll need

- 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>	Sample 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 sample 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

- 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`.

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

For example:

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

## After you finish

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, you need to extract 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.
```

## Overview of installation REST APIs

StorageGRID provides two REST APIs for performing installation tasks: the StorageGRID Installation API and the StorageGRID Appliance Installer API.

Both APIs use 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 (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs 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 in the event that 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.

### StorageGRID Appliance Installer API

The StorageGRID Appliance Installer API can be accessed over HTTPS from `Controller_IP:8443`.

To access the API documentation, go to the StorageGRID Appliance Installer on the appliance and select **Help > API Docs** from the menu bar.

The StorageGRID Appliance Installer API includes the following sections:

- **clone** — Operations to configure and control node cloning.
- **encryption** — Operations to manage encryption and view encryption status.
- **hardware configuration** — Operations to configure system settings on attached hardware.
- **installation** — Operations for starting the appliance installation and for monitoring installation status.
- **networking** — Operations related to the Grid, Admin, and Client Network configuration for a StorageGRID appliance and appliance port settings.
- **setup** — Operations to help with initial appliance installation setup including requests to get information about the system and update the primary Admin Node IP.
- **support** — Operations for rebooting the controller and getting logs.
- **upgrade** — Operations related to upgrading appliance firmware.
- **uploadsg** — Operations for uploading StorageGRID installation files.

## Troubleshooting the hardware installation

If you encounter issues during the installation, you might find it helpful to review troubleshooting information related to hardware setup and connectivity issues.

### Related information

[Hardware setup appears to hang](#)

## Hardware setup appears to hang

The StorageGRID Appliance Installer might not be available if hardware faults or cabling errors prevent the E5600SG controller from completing its boot-up processing.

### Steps

1. Check the Needs Attention LED on either controller and look for a flashing error code.

During power up, the Service Action Allowed and Service Action Required LEDs are turned on while the hardware is initializing. The upper decimal point of the lower digit, called the *diagnostic LED*, also illuminates. The seven-segment display runs through a sequence of codes that are common for both controllers. This is normal and is not an indication of an error. When the hardware boots successfully, the Service Action LEDs are turned off, and the displays are driven by the firmware.

2. Review the codes on the seven-segment display for the E5600SG controller.



The installation and provisioning take time. Some installation phases do not report updates to the StorageGRID Appliance Installer for several minutes.

If an error occurs, the seven-segment display flashes a sequence, such as HE.

3. To understand what these codes mean, see the following resources:

Controller	Reference
E5600SG controller	<ul style="list-style-type: none"><li>• “HE error: Error synchronizing with SANtricity OS Software”</li><li>• “E5600SG controller seven-segment display codes”</li></ul>
E2700 controller	E-Series documentation  <b>Note:</b> The codes described for the E-Series E5600 controller do not apply to the E5600SG controller in the appliance.

4. If this does not resolve the issue, contact technical support.

### Related information

[E5600SG controller seven-segment display codes](#)

[HE error: Error synchronizing with SANtricity OS Software](#)

[E2700 Controller-Drive Tray and Related Drive Trays Installation Guide](#)

[NetApp Documentation: E2700 Series](#)

**HE error: Error synchronizing with SANtricity OS Software**

The seven-segment display on the compute controller shows an HE error code if the

## StorageGRID Appliance Installer cannot synchronize with SANtricity OS Software.

### About this task

If an HE error code is displayed, perform this corrective action.

### Steps

1. Check the integrity of the two SAS interconnect cables, and confirm they are securely connected.
2. As required, replace one or both of the cables, and try again.
3. If this does not resolve the issue, contact technical support.

### Troubleshooting connection issues

If you encounter connection issues during the StorageGRID appliance installation, you should perform the corrective action steps listed.

#### Unable to connect to StorageGRID appliance over the network

If you cannot connect to the appliance, there might be a network issue, or the hardware installation might not have been completed successfully.

- **Issue**

You cannot connect to the appliance.

- **Cause**

This could occur if there is a network issue or the hardware installation did not complete successfully.

- **Corrective action**

- a. Ping the appliance:

```
ping E5600_controller_IP
```

- b. Access the StorageGRID Appliance Installer by opening a browser and entering the following:

```
https://Management_Port_IP:8443
```

For Management\_Port\_IP, enter the IP address for management port 1 on the E5600SG controller (provisioned during the physical installation).

- c. Click **Configure Admin network**, and check the IP.
- d. If you receive a response from the ping, check that port 8443 is open in the firewalls.
- e. Reboot the appliance.
- f. Refresh the installation web page.
- g. If this does not resolve the connection issue, contact technical support from the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com).

### Related information

[E5600SG controller seven-segment display codes](#)

## Rebooting the controller while the StorageGRID Appliance Installer is running

You might need to reboot the compute controller while the StorageGRID Appliance Installer is running. For example, you might need to reboot the controller if the installation fails.

### About this task

This procedure only applies when the compute controller is running the StorageGRID Appliance Installer. Once the installation is completed, this step no longer works because the StorageGRID Appliance Installer is no longer available.

### Steps

1. From the StorageGRID Appliance Installer, click **Advanced** > **Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
  - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



The SG6000-CN controller is rebooted.

## Maintaining the SG5600 appliance

You might need to upgrade the SANtricity OS Software on the E2700 controller, replace the E2700 controller or the E5600SG controller, or replace specific components. The procedures in this section assume that the appliance has already been deployed as a Storage Node in a StorageGRID system.

### Steps

- [Placing an appliance into maintenance mode](#)
- [Upgrading SANtricity OS on the storage controllers using the Grid Manager](#)
- [Upgrading SANtricity OS on the E2700 controller using maintenance mode](#)
- [Upgrading drive firmware using SANtricity Storage Manager](#)

- [Replacing the E2700 controller](#)
- [Replacing the E5600SG controller](#)
- [Replacing other hardware components](#)
- [Changing the link configuration of the E5600SG controller](#)
- [Changing the MTU setting](#)
- [Checking the DNS server configuration](#)
- [Monitoring node encryption in maintenance mode](#)

## Placing an appliance into maintenance mode

You must place the appliance into maintenance mode before performing specific maintenance procedures.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission. For details, see the instructions for administering StorageGRID.

### About this task

Placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.



The password and host key for a StorageGRID appliance in maintenance mode remain the same as they were when the appliance was in service.

### Steps

1. From the Grid Manager, select **Nodes**.
2. From the tree view of the Nodes page, select the appliance Storage Node.
3. Select **Tasks**.

The screenshot shows a navigation bar with the following items: Overview, Hardware, Network, Storage, Objects, ILM, Events, and Tasks. The 'Tasks' item is highlighted with a white border. Below the navigation bar, there are two sections:

- Reboot**: Shuts down and restarts the node. A blue button labeled 'Reboot' is visible.
- Maintenance Mode**: Places the appliance's compute controller into maintenance mode. A blue button labeled 'Maintenance Mode' is visible.

4. Select **Maintenance Mode**.

A confirmation dialog box appears.

## ⚠ Enter Maintenance Mode on SGA-106-15

You must place the appliance's compute controller into maintenance mode to perform certain maintenance procedures on the appliance.

Attention: All StorageGRID services on this node will be shut down. Wait a few minutes for the node to reboot into maintenance mode.

If you are ready to start, enter the provisioning passphrase and click OK.

Provisioning Passphrase

Cancel

OK

5. Enter the provisioning passphrase, and select **OK**.

A progress bar and a series of messages, including "Request Sent," "Stopping StorageGRID," and "Rebooting," indicate that the appliance is completing the steps for entering maintenance mode.

Overview

Hardware

Network

Storage

Objects

ILM

Events

Tasks

### Reboot

Shuts down and restarts the node.

Reboot

### Maintenance Mode

**Attention:** Your request has been sent, but the appliance might take 10-15 minutes to enter maintenance mode. Do not perform maintenance procedures until this tab indicates maintenance mode is ready, or data could become corrupted.



Request Sent

When the appliance is in maintenance mode, a confirmation message lists the URLs you can use to access the StorageGRID Appliance Installer.

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

This node is currently in maintenance mode. Navigate to one of the URLs listed below and perform any necessary maintenance procedures.

- <https://172.16.2.106:8443>
- <https://10.224.2.106:8443>
- <https://47.47.2.106:8443>
- <https://169.254.0.1:8443>

When you are done with any required maintenance procedures, you must exit maintenance mode by clicking Reboot Controller from the StorageGRID Appliance Installer.

6. To access the StorageGRID Appliance Installer, browse to any of the URLs displayed.

If possible, use the URL containing the IP address of the appliance's Admin Network port.




Accessing <https://169.254.0.1:8443> requires a direct connection to the local management port.

7. From the StorageGRID Appliance Installer, confirm that the appliance is in maintenance mode.


This node is in maintenance mode. Perform any required maintenance procedures. If you want to exit maintenance mode manually to resume normal operation, go to **Advanced > Reboot Controller** to **reboot** the controller.

8. Perform any required maintenance tasks.

9. After completing maintenance tasks, exit maintenance mode and resume normal node operation. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status  for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Upgrading SANtricity OS on the storage controllers using the Grid Manager

Use the Grid Manager to apply a SANtricity OS upgrade.

### What you'll need

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- You must have the Maintenance permission.
- You must be signed in to the Grid Manager using a supported browser.
- You must have the provisioning passphrase.
- You must have access to the NetApp downloads page for SANtricity OS.

### About this task

You cannot perform other software updates (StorageGRID software upgrade or a hotfix) until you have completed the SANtricity OS upgrade process. If you attempt to start a hotfix or a StorageGRID software upgrade before the SANtricity OS upgrade process has finished, you are redirected to the SANtricity OS upgrade page.

The procedure will not be complete until the SANtricity OS upgrade has been successfully applied to all applicable nodes. It might take more than 30 minutes to load the SANtricity OS on each node and up to 90 minutes to reboot each StorageGRID storage appliance.



The following steps are only applicable when you are using the Grid Manager to perform the upgrade.



This procedure will automatically upgrade the NVSRAM to the most recent version associated with the SANtricity OS upgrade. You do not need to apply a separate NVSRAM upgrade file.

## Steps

1. From a service laptop, download the new SANtricity OS file from the NetApp support site.

Be sure to choose the SANtricity OS version for the E2700 storage controller.

2. Sign in to the Grid Manager using a supported browser.
3. Select **Maintenance**. Then, in the System section of the menu, select **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.

- To perform a major version upgrade of StorageGRID, see the [instructions for upgrading StorageGRID](#), and then select **StorageGRID Upgrade**.
- To apply a hotfix to all nodes in your system, see "Hotfix procedure" in the [recovery and maintenance instructions](#), and then select **StorageGRID Hotfix**.
- To upgrade SANtricity OS software on a storage controller, see "Upgrading SANtricity OS Software on the storage controllers" in the installation and maintenance instructions for your storage appliance, and then select **SANtricity OS**.

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)



4. Click **SANtricity OS**.

The SANtricity OS page appears.

## SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

### SANtricity OS Upgrade File

---

SANtricity OS Upgrade File



Browse

### Passphrase

---

Provisioning Passphrase



Start

5. Select the SANtricity OS upgrade file you downloaded from the NetApp support site.
  - a. Click **Browse**.
  - b. Locate and select the file.
  - c. Click **Open**.

The file is uploaded and validated. When the validation process is done, the file name is shown in the Details field.



Do not change the file name since it is part of the verification process.

## SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

### SANtricity OS Upgrade File

SANtricity OS Upgrade File



Browse



RC\_XXXXXXXXXX\_XXXXX.dlp

Details



RC\_XXXXXXXXXX\_XXXXX.dlp

### Passphrase

Provisioning Passphrase



Start

6. Enter the provisioning passphrase.

The **Start** button is enabled.

## SANtricity OS

You can use this page to upgrade the SANtricity OS software on storage controllers in a storage appliance. Before installing the new software, confirm the storage controllers are Nominal (**Nodes > appliance node > Hardware**) and ready for an upgrade. A health check is automatically performed as part of the upgrade process and valid NVSRAM is automatically installed based on the appliance type and new software version. The software upgrade can take up to 30 minutes per appliance. When the upgrade is complete, the node will be automatically rebooted to activate the SANtricity OS on the storage controllers. If you have multiple types of appliances, repeat this procedure to install the appropriate OS software for each type.

### SANtricity OS Upgrade File

SANtricity OS Upgrade File



Browse



RC\_XXXXXXXXXX\_XXXXX.dlp

Details



RC\_XXXXXXXXXX\_XXXXX.dlp

### Passphrase

Provisioning Passphrase



Start

7. Click **Start**.

A warning box appears stating that your browser's connection might be lost temporarily as services on nodes that are upgraded are restarted.

**Warning**

**Nodes can disconnect and services might be affected**

The node will be automatically rebooted at the end of upgrade and services will be affected. Are you sure you want to start the SANtricity OS upgrade?

8. Click **OK** to stage the SANtricity OS upgrade file to the primary Admin Node.

When the SANtricity OS upgrade starts:

a. The health check is run. This process checks that no nodes have the status of Needs Attention.

 If any errors are reported, resolve them and click **Start** again.

b. The SANtricity OS Upgrade Progress table appears. This table shows all Storage Nodes in your grid and the current stage of the upgrade for each node.

 The table shows all Storage Nodes, including software-based Storage Nodes. You must approve the upgrade for all Storage Nodes, even though a SANtricity OS upgrade has no effect on software-based Storage Nodes. The upgrade message returned for software-based Storage Nodes is "SANtricity OS upgrade is not applicable to this node."

#### SANtricity OS Upgrade Progress

 **Storage Nodes** - 0 out of 4 completed

Search  

Site	Name	Progress	Stage	Details	Action
RTP Lab 1	DT-10-224-1-181-S1		Waiting for you to approve		<input type="button" value="Approve"/>
RTP Lab 1	DT-10-224-1-182-S2		Waiting for you to approve		<input type="button" value="Approve"/>
RTP Lab 1	DT-10-224-1-183-S3		Waiting for you to approve		<input type="button" value="Approve"/>
RTP Lab 1	NetApp-SGA-Lab2-002-024		Waiting for you to approve		<input type="button" value="Approve"/>

9. Optionally, sort the list of nodes 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.

You can scroll through the list of nodes by using the left and right arrows at the bottom right corner of the section.

10. Approve the grid nodes you are ready to add to the upgrade queue. Approved nodes of the same type are upgraded one at a time.



Do not approve the SANtricity OS upgrade for an appliance storage node unless you are sure the node is ready to be stopped and rebooted. When the SANtricity OS upgrade is approved on a node, the services on that node are stopped. Later, when the node is upgraded, the appliance node is rebooted. These operations might cause service interruptions for clients that are communicating with the node.

- Click either of the **Approve All** buttons to add all Storage Nodes to the SANtricity OS upgrade queue.



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).

- Click one or more **Approve** buttons to add one or more nodes to the SANtricity OS upgrade queue.



You can delay applying a SANtricity OS upgrade to a node, but the SANtricity OS upgrade process will not be complete until you approve the SANtricity OS upgrade on all the listed Storage Nodes.

After you click **Approve**, the upgrade process determines if the node can be upgraded. If a node can be upgraded, it is added to the upgrade queue. +

For some nodes, the selected upgrade file is intentionally not applied and you can complete the upgrade process without upgrading these specific nodes. For nodes intentionally not upgraded, the process will show stage of Complete with one of the following messages in the Details column: +

- Storage Node was already upgraded.
- SANtricity OS upgrade is not applicable to this node.
- SANtricity OS file is not compatible with this node.

The message “SANtricity OS upgrade is not applicable to this node” indicates that the node does not have a storage controller that can be managed by the StorageGRID system. This message will appear for non-appliance Storage Nodes. You can complete the SANtricity OS upgrade process without upgrading the node displaying this message.

The message “SANtricity OS file is not compatible with this node” indicates that the node requires a SANtricity OS file different than the one the process is attempting to install. After you complete the current SANtricity OS upgrade, download the SANtricity OS appropriate for the node and repeat the upgrade process.

11. If you need to remove a node or all nodes from the SANtricity OS upgrade queue, click **Remove** or **Remove All**.

As shown in the example, when the stage progresses beyond Queued, the **Remove** button is hidden and you can no longer remove the node from the SANtricity OS upgrade process.

Storage Nodes - 1 out of 9 completed Approve All Remove All

Search

Site	Name	Progress	Stage	Details	Action
Raleigh	RAL-S1-101-196	<div style="width: 0%;"></div>	Queued		Remove
Raleigh	RAL-S2-101-197	<div style="width: 100%; background-color: green;"></div>	Complete		
Raleigh	RAL-S3-101-198	<div style="width: 0%;"></div>	Queued		Remove
Sunnyvale	SVL-S1-101-199	<div style="width: 0%;"></div>	Queued		Remove
Sunnyvale	SVL-S2-101-93	<div style="width: 0%;"></div>	Waiting for you to approve		Approve
Sunnyvale	SVL-S3-101-94	<div style="width: 0%;"></div>	Waiting for you to approve		Approve
Vancouver	VTC-S1-101-193	<div style="width: 0%;"></div>	Waiting for you to approve		Approve
Vancouver	VTC-S2-101-194	<div style="width: 0%;"></div>	Waiting for you to approve		Approve
Vancouver	VTC-S3-101-195	<div style="width: 0%;"></div>	Waiting for you to approve		Approve

12. Wait while the SANtricity OS upgrade is applied to each approved grid node.



If any node shows a stage of Error while the SANtricity OS upgrade is being applied, the upgrade has failed for that node. The appliance might need to be placed in maintenance mode to recover from the failure. Contact technical support before continuing.

If the firmware on the node is too old to be upgraded with the Grid Manager, the node shows a stage of Error with the details: "You must use maintenance mode to upgrade SANtricity OS on this node. See the installation and maintenance instructions for your appliance. After the upgrade, you can use this utility for future upgrades." To resolve the error, do the following:

- a. Use maintenance mode to upgrade SANtricity OS on the node that shows a stage of Error.
- b. Use the Grid Manager to re-start and complete the SANtricity OS upgrade.

When the SANtricity OS upgrade is complete on all approved nodes, the SANtricity OS Upgrade Progress table closes and a green banner shows the date and time the SANtricity OS upgrade was completed.

SANtricity OS upgrade completed at 2020-04-07 13:26:02 EDT.

**SANtricity OS Upgrade File**

SANtricity OS Upgrade File

**Passphrase**

Provisioning Passphrase

13. Repeat this upgrade procedure for any nodes with a stage of Complete that require a different SANtricity OS upgrade file.



For any nodes with a status of Needs Attention, use maintenance mode to perform the upgrade.

### Related information

[Upgrading SANtricity OS on the E2700 controller using maintenance mode](#)

### Upgrading SANtricity OS on the E2700 controller using maintenance mode

If you are unable to upgrade the SANtricity OS Software using the Grid Manager, use the maintenance mode procedure to apply the upgrade.

#### What you'll need

- You have consulted the NetApp Interoperability Matrix Tool (IMT) to confirm that the SANtricity OS version you are using for the upgrade is compatible with your appliance.
- You must place the E5600SG controller into maintenance mode if you are not using the Grid Manager. Placing the controller into maintenance mode interrupts the connection to the E2700 controller. Before changing the link configuration, you must place the E5600SG controller into maintenance mode. Putting a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

[Placing an appliance into maintenance mode](#)

#### About this task

Do not upgrade the SANtricity OS or NVSRAM in the E-Series controller on more than one StorageGRID appliance at a time.



Upgrading more than one StorageGRID appliance at a time might cause data unavailability, depending on your deployment model and ILM policies.

#### Steps

1. From a service laptop, access SANtricity Storage Manager, and sign in.
2. Download the new SANtricity OS Software file and NVSRAM file to the management client.



The NVSRAM is specific to the StorageGRID appliance. Do not use the standard NVSRAM download.

3. Follow the instructions in the *E2700 and E5600 SANtricity Software and Firmware Upgrade instructions* or the SANtricity Storage Manager online help, and upgrade the E2700 controller's firmware, NVSRAM, or both.



If you need to upgrade the NVSRAM in the E2700 controller, you must confirm that the SANtricity OS file you downloaded was designated as compatible with StorageGRID appliances.



Activate the upgrade files immediately. Do not defer activation.


4. Once the upgrade operation has completed, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this

option if you are done working in maintenance mode and are ready to return the node to normal operation.

- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Upgrading drive firmware using SANtricity Storage Manager

You upgrade your drive firmware to make sure you have all the latest features and bug fixes.

### What you'll need

- The storage appliance has an Optimal status.
- All drives have an Optimal status.
- You have the latest version of SANtricity Storage Manager installed that is compatible with your StorageGRID version.

[Upgrading SANtricity OS on the storage controllers using the Grid Manager](#)

[Upgrading SANtricity OS on the E2700 controller using maintenance mode](#)

- You have placed the StorageGRID appliance in maintenance mode.

### Placing an appliance into maintenance mode



Maintenance mode interrupts the connection to the storage controller, stopping all I/O activity and placing all drives offline.



Do not upgrade the drive firmware on more than one StorageGRID appliance at a time. Doing so might cause data unavailability, depending on your deployment model and ILM policies.

### Steps

1. Open a web browser, and enter the IP address as the URL for SANtricity Storage Manager:  
**https://E2700\_Controller\_IP**
2. Enter the SANtricity Storage Manager administrator username and password, if required.
3. From SANtricity Enterprise Management, select the **Devices** tab.


The SANtricity Array Management window opens.

4. From SANtricity Array Management, double-click the Storage Array with the drives to upgrade.
5. Verify that both the Storage Array and drives have an Optimal status.
6. Verify the drive firmware version currently installed in the storage appliance:

- a. From SANtricity Enterprise Management, select **Upgrade > Drive Firmware**.

The Download Drive Firmware - Add Packages window displays the drive firmware files currently in use.

- b. Note current drive firmware revisions and drive identifiers under Current firmware packages.



In this example:

- The drive firmware revision is **NE00**.
- The drive identifier is **HUH728080AL5204**.

Select **View Associated Drives** to display where these drives are installed in your storage appliance.

7. Download and prepare the available drive firmware upgrade:

- a. Open your web browser, navigate to NetApp Support web site, and log in using your ID and password.

[NetApp Support](#)

- b. On the NetApp Support web site, select the **Downloads** tab, and then select **E-Series Disk Drive Firmware**.

The E-Series Disk Firmware page displays.

- c. Search for each **Drive Identifier** installed in your storage appliance and verify that each drive identifier has the latest firmware revision.
  - If the firmware revision is not a link, this drive identifier has the latest firmware revision.
  - If one or more drive part numbers are listed for a drive identifier, a firmware upgrade is available for these drives. You can select any link to download the firmware file.

NetApp | Support

PRODUCTS ▾ SYSTEMS ▾ DOCS & KNOWLEDGEBASE ▾ COMMUNITY ▾ DOWNLOADS ▾ TOOLS ▾ CASES ▾ PARTS ▾

Downloads > Firmware > E-Series Disk Firmware

## E-Series Disk Firmware

[Download all current E-Series Disk Firmware](#)

Drive Part Number	Descriptions	Drive Identifier	Firmware Rev. (Download)	Notes and Config Info	Release Date
<input type="text" value="Drive Part Number"/>	<input type="text" value="Descriptions"/>	<input type="text" value="HUH728080AL5204"/>	<input type="text" value="Firmware Rev. (Download)"/>		
E-X4073A	HDD, 8TB, SAS, 7.2K, PI	HUH728080AL5204	<a href="#">NE01</a>	NE01 Fixes Bug 1122414	26-Jul-2018
E-X4074A	HDD, 8TB, SAS, 7.2K, PI	HUH728080AL5204	<a href="#">NE01</a>	NE01 Fixes Bug 1122414	26-Jul-2018
E-X4127A	HDD, 8TB, SAS, 7.2K, PI	HUH728080AL5204	<a href="#">NE01</a>	NE01 Fixes Bug 1122414	26-Jul-2018
E-X4128A	HDD, 8TB, SAS, 7.2K, PI	HUH728080AL5204	<a href="#">NE01</a>	NE01 Fixes Bug 1122414	26-Jul-2018

- d. If a later firmware revision is listed, select the link in the Firmware Rev. (Download) column to download a .zip archive containing the firmware file.

- e. Extract (unzip) the drive firmware archive files you downloaded from the Support site.

8. Install the drive firmware upgrade:

- a. From the SANtricity Storage Manager Download Drive Firmware - Add Packages window, select **Add**.
- b. Navigate to the directory that contains the firmware files and select up to four firmware files.

Drive firmware files have a filename similar to

D\_HUC101212CSS600\_30602291\_MS01\_2800\_0002.dlp

Selecting more than one firmware file to upgrade the firmware of the same drive might result in a file conflict error. If a file conflict error occurs, an error dialog appears. To resolve this error, select **OK** and remove all other firmware files except the one that you want to use for upgrading the firmware of the drive. To remove a firmware file, select the firmware file in the Packages to Be Transferred information area, and select **Remove**. In addition, you can only select up to four drive firmware packages at one time.

- c. Select **OK**.

The system updates the Packages to be transferred information area with the firmware files you selected.

d. Select **Next**.

The Download Drive Firmware - Select Drives window opens.

- All drives in the appliance are scanned for configuration information and upgrade eligibility.
- You are presented with a selection (depending on what variety of drives you have in the storage array) of compatible drives that can be upgraded with the firmware you selected. The drives capable of being upgraded as an online operation are displayed by default.
- The selected firmware for the drive appears in the Proposed Firmware information area. If you must change the firmware, select **Back** to return to the previous dialog.

e. From the Drive upgrade capability, select the **Parallel** download operation or **All**.

You can use either of these upgrade methods because the appliance is in maintenance mode, where I/O activity is stopped for all drives and all volumes.

f. In Compatible Drives, select the drives for which you want to upgrade the selected firmware files.

- For one or more drives, select each drive you want to upgrade.
- For all compatible drives, select **Select all**.

The best practice is to upgrade all drives of the same model to the same firmware revision.

g. Select **Finish**; then, type `yes` and select **OK**.

- The drive firmware download and upgrade begins, with Download Drive Firmware - Progress indicating the status of the firmware transfer for all drives.
- The status of each drive participating in the upgrade appears in the Transfer Progress column of Devices updated.

A parallel drive firmware upgrade operation can take as much as 90 seconds to complete if all drives are upgraded on a 24-drive system. On a larger system, the execution time is slightly longer.

h. During the firmware upgrade process, you can: +

- Select **Stop** to stop the firmware upgrade in progress. Any firmware upgrade currently in progress are completed. Any drives that have attempted firmware upgrade show their individual status. Any remaining drives are listed with a status of Not attempted.



Stopping the drive firmware upgrade in process might result in data loss or unavailable drives.

- Select **Save As** to save a text report of the firmware upgrade progress summary. The report saves with a default `.log` file extension. If you want to change the file extension or directory, change the parameters in Save Drive Download Log.

i. Use Download Drive Firmware - Progress to monitor the progress of the drive firmware upgrades. The Drives Updated area contains a list of drives that are scheduled for firmware upgrade and the transfer status of each drive's download and upgrade.

The progress and status of each drive that is participating in the upgrade appears in the Transfer Progress column. Take the appropriate recommended action if any errors occur during the upgrade.

- **Pending**

This status is shown for an online firmware download operation that has been scheduled but has not yet started.

- **In progress**

The firmware is being transferred to the drive.

- **Reconstruction in progress**

This status is shown if a volume transfer takes place during the rapid reconstruction of a drive. This is typically due to a controller reset or failure and the controller owner transfers the volume.

The system will initiate a full reconstruction of the drive.

- **Failed - partial**

The firmware was only partially transferred to the drive before a problem prevented the rest of the file from being transferred.

- **Failed - invalid state**

The firmware is not valid.

- **Failed - other**

The firmware could not be downloaded, possibly because of a physical problem with the drive.

- **Not attempted**

The firmware was not downloaded, which may be due to a number of different reasons such as the download was stopped before it could occur, or the drive did not qualify for the upgrade, or the download could not occur due to an error.

- **Successful**


The firmware was downloaded successfully.

9. After the drive firmware upgrade completes:


- To close the Drive Firmware Download Wizard, select **Close**.
- To start the wizard again, select **Transfer More**.

10. Once the upgrade operation has completed, reboot the appliance. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Replacing the E2700 controller


You might need to replace the E2700 controller if it is not functioning optimally or it has failed.

### What you'll need

- You have a replacement controller with the same part number as the controller you are replacing.
- You have labels to identify each cable that is connected to the controller.
- You have antistatic protection.
- You must have the Maintenance or Root Access permission. For details, see the instructions for administering StorageGRID.

### About this task

You can determine if you have a failed controller by checking the amber Service Action Required LED on the controller (shown as 1 in the illustration). If this LED is on, the controller should be replaced.



The appliance Storage Node will not be accessible when you replace the controller. If the E2700 controller is functioning sufficiently, you can place the E5600SG controller into maintenance mode.

When you replace a controller, you must remove the battery from the original controller and install it in the replacement controller.

### Steps

1. Prepare to remove the controller.

You use SANtricity Storage Manager to perform these steps.

- a. Make a note of which version of SANtricity OS software is currently installed on the controller.
- b. Make a note of which version of NVSRAM is currently installed.
- c. If the Drive Security feature is enabled, be sure a saved key exists and that you know the pass phrase required to install it.



**Possible loss of data access** -- If all drives in the appliance are security enabled, the new controller will not be able to access the appliance until you unlock the secured drives using the Enterprise Management Window in SANtricity Storage Manager.

- d. Back up the configuration database.

If a problem occurs when you remove a controller, you can use the saved file to restore your configuration.

- e. Collect support data for the appliance.



Collecting support data before and after replacing a component ensures you can send a full set of logs to technical support in case the replacement does not resolve the problem.

2. If the StorageGRID appliance is running in a StorageGRID system, place the E5600SG controller into maintenance mode.

#### Placing an appliance into maintenance mode


3. If the E2700 controller is functioning sufficiently to allow for a controlled shutdown, confirm that all operations have completed.

- a. From the title bar of the Array Management Window, select **Monitor > Reports > Operations in**

## Progress.

- b. Confirm that all operations have completed.
4. Follow the instructions in the replacement procedure for a simplex E2700 controller to complete these steps:
  - a. Label the cables and then disconnect the cables.  

i To prevent degraded performance, do not twist, fold, pinch, or step on the cables.
  - b. Remove the failed controller from the appliance.
  - c. Remove the controller cover.
  - d. Unscrew the thumbscrew and remove the battery from the failed controller.
  - e. Install the battery in the replacement controller, and replace the controller cover.
  - f. Install the replacement controller into the appliance.
  - g. Replace the cables.
  - h. Wait for the E2700 controller to reboot. Verify that the seven-segment display shows a state of 99.
5. If the appliance uses secured drives, import the drive security key.
6. Return the appliance to normal operating mode. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.



During the reboot, the following screen appears:

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

This node is rebooting from maintenance mode to rejoin the grid. Monitor the node status to determine when the node has successfully rejoined the grid.

The appliance reboots and rejoins the grid. This process can take up to 20 minutes.

7. Confirm that the reboot is complete and that the node has rejoined the grid. In the Grid Manager, verify that the **Nodes** tab displays a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.

The screenshot shows the NetApp StorageGRID web interface. The top navigation bar includes 'Dashboard', 'Alerts', 'Nodes' (highlighted with a yellow box), 'Tenants', 'ILM', 'Configuration', 'Maintenance', and 'Support'. Below the navigation, the 'StorageGRID Deployment' section is expanded to show a tree view of 'Data Center 1' containing nodes: DC1-ADM1, DC1-ARC1, DC1-G1, DC1-S1, DC1-S2, and DC1-S3. Each node has a green checkmark next to it. To the right, there are tabs for 'Network', 'Storage', 'Objects', 'ILM', and 'Load Balancer'. Below these tabs, there are time range filters: '1 hour', '1 day', '1 week', '1 month', and 'Custom'. A 'Network Traffic' chart is displayed below the filters, showing a peak of 6.0 Mbps.

8. From SANtricity Storage Manager, confirm that the new controller is Optimal, and collect support data.

### Related information

[NetApp E-Series and EF-Series Hardware Replacement Procedures](#)

[NetApp Documentation: E2700 Series](#)

### Replacing the E5600SG controller

You might need to replace the E5600SG controller.

#### What you'll need

You must have access to the following resources:

- E-Series hardware replacement information on the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com)
- E5600 documentation on the Support Site
- The appliance has been placed maintenance mode.

### About this task

If both controllers are functioning sufficiently to allow for a controlled shutdown, you can shut down the E5600SG controller first to interrupt the connectivity to the E2700 controller.



If you are replacing the controller before installing StorageGRID software, you might not be able to access the StorageGRID Appliance Installer immediately after completing this procedure. While you can access the StorageGRID Appliance Installer from other hosts on the same subnet as the appliance, you cannot access it from hosts on other subnets. This condition should resolve itself within 15 minutes (when any ARP cache entries for the original controller time out), or you can clear the condition immediately by purging any old ARP cache entries manually from the local router or gateway.

### Steps

1. Use antistatic protection.
2. Label each cable that is attached to the E5600SG controller, so you can reconnect the cables correctly.



To prevent degraded performance, do not twist, fold, pinch, or step on the cables. Do not bend the cables tighter than a 5-cm (2-in) radius.

3. When the appliance has been placed maintenance mode, shut down the E5600SG controller.
  - 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. Shut down the E5600SG controller:  
**shutdown -h now**
4. Turn off the power to the enclosure, and wait until all LED and seven-segment display activity on the rear of the controller has stopped.
5. Remove the cables.
6. Remove the controller, as described in the E5600SG controller documentation.
7. Insert the new controller, as described in the E5600SG controller documentation.
8. Replace all cables.
9. Turn the power back on to the enclosure.
10. Monitor the seven-segment codes.
  - E2700 controller:  
  
The final LED state is 99.
  - E5600SG controller:

The final LED state is HA.

11. Monitor the status of the appliance Storage Node in the Grid Manager.

Verify that the appliance Storage Nodes returns to the expected status.

### Related information

[NetApp E-Series and EF-Series Hardware Replacement Procedures](#)

[NetApp Documentation: E5600 Series](#)

### Replacing other hardware components

You might need to replace a drive, fan, power supply, or battery in the StorageGRID appliance.

#### What you'll need

- You have the E-Series hardware replacement procedure.
- The appliance has been placed maintenance mode if the component replacement procedure requires that you shut down the appliance.

[Placing an appliance into maintenance mode](#)

#### About this task

To replace a drive, power-fan canister, fan canister, power canister, battery, or drive drawer, refer to the standard procedures for the E2700 and E5600 storage arrays. Focus on the step-by-step instructions for removing and replacing the hardware itself; many of the SANtricity Storage Manager procedures do not apply to an appliance.

#### SG5612 component replacement instructions

FRU	See
Drive	Follow the steps in the E-Series instructions for replacing a drive in the E2600, E2700, E5400, E5500, E5600 or 12-drive or 24-drive trays.
Power-fan canister	Follow the steps in the E-Series instructions for replacing a failed power-fan canister in the E5612 or the E5624 controller-drive tray.
Battery in the E2700 controller (requires removing the controller)	Follow the steps in <a href="#">Replacing the E2700 controller</a> , but install the new battery in the existing controller.

#### SG5660 component replacement instructions

FRU	See
Drive	Follow the steps in the E-Series instructions for replacing a drive in the E2660, E2760, E5460, E5560, or E5660 trays.
Power canister	Follow the steps in the E-Series instructions for replacing a failed power canister in the E5660 controller-drive tray.
Fan canister	Follow the steps in the E-Series instructions for replacing a failed fan canister in the E5660 controller-drive tray.
Battery in the E2700 controller (requires removing the controller)	Follow the steps in <a href="#">Replacing the E2700 controller</a> , but install the new battery in the existing controller.

### Related information

[NetApp E-Series and EF-Series Hardware Replacement Procedures](#)

[NetApp Documentation: E2700 Series](#)

[NetApp Documentation: E5600 Series](#)

### Changing the link configuration of the E5600SG controller

You can change the Ethernet link configuration of the E5600SG controller. You can change the port bond mode, the network bond mode, and the link speed.

#### What you'll need

- You must place the E5600SG controller into maintenance mode. Placing the controller into maintenance mode interrupts the connection to the E2700 controller. Putting a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

[Placing an appliance into maintenance mode](#)


#### About this task

Options for changing the Ethernet link configuration of the E5600SG controller include:

- Changing **Port bond mode** from Fixed to Aggregate, or from Aggregate to Fixed
- Changing **Network bond mode** from Active-Backup to LACP, or from LACP to Active-Backup
- Enabling or disabling VLAN tagging, or changing the value of a VLAN tag
- Changing the link speed from 10-GbE to 25-GbE, or from 25-GbE to 10-GbE

#### Steps

1. Select **Configure Networking > Link Configuration** from the menu.



2. Make the desired changes to the link configuration.

For more information on the options, see “Configuring network links.”

3. When you are satisfied with your selections, click **Save**.



You might lose your connection if you made changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:

**`https://E5600SG_Controller_IP:8443`**

If you made changes to the VLAN settings, the subnet for the appliance might have changed. If you need to change the IP addresses for the appliance, follow the instructions for configuring IP addresses.

### Setting the IP configuration


4. From the StorageGRID Appliance Installer, select **Configure Networking > Ping Test**.

5. Use the Ping Test tool to check connectivity to IP addresses on any networks that may have been affected by the link configuration changes you made in the [Change link configuration](#) step.


In addition to any other tests you choose to perform, confirm that you can ping the grid IP address of the primary Admin Node, and the grid IP address of at least one other Storage Node. If necessary, correct any link configuration issues.

6. Once you are satisfied that your link configuration changes are working, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Configuring network links \(SG5600\)](#)

## Changing the MTU setting

You can change the MTU setting that you assigned when you configured IP addresses for the appliance node.

### What you'll need

The appliance has been placed maintenance mode.

[Placing an appliance into maintenance mode](#)

### Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > IP Configuration**.
2. Make the desired changes to the MTU settings for the Grid Network, Admin Network, and Client Network.

## Grid Network

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

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

 All required Grid Network subnets must also be defined in the Grid Network Subnet List on the Primary Admin Node before starting installation.

Subnets (CIDR)  



MTU  



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.



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


- When you are satisfied with the settings, select **Save**.
- Reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this

option if you are done working in maintenance mode and are ready to return the node to normal operation.

- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Administer StorageGRID](#)

## Checking the DNS server configuration

You can check and temporarily change the domain name system (DNS) servers that are currently in use by this appliance node.

### What you'll need

The appliance has been placed in maintenance mode.

[Placing an appliance into maintenance mode](#)

### About this task

You might need to change the DNS server settings if an encrypted appliance cannot connect to the key management server (KMS) or KMS cluster because the hostname for the KMS was specified as a domain name instead of an IP address. Any changes that you make to the DNS settings for the appliance are temporary and are lost when you exit maintenance mode. To make these changes permanent, specify the DNS servers in Grid Manager (**Maintenance > Network > DNS Servers**).

- Temporary changes to the DNS configuration are necessary only for node-encrypted appliances where the KMS server is defined using a fully qualified domain name, instead of an IP address, for the hostname.
- When a node-encrypted appliance connects to a KMS using a domain name, it must connect to one of the DNS servers defined for the grid. One of these DNS servers then translates the domain name into an IP address.
- If the node cannot reach a DNS server for the grid, or if you changed the grid-wide DNS settings when a node-encrypted appliance node was offline, the node is unable to connect to the KMS. Encrypted data on the appliance cannot be decrypted until the DNS issue is resolved.

To resolve a DNS issue preventing KMS connection, specify the IP address of one or more DNS servers in the StorageGRID Appliance Installer. These temporary DNS settings allow the appliance to connect to the KMS and decrypt data on the node.

For example, if the DNS server for the grid changes while an encrypted node was offline, the node will not be able to reach the KMS when it comes back online, since it is still using the previous DNS values. Entering the new DNS server IP address in the StorageGRID Appliance Installer allows a temporary KMS connection to decrypt the node data.

### Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > DNS Configuration**.
2. Verify that the DNS servers specified are correct.

#### DNS Servers

 Configuration changes made on this page will not be passed to the StorageGRID software after appliance installation.

#### Servers

Server 1	<input type="text" value="10.224.223.135"/>	
Server 2	<input type="text" value="10.224.223.136"/>	 
<input type="button" value="Cancel"/>		<input type="button" value="Save"/>

3. If required, change the DNS servers.




Changes made to the DNS settings are temporary and are lost when you exit maintenance mode.

4. When you are satisfied with the temporary DNS settings, select **Save**.

The node uses the DNS server settings specified on this page to reconnect to the KMS, allowing data on the node to be decrypted.


5. After node data is decrypted, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



When the node reboots and rejoins the grid, it uses the system-wide DNS servers listed in the Grid Manager. After rejoining the grid, the appliance will no longer use the temporary DNS servers specified in the StorageGRID Appliance Installer while the appliance was in maintenance mode.

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



### Monitoring node encryption in maintenance mode

If you enabled node encryption for the appliance during installation, you can monitor the node-encryption status of each appliance node, including the node-encryption state and key management server (KMS) details.

## What you'll need

- Node encryption must have been enabled for the appliance during installation. You cannot enable node encryption after the appliance is installed.
- The appliance has been placed into maintenance mode.

### [Placing an appliance into maintenance mode](#)

## Steps

1. From the StorageGRID Appliance Installer, select **Configure Hardware > Node Encryption**.

### Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

### Encryption Status

 You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

Save

### Key Management Server Details

View the status and configuration details for the KMS that manages the encryption key for this appliance. You must use the Grid Manager to make configuration changes.

KMS display name	thales
External key UID	41b0306abcce451facfe01b1b4870ae1c1ec6bd5e3849d790223766baf35c57
Hostnames	10.96.99.164 10.96.99.165
Port	5696

Server certificate >

Client certificate >

### Clear KMS Key

 Do not clear the KMS key if you need to access or preserve any data on this appliance.

If you want to reinstall this appliance node (for example, in another grid), you must clear the KMS key. When the KMS key is cleared, all data on this appliance is deleted.

Clear KMS Key and Delete Data

The Node Encryption page includes these three sections:

- Encryption Status shows whether node encryption is enabled or disabled for the appliance.
- Key Management Server Details shows information about the KMS being used to encrypt the appliance. You can expand the server and client certificate sections to view certificate details and status.

- To address issues with the certificates themselves, such as renewing expired certificates, see the information about KMS in the instructions for administering StorageGRID.
- If there are unexpected problems connecting to KMS hosts, verify that the domain name system (DNS) servers are correct and that appliance networking is correctly configured.

### Checking the DNS server configuration

- If you are unable to resolve your certificate issues, contact technical support.
- Clear KMS Key disables node encryption for the appliance, removes the association between the appliance and the key management server that was configured for the StorageGRID site, and deletes all data from the appliance. You must clear the KMS key before you can install the appliance into another StorageGRID system.


### Clearing the key management server configuration




Clearing the KMS configuration deletes data from the appliance, rendering it permanently inaccessible. This data is not recoverable.

2. When you are done checking node-encryption status, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Administer StorageGRID](#)

### Clearing the key management server configuration

Clearing the key management server (KMS) configuration disables node encryption on your appliance. After clearing the KMS configuration, the data on your appliance is permanently deleted and is no longer accessible. This data is not recoverable.

### What you'll need

If you need to preserve data on the appliance, you must perform a node decommission procedure before you clear the KMS configuration.



When KMS is cleared, data on the appliance will be permanently deleted and no longer accessible. This data is not recoverable.

Decommission the node to move any data it contains to other nodes in StorageGRID. See the recovery and maintenance instructions for grid node decommissioning.

### About this task

Clearing the appliance KMS configuration disables node encryption, removing the association between the appliance node and the KMS configuration for the StorageGRID site. Data on the appliance is then deleted and the appliance is left in a pre-install state. This process cannot be reversed.

You must clear the KMS configuration:

- Before you can install the appliance into another StorageGRID system, that does not use a KMS or that uses a different KMS.



Do not clear the KMS configuration if you plan to reinstall an appliance node in a StorageGRID system that uses the same KMS key.

- Before you can recover and reinstall a node where the KMS configuration was lost and the KMS key is not recoverable.
- Before returning any appliance that was previously in use at your site.
- After decommissioning a appliance that had node encryption enabled.



Decommission the appliance before clearing KMS to move its data to other nodes in your StorageGRID system. Clearing KMS before decommissioning the appliance will result in data loss and might render the appliance inoperable.

## Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > Node Encryption**.

### Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

### Encryption Status

You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

Save

### Key Management Server Details

View the status and configuration details for the KMS that manages the encryption key for this appliance. You must use the Grid Manager to make configuration changes.

KMS display name	thales
External key UID	41b0306abcce451facfe01b1b4870ae1c1ec6bd5e3849d790223766baf35c57
Hostnames	10.96.99.164 10.96.99.165
Port	5696

Server certificate >

Client certificate >

### Clear KMS Key

Do not clear the KMS key if you need to access or preserve any data on this appliance.


If you want to reinstall this appliance node (for example, in another grid), you must clear the KMS key. When the KMS key is cleared, all data on this appliance is deleted.

Clear KMS Key and Delete Data



If the KMS configuration is cleared, data on the appliance will be permanently deleted. This data is not recoverable.

3. At the bottom of the window, select **Clear KMS Key and Delete Data**.
4. If you are sure that you want to clear the KMS configuration, type **clear** and select **Clear KMS Key and Delete Data**.



The KMS encryption key and all data are deleted from the node, and the appliance reboots. This can take up to 20 minutes.

5. Open a browser, and enter one of the IP addresses for the appliance's compute controller.  
**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

6. Select **Configure Hardware > Node Encryption**.
7. Verify that node encryption is disabled and that the key and certificate information in **Key Management Server Details** and the **Clear KMS Key and Delete Data** control are removed from the window.

Node encryption cannot be reenabled on the appliance until it is reinstalled in a grid.

#### After you finish

After the appliance reboots and you have verified that KMS has been cleared and that the appliance is in a pre-install state, you can physically remove the appliance from your StorageGRID system. See the recovery and maintenance instructions for information about preparing an appliance for reinstallation.

#### Related information

[Administer StorageGRID](#)

## SG100 & SG1000 services appliances

Learn how to install and maintain the StorageGRID SG100 and SG1000 appliances.

- [SG100 and SG1000 appliances overview](#)
- [SG100 and SG1000 applications](#)
- [Installation and deployment overview](#)
- [Preparing for installation](#)
- [Installing the hardware](#)
- [Configuring StorageGRID connections](#)
- [Configuring the BMC interface](#)
- [Optional: Enabling node encryption](#)
- [Deploying a services appliance node](#)
- [Troubleshooting the hardware installation](#)
- [Maintaining the appliance](#)

### SG100 and SG1000 appliances overview

The StorageGRID SG100 services appliance and the SG1000 services appliance can operate as a Gateway Node and as an Admin Node to provide high availability load balancing services in a StorageGRID system. Both appliances can operate as Gateway Nodes and Admin Nodes (primary or non-primary) at the same time.

#### Appliance features

Both models of the services appliance provide the following features:

- Gateway Node or Admin Node functions for a StorageGRID system.
- The StorageGRID Appliance Installer to simplify node deployment and configuration.
- When deployed, can access StorageGRID software from an existing Admin Node or from software downloaded to a local drive. To further simplify the deployment process, a recent version of the software is preloaded onto the appliance during manufacturing.
- A baseboard management controller (BMC) for monitoring and diagnosing some of the appliance hardware.
- The ability to connect to all three StorageGRID networks, including the Grid Network, the Admin Network, and the Client Network:
  - The SG100 supports up to four 10- or 25-GbE connections to the Grid Network and Client Network.
  - The SG1000 supports up to four 10-, 25-, 40-, or 100-GbE connections to the Grid Network and Client Network.

#### SG100 and SG1000 diagrams

This figure shows the front of the SG100 and the SG1000 with the bezel removed.




From the front, the two appliances are identical except for the product name on the bezel.

The two solid-state drives (SSDs), indicated by the orange outline, are used for storing the StorageGRID operating system and are mirrored using RAID1 for redundancy. When the SG100 or SG1000 services appliance is configured as an Admin Node, these drives are used to store audit logs, metrics, and database tables.

The remaining drive slots are blank.

### Connectors on the rear of the SG100

This figure shows the connectors on the back of the SG100.




	Port	Type	Use
1	Network ports 1-4	10/25-GbE, based on cable or SFP transceiver type (SFP28 and SFP+ modules are supported), switch speed, and configured link speed	Connect to the Grid Network and the Client Network for StorageGRID.
2	BMC management port	1-GbE (RJ-45)	Connect to the appliance baseboard management controller.
3	Diagnostic and support ports	<ul style="list-style-type: none"> <li>• VGA</li> <li>• Serial, 115200 8-N-1</li> <li>• USB</li> </ul>	Reserved for technical support use.
4	Admin Network port 1	1-GbE (RJ-45)	Connect the appliance to the Admin Network for StorageGRID.

	Port	Type	Use
5	Admin Network port 2	1-GbE (RJ-45)	<p>Options:</p> <ul style="list-style-type: none"> <li>• Bond with management port 1 for a redundant connection to the Admin Network for StorageGRID.</li> <li>• Leave disconnected and available for temporary local access (IP 169.254.0.1).</li> <li>• During installation, use port 2 for IP configuration if DHCP-assigned IP addresses are not available.</li> </ul>

### Connectors on the rear of the SG1000

This figure shows the connectors on the back of the SG1000.



	Port	Type	Use
1	Network ports 1-4	10/25/40/100-GbE, based on cable or transceiver type, switch speed, and configured link speed. QSFP28 and QSFP+ (40/100GbE) are supported natively and SFP28/SFP+ transceivers can be used with a QSA (sold separately) to use 10/25GbE speeds.	Connect to the Grid Network and the Client Network for StorageGRID.
2	BMC management port	1-GbE (RJ-45)	Connect to the appliance baseboard management controller.
3	Diagnostic and support ports	<ul style="list-style-type: none"> <li>• VGA</li> <li>• Serial, 115200 8-N-1</li> <li>• USB</li> </ul>	Reserved for technical support use.
4	Admin Network port 1	1-GbE (RJ-45)	Connect the appliance to the Admin Network for StorageGRID.

	Port	Type	Use
5	Admin Network port 2	1-GbE (RJ-45)	<p>Options:</p> <ul style="list-style-type: none"> <li>• Bond with management port 1 for a redundant connection to the Admin Network for StorageGRID.</li> <li>• Leave disconnected and available for temporary local access (IP 169.254.0.1).</li> <li>• During installation, use port 2 for IP configuration if DHCP-assigned IP addresses are not available.</li> </ul>

## SG100 and SG1000 applications

You can configure the StorageGRID services appliances in various ways to provide gateway services as well as redundancy of some grid administration services.

Appliances can be deployed in the following ways:

- Add to a new or existing grid as a Gateway Node
- Add to a new grid as a primary or non-primary Admin Node, or to an existing grid as a non-primary Admin Node
- Operate as a Gateway Node and Admin Node (primary or non-primary) at the same time

The appliance facilitates the use of high availability (HA) groups and intelligent load balancing for S3 or Swift data path connections.

The following examples describe how you can maximize the capabilities of the appliance:

- Use two SG100 or two SG1000 appliances to provide gateway services by configuring them as Gateway Nodes.



Do not deploy the SG100 and SG1000 service appliances in the same site. Unpredictable performance might result.

- Use two SG100 or two SG1000 appliances to provide redundancy of some grid administration services. Do this by configuring each appliance as Admin Nodes.
- Use two SG100 or two SG1000 appliances to provide highly available load balancing and traffic shaping services accessed through one or more virtual IP addresses. Do this by configuring the appliances as any combination of Admin Nodes or Gateway Nodes and adding both nodes to the same HA group.



If you use Admin Nodes and Gateway Nodes in the same HA group, CLB (Connection Load Balancer) ports and Admin Node-only ports will not fail over. For instructions for configuring HA groups, see the instructions for administering StorageGRID.



The CLB service is deprecated.

When used with StorageGRID storage appliances, both the SG100 and the SG1000 services appliances enable deployment of appliance-only grids with no dependencies on external hypervisors or compute hardware.

### Related information

[Administer StorageGRID](#)

## Installation and deployment overview

You can install one or more StorageGRID services appliances when you first deploy StorageGRID, or you can add services appliance nodes later as part of an expansion.

### What you'll need

Your StorageGRID system is using the required version of StorageGRID software.

Appliance	Required StorageGRID version
SG100	11.4 or later (latest hotfix recommended)
SG1000	11.3 or later (latest hotfix recommended)

### Installation and deployment tasks

Preparing and adding a StorageGRID appliance to the grid includes four primary steps:

1. Preparing for installation:
  - Preparing the installation site
  - Unpacking the boxes and checking the contents
  - Obtaining additional equipment and tools
  - Verifying network configuration
  - Optional: Configuring an external key management server (KMS) if you plan to encrypt all appliance data. See details about external key management in the instructions for administering StorageGRID.
2. Installing the hardware:
  - Registering the hardware
  - Installing the appliance into a cabinet or rack
  - Cabling the appliance
  - Connecting the power cord and applying power
  - Viewing boot-up status codes
3. Configuring the hardware:
  - Accessing StorageGRID Appliance Installer and configuring the link and network IP settings required to connect to StorageGRID networks
  - Accessing the baseboard management controller (BMC) interface on the appliance.
  - Optional: Enabling node encryption if you plan to use an external KMS to encrypt appliance data.
4. Deploying an appliance Gateway or Admin Node

After the appliance hardware has been installed and configured, you can deploy the appliance as a Gateway Node and an Admin Node in a StorageGRID system. Both the SG100 and the SG1000 appliances can operate as Gateway Nodes and Admin Nodes (primary and non-primary) at the same time.

Task	Instructions
Deploying an appliance Gateway or Admin Node in a new StorageGRID system	<a href="#">Deploying a services appliance node</a>
Adding an appliance Gateway or Admin Node to an existing StorageGRID system	<a href="#">Instructions for expanding a StorageGRID system</a>
Deploying an appliance Gateway or Admin Node as part of a node recovery operation	<a href="#">Instructions for recovery and maintenance</a>

### Related information

[Preparing for installation](#)

[Installing the hardware](#)

[Configuring StorageGRID connections](#)

[Expand your grid](#)

[Maintain & recover](#)

[Administer StorageGRID](#)

## Preparing for installation

Preparing to install a StorageGRID appliance entails preparing the site and obtaining all required hardware, cables, and tools. You should also gather IP addresses and network information.

### Steps

- [Preparing the site \(SG100 and SG1000\)](#)
- [Unpacking the boxes \(SG100 and SG1000\)](#)
- [Obtaining additional equipment and tools \(SG100 and SG1000\)](#)
- [Web browser requirements](#)
- [Reviewing appliance network connections](#)
- [Gathering installation information \(SG100 and SG1000\)](#)

### Preparing the site (SG100 and SG1000)

Before installing the appliance, you must make sure that the site and the cabinet or rack you plan to use meet the specifications for a StorageGRID appliance.

### Steps

1. Confirm that the site meets the requirements for temperature, humidity, altitude range, airflow, heat dissipation, wiring, power, and grounding. See the NetApp Hardware Universe for more information.
2. Confirm that your location provides the correct voltage of AC power (in the range of 120 to 240 volts AC).
3. Obtain a 19-inch (48.3-cm) cabinet or rack to fit shelves of this size (without cables):

Height	Width	Depth	Maximum weight
1.70 in. (4.32 cm)	17.32 in. (44.0 cm)	32.0 in. (81.3 cm)	39 lb. (17.7 kg)

4. Decide where you are going to install the appliance.

#### Related information

[NetApp Hardware Universe](#)

[NetApp Interoperability Matrix Tool](#)

#### Unpacking the boxes (SG100 and SG1000)

Before installing the StorageGRID appliance, unpack all boxes and compare the contents to the items on the packing slip.

#### Appliance hardware

- **SG100 or SG1000**



- **Rail kit with instructions**



#### Power cords

The shipment for the StorageGRID appliance includes the following power cords:

- **Two power cords for your country**




Your cabinet might have special power cords that you use instead of the power cords that ship with the appliance.

## Obtaining additional equipment and tools (SG100 and SG1000)

Before installing the StorageGRID appliance, confirm you have all of the additional equipment and tools that you need.

You need the following additional equipment to install and configure the hardware:

- **Screwdrivers**




Phillips No. 2 screwdriver

Medium flat-blade screwdriver

- **ESD wrist strap**



- **Optical cables and transceivers**



- Cable

- TwinAx/Copper (1 to 4)

or

- Fibre/Optical (1 to 4)

- 1 to 4 of each of these transceivers/adapters based on link speed (mixed speeds are not supported)

- SG100:

Link speed (GbE)	Required equipment
10	SFP+ transceiver

Link speed (GbE)	Required equipment
25	SFP28 transceiver

- SG1000:

Link speed (GbE)	Required equipment
10	QSFP-to-SFP adapter (QSA) and SFP+ transceiver
25	QSFP-to-SFP adapter (QSA) and SFP28 transceiver
40	QSFP+ transceiver
100	QFSP28 transceiver

- **RJ-45 (Cat5/Cat5e/Cat6/Cat6a) Ethernet cables**



- **Service laptop**



Supported web browser

1-GbE (RJ-45) port



Some ports might not support 10/100 Ethernet speeds.

- **Optional tools**



Power drill with Phillips head bit

Flashlight

## Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## Reviewing appliance network connections

Before installing the StorageGRID appliance, you should understand which networks can be connected to the appliance.

When you deploy a StorageGRID appliance as a node in a StorageGRID system, you can connect it to the following networks:

- **Grid Network for StorageGRID:** The Grid Network is used for all internal StorageGRID traffic. It provides connectivity between all nodes in the grid, across all sites and subnets. The Grid Network is required.
- **Admin Network for StorageGRID:** The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites. The Admin Network is optional.
- **Client Network for StorageGRID:** The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network provides client protocol access to the grid, so the Grid Network can be isolated and secured. You can configure the Client Network so that the appliance can be accessed over this network using only the ports you choose to open. The Client Network is optional.
- **BMC management network for the services appliance:** This network provides access to the baseboard management controller in the SG100 and SG1000, appliances allowing you to monitor and manage the hardware components in the appliance. This management network can be the same as the Admin Network for StorageGRID, or it can be an independent management network.

## Related information

[Gathering installation information \(SG100 and SG1000\)](#)

## Cabling the appliance SG100 and SG1000)

### Network guidelines

### Grid primer

#### Port bond modes for the SG100 and SG1000 appliances

When configuring network links for the SG100 and SG1000 appliances, you can use port bonding for the ports that connect to the Grid Network and optional Client Network, and the 1-GbE management ports that connect to the optional Admin Network. Port bonding helps protect your data by providing redundant paths between StorageGRID networks and the appliance.


#### Network bond modes

The networking ports on the services appliance support Fixed port bond mode or Aggregate port bond mode for the Grid Network and Client Network connections.


#### Fixed port bond mode

Fixed port bond mode is the default configuration for the networking ports.

#### SG100 fixed port bond mode



#### SG1000 fixed port bond mode



	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

When using Fixed port bond mode, the ports can be bonded using active-backup mode or Link Aggregation Control Protocol mode (LACP 802.3ad).

- In active-backup mode (default), only one port is active at a time. If the active port fails, its backup port

automatically provides a failover connection. Port 4 provides a backup path for port 2 (Grid Network), and port 3 provides a backup path for port 1 (Client Network).

- In LACP mode, each pair of ports forms a logical channel between the services appliance and the network, allowing for higher throughput. If one port fails, the other port continues to provide the channel. Throughput is reduced, but connectivity is not impacted.




If you do not need redundant connections, you can use only one port for each network. However, be aware that the **Services appliance link down** alert might be triggered in the Grid Manager after StorageGRID is installed, indicating that a cable is unplugged. You can safely disable this alert rule.


### Aggregate port bond mode

Aggregate port bond mode significantly increases the throughput for each StorageGRID network and provides additional failover paths.

#### SG100 aggregate port bond mode



#### SG1000 aggregate port bond mode



	Which ports are bonded
1	All connected ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

If you plan to use aggregate port bond mode:

- You must use LACP network bond mode.
- You must specify a unique VLAN tag for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.
- The ports must be connected to switches that can support VLAN and LACP. If multiple switches are participating in the LACP bond, the switches must support multi-chassis link aggregation groups (MLAG), or equivalent.
- You must understand how to configure the switches to use VLAN, LACP, and MLAG, or equivalent.

If you do not want to use all four ports, you can use one, two, or three ports. Using more than one port maximizes the chance that some network connectivity will remain available if one of the ports fails.



If you choose to use fewer than four network ports, be aware that a **Services appliance link down** alert might be triggered in the Grid Manager after the appliance node is installed, indicating that a cable is unplugged. You can safely disable this alert rule for the triggered alert.

### Network bond modes for the management ports

For the two 1-GbE management ports on the services appliance, you can choose Independent network bond mode or Active-Backup network bond mode to connect to the optional Admin Network.

#### SG100 network management ports



#### SG1000 network management ports



In Independent mode, only the management port on the left is connected to the Admin Network. This mode does not provide a redundant path. The management port on the right is unconnected and available for temporary local connections (uses IP address 169.254.0.1)

In Active-Backup mode, both management ports are connected to the Admin Network. Only one port is active at a time. If the active port fails, its backup port automatically provides a failover connection. Bonding these two physical ports into one logical management port provides a redundant path to the Admin Network.



If you need to make a temporary local connection to the services appliance when the 1-GbE management ports are configured for Active-Backup mode, remove the cables from both management ports, plug your temporary cable into the management port on the right, and access the appliance using IP address 169.254.0.1.

	Network bond mode
A	Active-Backup mode. Both management ports are bonded into one logical management port connected to the Admin Network.

	<b>Network bond mode</b>
I	Independent mode. The port on the left is connected to the Admin Network. The port on the right is available for temporary local connections (IP address 169.254.0.1).

**Gathering installation information (SG100 and SG1000)**

As you install and configure the StorageGRID appliance, you must make decisions and gather information about Ethernet switch ports, IP addresses, and port and network bond modes. Record the required information for each network you connect to the appliance. These values are required to install and configure the hardware.

**Administration and maintenance ports**

The Admin Network for StorageGRID is an optional network, used for system administration and maintenance. The appliance connects to the Admin Network using the following 1-GbE management ports on the appliance.

**SG100 RJ-45 ports**



**SG1000 RJ-45 ports**



**Administration and maintenance connections**

Information needed	Your value
Admin Network enabled	Choose one: <ul style="list-style-type: none"> <li>No</li> <li>Yes (default)</li> </ul>
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>Independent (default)</li> <li>Active-Backup</li> </ul>
Switch port for the left port circled in the diagram (default active port for Independent network bond mode)	
Switch port for the right port circled in the diagram (Active-Backup network bond mode only)	

Information needed	Your value
<p>MAC address for the Admin Network port</p> <p><b>Note:</b> The MAC address label on the front of the appliance lists the MAC address for the BMC management port. To determine the MAC address for the Admin Network port, you must add <b>2</b> to the hexadecimal number on the label. For example, if the MAC address on the label ends in <b>09</b>, the MAC address for the Admin Port would end in <b>0B</b>. If the MAC address on the label ends in <b>(y)FF</b>, the MAC address for the Admin Port would end in <b>(y+1)01</b>. You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing <b>+ 2 =</b>.</p>	
<p>DHCP-assigned IP address for the Admin Network port, if available after power on</p> <p><b>Note:</b> You can determine the DHCP-assigned IP address by using the MAC address to look up the assigned IP.</p>	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
<p>Static IP address you plan to use for the appliance node on the Admin Network</p> <p><b>Note:</b> If your network does not have a gateway, specify the same static IPv4 address for the gateway.</p>	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
<p>Admin Network subnets (CIDR)</p>	

### Networking ports

The four networking ports on the appliance connect to the StorageGRID Grid Network and the optional Client Network.

### Networking connections

Information needed	Your value
Link speed	<p>For the SG100, choose one of the following:</p> <ul style="list-style-type: none"> <li>• Auto (default)</li> <li>• 10 GbE</li> <li>• 25 GbE</li> </ul> <p>For the SG1000, choose one of the following:</p> <ul style="list-style-type: none"> <li>• Auto (default)</li> <li>• 10 GbE</li> <li>• 25 GbE</li> <li>• 40 GbE</li> <li>• 100 GbE</li> </ul> <p><b>Note:</b> For the SG1000, 10- and 25-GbE speeds require the use of QSA adapters.</p>
Port bond mode	<p>Choose one:</p> <ul style="list-style-type: none"> <li>• Fixed (default)</li> <li>• Aggregate</li> </ul>
Switch port for port 1 (Client Network for Fixed mode)	
Switch port for port 2 (Grid Network for Fixed mode)	
Switch port for port 3 (Client Network for Fixed mode)	
Switch port for port 4 (Grid Network for Fixed mode)	

### Grid Network ports

The Grid Network for StorageGRID is a required network, used for all internal StorageGRID traffic. The appliance connects to the Grid Network using the four network ports.

### Grid Network connections

Information needed	Your value
Network bond mode	<p>Choose one:</p> <ul style="list-style-type: none"> <li>• Active-Backup (default)</li> <li>• LACP (802.3ad)</li> </ul>

Information needed	Your value
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>
VLAN tag(if VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Grid Network, if available after power on	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Static IP address you plan to use for the appliance node on the Grid Network  <b>Note:</b> If your network does not have a gateway, specify the same static IPv4 address for the gateway.	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Grid Network subnets (CIDRs)	
Maximum transmission unit (MTU) setting (optional)You can use the default value of 1500, or set the MTU to a value suitable for jumbo frames, such as 9000.	

### Client Network ports

The Client Network for StorageGRID is an optional network, typically used to provide client protocol access to the grid. The appliance connects to the Client Network using the four network ports.

### Client Network connections

Information needed	Your value
Client Network enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>
Network bond mode	Choose one: <ul style="list-style-type: none"> <li>• Active-Backup (default)</li> <li>• LACP (802.3ad)</li> </ul>
VLAN tagging enabled	Choose one: <ul style="list-style-type: none"> <li>• No (default)</li> <li>• Yes</li> </ul>

Information needed	Your value
VLAN tag(If VLAN tagging is enabled)	Enter a value between 0 and 4095:
DHCP-assigned IP address for the Client Network, if available after power on	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Static IP address you plan to use for the appliance node on the Client Network  <b>Note:</b> If the Client Network is enabled, the default route on the appliance will use the gateway specified here.	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>

### BMC management network ports

You can access the BMC interface on the services appliance using the 1-GbE management port circled in the diagram. This port supports remote management of the controller hardware over Ethernet using the Intelligent Platform Management Interface (IPMI) standard.

### SG100 BMC management port



### SG1000 BMC management port



### BMC management network connections

Information needed	Your value
Ethernet switch port you will connect to the BMC management port (circled in the diagram)	
DHCP-assigned IP address for the BMC management network, if available after power on	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>
Static IP address you plan to use for the BMC management port	<ul style="list-style-type: none"> <li>• IPv4 address (CIDR):</li> <li>• Gateway:</li> </ul>

### Related information

[SG100 and SG1000 appliances overview](#)

[Cabling the appliance \(SG100 and SG1000\)](#)

## Installing the hardware

Hardware installation entails installing the appliance into a cabinet or rack, connecting the cables, and applying power.

### Steps

- [Registering the hardware](#)
- [Installing the appliance into a cabinet or rack \(SG100 and SG1000\)](#)
- [Cabling the appliance SG100 and SG1000\)](#)
- [Connecting power cords and applying power \(SG100 and SG1000\)](#)
- [Viewing status indicators on the SG100 and SG1000 appliances](#)

## Registering the hardware

Registering the appliance hardware provides support benefits.

### Steps

1. Locate the chassis serial number for the appliance.

You can find the number on the packing slip, in your confirmation email, or on the appliance after you unpack it.



2. Go to the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com).
3. Determine whether you need to register the hardware:

If you are a...	Follow these steps...
Existing NetApp customer	<ol style="list-style-type: none"><li>a. Sign in with your username and password.</li><li>b. Select <b>Products &gt; My Products</b>.</li><li>c. Confirm that the new serial number is listed.</li><li>d. If it is not, follow the instructions for new NetApp customers.</li></ol>
New NetApp customer	<ol style="list-style-type: none"><li>a. Click <b>Register Now</b>, and create an account.</li><li>b. Select <b>Products &gt; Register Products</b>.</li><li>c. Enter the product serial number and requested details.</li></ol> <p>After your registration is approved, you can download any required software. The approval process might take up to 24 hours.</p>

## Installing the appliance into a cabinet or rack (SG100 and SG1000)


You must install a set of rails for the appliance in your cabinet or rack, and then slide the appliance onto the rails.

### What you'll need

- You have reviewed the Safety Notices document included in the box, and understand the precautions for moving and installing hardware.
- You have the instructions packaged with the rail kit.


### Steps

1. Carefully follow the instructions for the rail kit to install the rails in your cabinet or rack.
2. On the two rails installed in the cabinet or rack, extend the movable parts of the rails until you hear a click.



3. Insert the appliance into the rails.
4. Slide the appliance into the cabinet or rack.

When you cannot move the appliance any further, pull the blue latches on both sides of the chassis to slide the appliance all the way in.



Do not attach the front bezel until after you power on the appliance.

### Cabling the appliance SG100 and SG1000

You must connect the management port on the appliance to the service laptop and connect the network ports on the appliance to the Grid Network and optional Client Network for StorageGRID.

#### What you'll need

- You have an RJ-45 Ethernet cable for connecting the management port.
- You have one of the following options for the network ports. These items are not provided with the appliance.
  - One to four TwinAx cables for connecting the four network ports.
  - For the SG100, one to four SFP+ or SFP28 transceivers if you plan to use optical cables for the ports.
  - For the SG1000, one to four QSFP+ or QSFP28 transceivers if you plan to use optical cables for the ports.




**Risk of exposure to laser radiation** — Do not disassemble or remove any part of a SFP or QSFP transceiver. You might be exposed to laser radiation.


#### About this task

The following figures show the ports on the back of the appliance.

## SG100 port connections



## SG1000 port connections



	Port	Type of port	Function
1	BMC management port on the appliance	1-GbE (RJ-45)	Connects to the network where you access the BMC interface.
2	Four network ports on the appliance	<ul style="list-style-type: none"> <li>For the SG100: 10/25-GbE</li> <li>For the SG1000: 10/25/40/100-GbE</li> </ul>	Connect to the Grid Network and the Client Network for StorageGRID.
3	Admin Network port on the appliance (labelled P1 in the figures)	1-GbE (RJ-45) <b>Important:</b> This port operates only at 1000 baseT/full and does not support 10- or 100-megabit speeds.	Connects the appliance to the Admin Network for StorageGRID.
3	Rightmost RJ-45 port on the appliance	1-GbE (RJ-45) <b>Important:</b> This port operates only at 1000 baseT/full and does not support 10- or 100-megabit speeds.	<ul style="list-style-type: none"> <li>Can be bonded with management port 1 if you want a redundant connection to the Admin Network.</li> <li>Can be left disconnected and available for temporary local access (IP 169.254.0.1).</li> <li>During installation, can be used to connect the appliance to a service laptop if DHCP-assigned IP addresses are not available.</li> </ul>

## Steps

1. Connect the BMC management port on the appliance to the management network, using an Ethernet cable.

Although this connection is optional, it is recommended to facilitate support.

2. Connect the network ports on the appliance to the appropriate network switches, using TwinAx cables or optical cables and transceivers.



The four network ports must use the same link speed. See the following tables for the equipment required based on your hardware and link speed.

SG100 link speed (GbE)	Required equipment
10	SFP+ transceiver
25	SFP28 transceiver
SG1000 link speed (GbE)	Required equipment
10	QSA and SFP+ transceiver
25	QSA and SFP28 transceiver
40	QSFP+ transceiver
100	QFSP28 transceiver

- If you plan to use Fixed port bond mode (default), connect the ports to the StorageGRID Grid and Client Networks, as shown in the table.

Port	Connects to...
Port 1	Client Network (optional)
Port 2	Grid Network
Port 3	Client Network (optional)
Port 4	Grid Network

- If you plan to use the Aggregate port bond mode, connect one or more of the network ports to one or more switches. You should connect at least two of the four ports to avoid having a single point of failure. If you use more than one switch for a single LACP bond, the switches must support MLAG or equivalent.

3. If you plan to use the Admin Network for StorageGRID, connect the Admin Network port on the appliance to the Admin Network, using an Ethernet cable.

## Connecting power cords and applying power (SG100 and SG1000)

After connecting the network cables, you are ready to apply power to the appliance.

### Steps

1. Connect a power cord to each of the two power supply units in the appliance.
2. Connect these two power cords to two different power distribution units (PDUs) in the cabinet or rack.
3. If the power button on the front of the appliance is not currently illuminated blue, press the button to turn on power to the appliance.

Do not press the power button again during the power-on process.

4. If errors occur, correct any issues.
5. Attach the front bezel to the appliance.

### Related information

[Viewing status indicators on the SG100 and SG1000 appliances](#)

### Viewing status indicators on the SG100 and SG1000 appliances

The appliance includes indicators that help you determine the status of the appliance controller and the two SSDs.

### Appliance indicators and buttons



	Display	State
1	Power button	<ul style="list-style-type: none"><li>• Blue: the appliance is powered on.</li><li>• Off: the appliance is powered off.</li></ul>
2	Reset button	Use this button to perform a hard reset of the controller.
3	Identify button	This button can be set to Blink, On (Solid), or Off. <ul style="list-style-type: none"><li>• Blue, blinking: Identifies the appliance in the cabinet or rack.</li><li>• Blue, solid: Identifies the appliance in the cabinet or rack.</li><li>• Off: The appliance is not visually identifiable in the cabinet or rack.</li></ul>

	Display	State
4	Alarm LED	<ul style="list-style-type: none"> <li>Amber, solid: An error has occurred.</li> </ul> <p><b>Note:</b> To view the boot-up and error codes, you must access the BMC interface.</p> <ul style="list-style-type: none"> <li>Off: No errors are present.</li> </ul>


### General boot-up codes

During boot-up or after a hard reset of the appliance, the following occurs:

1. The baseboard management controller (BMC) logs codes for the boot-up sequence, including any errors that occur.
2. The power button lights up.
3. If any errors occur during boot-up, the alarm LED lights up.

To view the boot-up and error codes, you must access the BMC interface.

### SSD indicators



LED	Display	State
1	Drive status/fault	<ul style="list-style-type: none"> <li>Blue (solid): drive is online</li> <li>Amber (blinking): drive failure</li> <li>Off: slot is empty</li> </ul>
2	Drive active	Blue (blinking): drive is being accessed

### Related information

[Troubleshooting the hardware installation](#)

[Configuring the BMC interface](#)

## Configuring StorageGRID connections

Before you can deploy the services appliance as a node in a StorageGRID system, you must configure the connections between the appliance and the networks you plan to use. You can configure networking by browsing to the StorageGRID Appliance Installer, which

is pre-installed on the services appliance.

### Steps

- [Accessing the StorageGRID Appliance Installer](#)
- [Verifying and upgrading the StorageGRID Appliance Installer version](#)
- [Configuring network links \(SG100 and SG1000\)](#)
- [Configuring StorageGRID IP addresses](#)
- [Verifying network connections](#)
- [Verifying port-level network connections](#)

### Accessing the StorageGRID Appliance Installer

You must access the StorageGRID Appliance Installer to configure the connections between the appliance and the three StorageGRID networks: the Grid Network, the Admin Network (optional), and the Client Network (optional).

### What you'll need

- You are using any management client that can connect to the StorageGRID Admin Network.
- The client has a supported web browser.
- The services appliance is connected to all of the StorageGRID networks you plan to use.
- You know the IP address, gateway, and subnet for the services appliance on these networks.
- You have configured the network switches you plan to use.

### About this task

To initially access the StorageGRID Appliance Installer, you can use the DHCP-assigned IP address for the Admin Network port on the services appliance (assuming it is connected to the Admin Network), or you can connect a service laptop directly to the services appliance.

### Steps

1. If possible, use the DHCP address for the Admin Network port on the services appliance to access the StorageGRID Appliance Installer.

#### SG100 Admin Network port



#### SG1000 Admin Network port



- a. Locate the MAC address label on the front of the services appliance, and determine the MAC address for the Admin Network port.

The MAC address label lists the MAC address for the BMC management port.

To determine the MAC address for the Admin Network port, you must add **2** to the hexadecimal number on the label. For example, if the MAC address on the label ends in **09**, the MAC address for the Admin Port would end in **0B**. If the MAC address on the label ends in **(y)FF**, the MAC address for the Admin Port would end in **(y+1)01**. You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing **+ 2 =**.

b. Provide the MAC address to your network administrator, so they can look up the DHCP address for the appliance on the Admin Network.

c. From the client, enter this URL for the StorageGRID Appliance Installer:

**`https://services-appliance_IP:8443`**

For *services-appliance\_IP*, use the DHCP address.

d. If you are prompted with a security alert, view and install the certificate using the browser's installation wizard.

The alert will not appear the next time you access this URL.

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected to StorageGRID networks. Error messages might appear that will be resolved in later steps.

2. Alternatively, if you cannot obtain an IP address using DHCP, use a link-local connection to access the StorageGRID Appliance Installer.

a. Connect a service laptop directly to the rightmost RJ-45 port on the services appliance, using an Ethernet cable.

#### SG100 link-local connection



#### SG1000 link-local connection



b. Open a web browser.

c. Enter this URL for the StorageGRID Appliance Installer:

**`https://169.254.0.1:8443`**

The StorageGRID Appliance Installer Home page appears. The information and messages shown when you first access this page depend on how your appliance is currently connected to StorageGRID networks. Error messages might appear that will be resolved in later steps.



If you cannot access the Home page over a link-local connection, configure the service laptop IP address as 169.254.0.2, and try again.

3. Review any messages displayed on the Home page and configure the link configuration and the IP configuration, as required.

## NetApp® StorageGRID® Appliance Installer

Home

Configure Networking ▾

Configure Hardware ▾

Monitor Installation

Advanced ▾

Home

This Node

Node type	<input type="text" value="Gateway"/> ▾
Node name	<input type="text" value="xlr8r-10"/>
	<input type="button" value="Cancel"/> <input type="button" value="Save"/>

Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

Connection state Connection to 192.168.7.44 ready

Installation

Current state Ready to start installation of xlr8r-10 into grid with Admin Node 192.168.7.44 running StorageGRID 11.4.0, using StorageGRID software downloaded from the Admin Node.

### Related information

[Web browser requirements](#)

### Verifying and upgrading the StorageGRID Appliance Installer version

The StorageGRID Appliance Installer version on the appliance must match the software version installed on your StorageGRID system to ensure that all StorageGRID features are supported.

### What you'll need

You have accessed the StorageGRID Appliance Installer.

### About this task

StorageGRID appliances come from the factory preinstalled with the StorageGRID Appliance Installer. If you are adding an appliance to a recently upgraded StorageGRID system, you might need to manually upgrade the StorageGRID Appliance Installer before installing the appliance as a new node.

The StorageGRID Appliance Installer automatically upgrades when you upgrade to a new StorageGRID version. You do not need to upgrade the StorageGRID Appliance Installer on installed appliance nodes. This procedure is only required when you are installing an appliance that contains an earlier version of the StorageGRID Appliance Installer.

## Steps

1. From the StorageGRID Appliance Installer, select **Advanced > Upgrade Firmware**.
2. Compare the Current Firmware version to the software version installed on your StorageGRID system (from the Grid Manager select **Help > About**).

The second digit in the two versions should match. For example, if your StorageGRID system is running version 11.5.x.y, the StorageGRID Appliance Installer version should be 3.5.z.

3. If the appliance has a down-level version of the StorageGRID Appliance Installer, go to the NetApp Downloads page for StorageGRID.

### [NetApp Downloads: StorageGRID](#)

Sign in with the username and password for your NetApp account.

4. Download the appropriate version of the **Support file for StorageGRID Appliances** and the corresponding checksum file.

The Support file for StorageGRID Appliances file is a .zip archive that contains the current and previous firmware versions for all StorageGRID appliance models, in subdirectories for each controller type.

After downloading the Support file for StorageGRID Appliances file, extract the .zip archive and see the README file for important information about installing the StorageGRID Appliance Installer.

5. Follow the instructions on the Upgrade Firmware page of the StorageGRID Appliance Installer to perform these steps:
  - a. Upload the appropriate support file (firmware image) for your controller type and the checksum file.
  - b. Upgrade the inactive partition.
  - c. Reboot and swap partitions.
  - d. Upgrade the second partition.

## Related information

[Accessing the StorageGRID Appliance Installer](#)

## Configuring network links (SG100 and SG1000)

You can configure network links for the ports used to connect the appliance to the Grid Network, the Client Network, and the Admin Network. You can set the link speed as well as the port and network bond modes.

## What you'll need

- You have obtained the additional equipment required for your cable type and link speed.
- You have connected the network ports to switches that support your chosen speed.


If you plan to use Aggregate port bond mode, LACP network bond mode, or VLAN tagging:

- You have connected the network ports on the appliance to switches that can support VLAN and LACP.
- If multiple switches are participating in the LACP bond, the switches support multi-chassis link aggregation groups (MLAG), or equivalent.
- You understand how to configure the switches to use VLAN, LACP, and MLAG or equivalent.
- You know the unique VLAN tag to use for each network. This VLAN tag will be added to each network packet to ensure that network traffic is routed to the correct network.


**About this task**

The figures show how the four network ports are bonded in fixed port bond mode (default configuration).

**SG100 fixed port bond mode**




**SG1000 fixed port bond mode**




	Which ports are bonded
C	Ports 1 and 3 are bonded together for the Client Network, if this network is used.
G	Ports 2 and 4 are bonded together for the Grid Network.

This figure shows how the four network ports are bonded in aggregate port bond mode.

**SG100 aggregate port bond mode**



**SG1000 aggregate port bond mode**



	<b>Which ports are bonded</b>
1	All four ports are grouped in a single LACP bond, allowing all ports to be used for Grid Network and Client Network traffic.

The table summarizes the options for configuring the four network ports. The default settings are shown in bold. You only need to configure the settings on the Link Configuration page if you want to use a non-default setting.



The LACP transmit hash policy defaults to layer2+3 mode. If necessary, you can use the Grid Management API to change it to layer3+4 mode.

• **Fixed (default) port bond mode**

<b>Network bond mode</b>	<b>Client Network disabled (default)</b>	<b>Client Network enabled</b>
<b>Active-Backup (default)</b>	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an active-backup bond for the Grid Network.</li> <li>• Ports 1 and 3 are not used.</li> <li>• A VLAN tag is optional.</li> </ul>	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an active-backup bond for the Grid Network.</li> <li>• Ports 1 and 3 use an active-backup bond for the Client Network.</li> <li>• VLAN tags can be specified for both networks for the convenience of the network administrator.</li> </ul>
LACP (802.3ad)	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an LACP bond for the Grid Network.</li> <li>• Ports 1 and 3 are not used.</li> <li>• A VLAN tag is optional.</li> </ul>	<ul style="list-style-type: none"> <li>• Ports 2 and 4 use an LACP bond for the Grid Network.</li> <li>• Ports 1 and 3 use an LACP bond for the Client Network.</li> <li>• VLAN tags can be specified for both networks for the convenience of the network administrator.</li> </ul>

• **Aggregate port bond mode**

Network bond mode	Client Network disabled (default)	Client Network enabled
LACP (802.3ad) only	<ul style="list-style-type: none"> <li>Ports 1-4 use a single LACP bond for the Grid Network.</li> <li>A single VLAN tag identifies Grid Network packets.</li> </ul>	<ul style="list-style-type: none"> <li>Ports 1-4 use a single LACP bond for the Grid Network and the Client Network.</li> <li>Two VLAN tags allow Grid Network packets to be segregated from Client Network packets.</li> </ul>

For additional details, see the article about GbE port connections for the services appliance.

This figure shows how the two 1-GbE management ports on the SG100 are bonded in Active-Backup network bond mode for the Admin Network.

These figures shows how the two 1-GbE management ports on the appliance are bonded in Active-Backup network bond mode for the Admin Network.

### SG100 Admin Network ports bonded



### SG1000 Admin Network ports bonded




### Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Link Configuration**.

The Network Link Configuration page displays a diagram of your appliance with the network and management ports numbered.

### SG100 ports


## Network Link Configuration



⚠ You might lose your connection if you make changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL using one of the other IP addresses assigned to the appliance.

## SG1000 ports

### Network Link Configuration



⚠ You might lose your connection if you make changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL using one of the other IP addresses assigned to the appliance.

The Link Status table lists the link state and speed of the numbered ports (SG1000 shown).

### Link Status

Link	State	Speed (Gbps)
1	Up	100
2	Down	N/A
3	Down	N/A
4	Down	N/A
5	Up	1
6	Up	1

The first time you access this page:

- **Link Speed** is set to **Auto**.
- **Port bond mode** is set to **Fixed**.

- **Network bond mode** is set to **Active-Backup** for the Grid Network.
- The **Admin Network** is enabled, and the network bond mode is set to **Independent**.
- The **Client Network** is disabled.

### Link Settings

Link speed

Port bond mode  Fixed  Aggregate

Choose Fixed port bond mode if you want to use ports 2 and 4 for the Grid Network and ports 1 and 3 for the Client Network (if enabled). Choose Aggregate port bond mode if you want all connected ports to share a single LACP bond for both the Grid and Client Networks.

### Grid Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

Enable VLAN (802.1q) tagging

MAC Addresses 50:6b:4b:42:d7:00 50:6b:4b:42:d7:01 50:6b:4b:42:d7:24 50:6b:4b:42:d7:25

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

### Admin Network

Enable network

Network bond mode  Independent  Active-Backup

Connect the Admin Network to port 5. Leave port 6 unconnected. If necessary, you can make a temporary direct Ethernet connection to port 6 and use link-local IP address 169.254.0.1 for access.

MAC Addresses d8:c4:97:2a:e4:95

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

### Client Network

Enable network

Enabling the Client Network causes the default gateway for this node to move to the Client Network. Before enabling the Client Network, ensure that you've added all necessary subnets to the Grid Network Subnet List. Otherwise, the connection to the node might be lost.

2. Select the link speed for the network ports from the **Link speed** drop-down list.

The network switches you are using for the Grid Network and the Client Network must also support and be

configured for this speed. You must use the appropriate adapters or transceivers for the configured link speed. Use Auto link speed when possible because this option negotiates both link speed and Forward Error Correction (FEC) mode with the link partner.

3. Enable or disable the StorageGRID networks you plan to use.

The Grid Network is required. You cannot disable this network.

- a. If the appliance is not connected to the Admin Network, unselect the **Enable network** check box for the Admin Network.

#### Admin Network

---

Enable network

- b. If the appliance is connected to the Client Network, select the **Enable network** check box for the Client Network.

The Client Network settings for the data NIC ports are now shown.

4. Refer to the table, and configure the port bond mode and the network bond mode.

This example shows:

- **Aggregate** and **LACP** selected for the Grid and the Client networks. You must specify a unique VLAN tag for each network. You can select values between 0 and 4095.
- **Active-Backup** selected for the Admin Network.

## Link Settings

Link speed

Port bond mode  Fixed  Aggregate

Choose Fixed port bond mode if you want to use ports 2 and 4 for the Grid Network and ports 1 and 3 for the Client Network (if enabled). Choose Aggregate port bond mode if you want all connected ports to share a single LACP bond for both the Grid and Client Networks.

## Grid Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

MAC Addresses 50:6b:4b:42:d7:00 50:6b:4b:42:d7:01 50:6b:4b:42:d7:24 50:6b:4b:42:d7:25

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

## Admin Network

Enable network

Network bond mode  Independent  Active-Backup

Connect the Admin Network to ports 5 and 6. If necessary, you can make a temporary direct Ethernet connection by disconnecting ports 5 and 6, then connecting to port 6 and using link-local IP address 169.254.0.1 for access.

MAC Addresses d8:c4:97:2a:e4:95

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

## Client Network

Enable network

Network bond mode  Active-Backup  LACP (802.3ad)

If the port bond mode is Aggregate, all bonds must be in LACP (802.3ad) mode.

Enable VLAN (802.1q) tagging

VLAN (802.1q) tag

MAC Addresses 50:6b:4b:42:d7:00 50:6b:4b:42:d7:01 50:6b:4b:42:d7:24 50:6b:4b:42:d7:25

If you are using DHCP, it is recommended that you configure a permanent DHCP reservation. Use all of these MAC addresses in the reservation to assign one IP address to this network interface.

5. When you are satisfied with your selections, click **Save**.



You might lose your connection if you made changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:

**`https://services_appliance_IP:8443`**

### Related information

[Obtaining additional equipment and tools \(SG100 and SG1000\)](#)

### Configuring StorageGRID IP addresses

You use the StorageGRID Appliance Installer to configure the IP addresses and routing information used for the services appliance on the StorageGRID Grid, Admin, and Client Networks.

#### About this task

You must either assign a static IP for the appliance on each connected network or assign a permanent lease for the address on the DHCP server.

If you want to change the link configuration, see the instructions for changing the link configuration of the services appliance.

#### Steps

1. In the StorageGRID Appliance Installer, select **Configure Networking > IP Configuration**.

The IP Configuration page appears.

2. To configure the Grid Network, select either **Static** or **DHCP** in the **Grid Network** section of the page.

## Grid Network

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

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

 All required Grid Network subnets must also be defined in the Grid Network Subnet List on the Primary Admin Node before starting installation.

Subnets (CIDR)  



MTU  

3. If you selected **Static**, follow these steps to configure the Grid Network:

- Enter the static IPv4 address, using CIDR notation.
- Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

- If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

**`https://services_appliance_IP:8443`**

e. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

f. Click **Save**.

4. If you selected **DHCP**, follow these steps to configure the Grid Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Grid Network subnets is correct.

If you have grid subnets, the Grid Network gateway is required. All grid subnets specified must be reachable through this gateway. These Grid Network subnets must also be defined in the Grid Network Subnet List on the primary Admin Node when you start StorageGRID installation.



The default route is not listed. If the Client Network is not enabled, the default route will use the Grid Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.



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

d. Click **Save**.

5. To configure the Admin Network, select either **Static** or **DHCP** in the Admin Network section of the page.



To configure the Admin Network, you must enable the Admin Network on the Link Configuration page.

### Admin Network

The Admin Network is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites.

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

Subnets (CIDR)  +

MTU

6. If you selected **Static**, follow these steps to configure the Admin Network:

a. Enter the static IPv4 address, using CIDR notation, for Management Port 1 on the appliance.

Management Port 1 is the left of the two 1-GbE RJ45 ports on the right end of the appliance.

b. Enter the gateway.

If your network does not have a gateway, re-enter the same static IPv4 address.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as

9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

When you change the IP address, the gateway and list of subnets might also change.

If you lose your connection to the StorageGRID Appliance Installer, re-enter the URL using the new static IP address you just assigned. For example,

**`https://services_appliance:8443`**

e. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route cannot be made to use the Admin Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

f. Click **Save**.

7. If you selected **DHCP**, follow these steps to configure the Admin Network:

a. After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address**, **Gateway**, and **Subnets** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

b. Confirm that the list of Admin Network subnets is correct.

You must verify that all subnets can be reached using the gateway you provided.



The default route cannot be made to use the Admin Network gateway.

- To add a subnet, click the insert icon **+** to the right of the last entry.
- To remove an unused subnet, click the delete icon **x**.

c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

d. Click **Save**.

8. To configure the Client Network, select either **Static** or **DHCP** in the **Client Network** section of the page.



To configure the Client Network, you must enable the Client Network on the Link Configuration page.

## Client Network

The Client Network is an open network used to provide access to client applications, including S3 and Swift. The Client Network enables grid nodes to communicate with any subnet reachable through the Client Network gateway. The Client Network does not become operational until you complete the StorageGRID configuration steps.

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

MTU

9. If you selected **Static**, follow these steps to configure the Client Network:

- Enter the static IPv4 address, using CIDR notation.
- Click **Save**.
- Confirm that the IP address for the Client Network gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and cannot be moved to another interface while the Client Network is enabled.

- If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

- Click **Save**.

10. If you selected **DHCP**, follow these steps to configure the Client Network:

- After you select the **DHCP** radio button, click **Save**.

The **IPv4 Address** and **Gateway** fields are automatically populated. If the DHCP server is set up to assign an MTU value, the **MTU** field is populated with that value, and the field becomes read-only.

Your web browser is automatically redirected to the new IP address for the StorageGRID Appliance Installer.

- b. Confirm that the gateway is correct.



If the Client Network is enabled, the default route is displayed. The default route uses the Client Network gateway and cannot be moved to another interface while the Client Network is enabled.

- c. If you want to use jumbo frames, change the MTU field to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value of 1500.



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.

## Related information

[Changing the link configuration of the services appliance](#)

## Verifying network connections

You should confirm you can access the StorageGRID networks you are using from the appliance. To validate routing through network gateways, you should test connectivity between the StorageGRID Appliance Installer and IP addresses on different subnets. You can also verify the MTU setting.

### Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Configure Networking > Ping and MTU Test**.

The Ping and MTU Test page appears.

### Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

#### Ping and MTU Test

Network

Destination IPv4 Address or FQDN

Test MTU

2. From the **Network** drop-down box, select the network you want to test: Grid, Admin, or Client.

3. Enter the IPv4 address or fully qualified domain name (FQDN) for a host on that network.

For example, you might want to ping the gateway on the network or the primary Admin Node.

4. Optionally, select the **Test MTU** check box to verify the MTU setting for the entire path through the network to the destination.

For example, you can test the path between the appliance node and a node at a different site.

5. Click **Test Connectivity**.

If the network connection is valid, the "Ping test passed" message appears, with the ping command output listed.

### Ping and MTU Test

Use a ping request to check the appliance's connectivity to a remote host. Select the network you want to check connectivity through, and enter the IP address of the host you want to reach. To verify the MTU setting for the entire path through the network to the destination, select Test MTU.

#### Ping and MTU Test

Network	<input type="text" value="Grid"/>
Destination IPv4 Address or FQDN	<input type="text" value="10.96.104.223"/>
Test MTU	<input checked="" type="checkbox"/>
<input type="button" value="Test Connectivity"/>	

Ping test passed

#### Ping command output

```
PING 10.96.104.223 (10.96.104.223) 1472(1500) bytes of data.  
1480 bytes from 10.96.104.223: icmp_seq=1 ttl=64 time=0.318 ms  
  
--- 10.96.104.223 ping statistics ---  
1 packets transmitted, 1 received, 0% packet loss, time 0ms  
rtt min/avg/max/mdev = 0.318/0.318/0.318/0.000 ms  
  
Found MTU 1500 for 10.96.104.223 via br0
```

### Related information

[Configuring network links \(SG100 and SG1000\)](#)

[Changing the MTU setting](#)

### Verifying port-level network connections

To ensure that access between the StorageGRID Appliance Installer and other nodes is

not obstructed by firewalls, confirm that the StorageGRID Appliance Installer can connect to a specific TCP port or set of ports at the specified IP address or range of addresses.

### About this task

Using the list of ports provided in the StorageGRID Appliance Installer, you can test the connectivity between the appliance and the other nodes in your Grid Network.

Additionally, you can test connectivity on the Admin and Client Networks and on UDP ports, such as those used for external NFS or DNS servers. For a list of these ports, see the port reference in the StorageGRID networking guidelines.



The Grid Network ports listed in the port connectivity table are valid only for StorageGRID version 11.5.0. To verify which ports are correct for each node type, you should always consult the networking guidelines for your version of StorageGRID.

### Steps

1. From the StorageGRID Appliance Installer, click **Configure Networking > Port Connectivity Test (nmap)**.

The Port Connectivity Test page appears.

The port connectivity table lists node types that require TCP connectivity on the Grid Network. For each node type, the table lists the Grid Network ports that should be accessible to your appliance.

The following node types require TCP connectivity on the Grid Network.

Node Type	Grid Network Ports
Admin Node	22,80,443,1504,1505,1506,1508,7443,9999
Storage Node without ADC	22,1139,1502,1506,1511,7001,9042,9999,18002,18017,18019,18082,18083,18200
Storage Node with ADC	22,1139,1501,1502,1506,1511,7001,9042,9999,18000,18001,18002,18003,18017,18019,18082,18083,18200,19000
API Gateway	22,1506,1507,9999
Archive Node	22,1506,1509,9999,11139

You can test the connectivity between the appliance ports listed in the table and the other nodes in your Grid Network.

2. From the **Network** drop-down, select the network you want to test: **Grid**, **Admin**, or **Client**.
3. Specify a range of IPv4 addresses for the hosts on that network.

For example, you might want to probe the gateway on the network or the primary Admin Node.

Specify a range using a hyphen, as shown in the example.

4. Enter a TCP port number, a list of ports separated by commas, or a range of ports.

The following node types require TCP connectivity on the Grid Network.

Node Type	Grid Network Ports
Admin Node	22,80,443,1504,1505,1506,1508,7443,9999
Storage Node without ADC	22,1139,1502,1506,1511,7001,9042,9999,18002,18017,18019,18082,18083,18200
Storage Node with ADC	22,1139,1501,1502,1506,1511,7001,9042,9999,18000,18001,18002,18003,18017,18019,18082,18083,18200,19000
API Gateway	22,1506,1507,9999
Archive Node	22,1506,1509,9999,11139

### Port Connectivity Test

Network

IPv4 Address Ranges

Port Ranges

Protocol  TCP  UDP

### 5. Click **Test Connectivity**.

- If the selected port-level network connections are valid, the “Port connectivity test passed” message appears in a green banner. The nmap command output is listed below the banner.

Port connectivity test passed

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Fri Nov 13 18:32:03 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,2022 10.224.6.160-161
Nmap scan report for 10.224.6.160
Host is up (0.00072s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp  open  down

Nmap scan report for 10.224.6.161
Host is up (0.00060s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
2022/tcp  open  down

# Nmap done at Fri Nov 13 18:32:04 2020 -- 2 IP addresses (2 hosts up) scanned in 0.55 seconds
```

- If a port-level network connection is made to the remote host, but the host is not listening on one or more of the selected ports, the “Port connectivity test failed” message appears in a yellow banner. The nmap command output is listed below the banner.

Any remote port the host is not listening to has a state of “closed.” For example, you might see this yellow banner when the node you are trying to connect to is in a pre-installed state and the StorageGRID NMS service is not yet running on that node.

 Port connectivity test failed  
Connection not established. Services might not be listening on target ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:07:02 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,80,443,1504,1505,1506,1508,7443,9999
Nmap scan report for 172.16.4.71
Host is up (0.00020s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
80/tcp    open  http
443/tcp   open  https
1504/tcp   closed evb-elm
1505/tcp   open  funkproxy
1506/tcp   open  utcd
1508/tcp   open  diagmond
7443/tcp   open  oracleas-https
9999/tcp   open  abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:07:03 2020 -- 1 IP address (1 host up) scanned in 0.59 seconds
```

- If a port-level network connection cannot be made for one or more selected ports, the “Port connectivity test failed” message appears in a red banner. The nmap command output is listed below the banner.

The red banner indicates that a TCP connection attempt to a port on the remote host was made, but nothing was returned to the sender. When no response is returned, the port has a state of “filtered” and is likely blocked by a firewall.



Ports with “closed” are also listed.

 Port connectivity test failed  
Connection failed to one or more ports.

Nmap command output. Note: Unreachable hosts will not appear in the output.

```
# Nmap 7.70 scan initiated Sat May 16 17:11:01 2020 as: /usr/bin/nmap -n -oN - -e br0 -p 22,79,80,443,1504,1505,1506,1508,7443,9999 172.16.4.71
Nmap scan report for 172.16.4.71
Host is up (0.00029s latency).

PORT      STATE SERVICE
22/tcp    open  ssh
79/tcp    filtered finger
80/tcp    open  http
443/tcp   open  https
1504/tcp   closed evb-elm
1505/tcp   open  funkproxy
1506/tcp   open  utcd
1508/tcp   open  diagmond
7443/tcp   open  oracleas-https
9999/tcp   open  abyss
MAC Address: 00:50:56:87:39:AE (VMware)

# Nmap done at Sat May 16 17:11:02 2020 -- 1 IP address (1 host up) scanned in 1.60 seconds
```

## Related information

[Network guidelines](#)

## Configuring the BMC interface

The user interface for the baseboard management controller (BMC) on the services appliance provides status information about the hardware and allows you to configure

SNMP settings and other options for the services appliance.

### Steps

- [Changing the root password for the BMC interface](#)
- [Setting the IP address for the BMC management port](#)
- [Accessing the BMC interface](#)
- [Configuring SNMP settings for the services appliance](#)
- [Setting up email notifications for alerts](#)

### Changing the root password for the BMC interface

For security, you must change the password for the BMC's root user.

#### What you'll need

The management client is using a supported web browser.

#### About this task

When you first install the appliance, the BMC uses a default password for the root user (`root/calvin`). You must change the password for the root user to secure your system.


### Steps

1. From the client, enter the URL for the StorageGRID Appliance Installer:  
**`https://services_appliance_IP:8443`**

For `services_appliance_IP`, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > BMC Configuration**.



The Baseboard Management Controller Configuration page appears.

3. Enter a new password for the root account in the two fields provided.

## Baseboard Management Controller Configuration

### User Settings

Root Password

.....

Confirm Root Password

.....

4. Click **Save**.

### Setting the IP address for the BMC management port

Before you can access the BMC interface, you must configure the IP address for the BMC management port on the services appliance.

#### What you'll need

- The management client is using a supported web browser.
- You are using any management client that can connect to a StorageGRID network.
- The BMC management port is connected to the management network you plan to use.

#### SG100 BMC management port



#### SG1000 BMC management port



#### About this task



For support purposes, the BMC management port allows low-level hardware access. You should only connect this port to a secure, trusted, internal management network. If no such network is available, leave the BMC port unconnected or blocked, unless a BMC connection is requested by technical support.


#### Steps

1. From the client, enter the URL for the StorageGRID Appliance Installer:  
**`https://services_appliance_IP:8443`**

For *services\_appliance\_IP*, use the IP address for the appliance on any StorageGRID network.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > BMC Configuration**.



The Baseboard Management Controller Configuration page appears.

3. Make a note of the IPv4 address that is automatically displayed.

DHCP is the default method for assigning an IP address to this port.



It might take a few minutes for the DHCP values to appear.

#### Baseboard Management Controller Configuration

##### LAN IP Settings

IP Assignment	<input type="radio"/> Static	<input checked="" type="radio"/> DHCP
MAC Address	<input type="text" value="d8:c4:97:28:50:62"/>	
IPv4 Address (CIDR)	<input type="text" value="10.224.3.225/21"/>	
Default gateway	<input type="text" value="10.224.0.1"/>	

<input type="button" value="Cancel"/>	<input type="button" value="Save"/>
---------------------------------------	-------------------------------------

4. Optionally, set a static IP address for the BMC management port.



You should either assign a static IP for the BMC management port or assign a permanent lease for the address on the DHCP server.

- a. Select **Static**.
- b. Enter the IPv4 address, using CIDR notation.
- c. Enter the default gateway.

## Baseboard Management Controller Configuration

### LAN IP Settings

IP Assignment	<input checked="" type="radio"/> Static <input type="radio"/> DHCP
MAC Address	d8:c4:97:28:50:62
IPv4 Address (CIDR)	10.224.3.225/21
Default gateway	10.224.0.1

d. Click **Save**.

It might take a few minutes for your changes to be applied.

### Accessing the BMC interface

You can access the BMC interface on the services appliance using the DHCP or static IP address for the BMC management port.

#### What you'll need

- The management client is using a supported web browser.
- The BMC management port on the services appliance is connected to the management network you plan to use.

#### SG100 BMC management port



#### SG1000 BMC management port



### Steps

1. Enter the URL for the BMC interface:

**`https://BMC_Port_IP`**

For *BMC\_Port\_IP*, use the DHCP or static IP address for the BMC management port.

The BMC sign-in page appears.

2. Enter the root username and password, using the password you set when you changed the default root password:

root

password



# NetApp®

root

.....|


Remember Username

Sign me in

[I forgot my password](#)

3. Click **Sign me in**

The BMC dashboard appears.



4. Optionally, create additional users by selecting **Settings > User Management** and clicking on any “disabled” user.



When users sign in for the first time, they might be prompted to change their password for increased security.

## Related information

[Changing the root password for the BMC interface](#)

## Configuring SNMP settings for the services appliance

If you are familiar with configuring SNMP for hardware, you can use the BMC interface to configure the SNMP settings for the services appliance. You can provide secure community strings, enable SNMP Trap, and specify up to five SNMP destinations.

### What you'll need

- You know how to access the BMC dashboard.
- You have experience in configuring SNMP settings for SNMPv1-v2c equipment.

### Steps

1. From the BMC dashboard, select **Settings > SNMP Settings**.
2. On the SNMP Settings page, select **Enable SNMP V1/V2**, and then provide a Read-Only Community String and a Read-Write Community String.

The Read-Only Community String is like a user ID or password. You should change this value to prevent intruders from getting information about your network setup. The Read-Write Community String protects the device against unauthorized changes.

3. Optionally, select **Enable Trap**, and enter the required information.



Enter the Destination IP for each SNMP trap using an IP address. Fully qualified domain names are not supported.

Enable traps if you want the services appliance to send immediate notifications to an SNMP console when it is in an unusual state. Traps might indicate link up/down conditions, temperatures exceeding certain thresholds, or high traffic.

4. Optionally, click **Send Test Trap** to test your settings.
5. If the settings are correct, click **Save**.

## Setting up email notifications for alerts


If you want email notifications to be sent when alerts occur, you must use the BMC interface to configure SMTP settings, users, LAN destinations, alert policies, and event filters.

### What you'll need

You know how to access the BMC dashboard.

### About this task

In the BMC interface, you use the **SMTP Settings**, **User Management**, and **Platform Event Filters** options on the Settings page to configure email notifications.



## Steps

1. Configure the SMTP settings.

- a. Select **Settings > SMTP Settings**.
- b. For Sender Email ID, enter a valid email address.

This email address is provided as the From address when the BMC sends email.

2. Set up users to receive alerts.

- a. From the BMC dashboard, select **Settings > User Management**.
- b. Add at least one user to receive alert notifications.

The email address you configure for a user is the address the BMC sends alert notifications to. For example, you could add a generic user, such as “notification-user,” and use the email address of a technical support team email distribution list.

3. Configure the LAN destination for alerts.

- a. Select **Settings > Platform Event Filters > LAN Destinations**.
- b. Configure at least one LAN destination.
  - Select **Email** as the Destination Type.
  - For BMC Username, select a user name that you added earlier.
  - If you added multiple users and want all of them to receive notification emails, you must add a LAN Destination for each user.

c. Send a test alert.

4. Configure alert policies so you can define when and where the BMC sends alerts.

- a. Select **Settings > Platform Event Filters > Alert Policies**.
- b. Configure at least one alert policy for each LAN destination.
  - For Policy Group Number, select **1**.
  - For Policy Action, select **Always send alert to this destination**.
  - For LAN Channel, select **1**.
  - In the Destination Selector, select the LAN destination for the policy.

5. Configure event filters to direct alerts for different event types to the appropriate users.

- a. Select **Settings > Platform Event Filters > Event Filters**.
- b. For Alert Policy Group Number, enter **1**.
- c. Create filters for every event you want the Alert Policy Group to be notified about.
  - You can create event filters for power actions, specific sensor events, or all events.
  - If you are uncertain which events to monitor, select **All Sensors** for Sensor Type and **All Events** for Event Options. If you receive unwanted notifications, you can change your selections later.

## Optional: Enabling node encryption

If you enable node encryption, the disks in your appliance can be protected by secure key management server (KMS) encryption against physical loss or removal from the site. You must select and enable node encryption during appliance installation and cannot unselect node encryption once the KMS encryption process starts.

### What you'll need

Review the information about KMS in the instructions for administering StorageGRID.

### About this task

An appliance that has node encryption enabled connects to the external key management server (KMS) that is configured for the StorageGRID site. Each KMS (or KMS cluster) manages the encryption keys for all appliance nodes at the site. These keys encrypt and decrypt the data on each disk in an appliance that has node encryption enabled.

A KMS can be set up in Grid Manager before or after the appliance is installed in StorageGRID. See the information about KMS and appliance configuration in the instructions for administering StorageGRID for additional details.

- If a KMS is set up before installing the appliance, KMS-controlled encryption begins when you enable node encryption on the appliance and add it to a StorageGRID site where KMS is configured.
- If a KMS is not set up before you install the appliance, KMS-controlled encryption is performed on each appliance that has node encryption enabled as soon as a KMS is configured and available for the site that contains the appliance node.



Any data that exists before an appliance that has node encryption enabled connects to the configured KMS is encrypted with a temporary key that is not secure. The appliance is not protected from removal or theft until the key is set to a value provided by the KMS.

Without the KMS key needed to decrypt the disk, data on the appliance cannot be retrieved and is effectively lost. This is the case whenever the decryption key cannot be retrieved from the KMS. The key becomes inaccessible if a customer clears the KMS configuration, a KMS key expires, connection to the KMS is lost, or the appliance is removed from the StorageGRID system where its KMS keys are installed.

### Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.



After the appliance has been encrypted with a KMS key, the appliance disks cannot be decrypted without using the same KMS key.

## 2. Select **Configure Hardware > Node Encryption**.

The screenshot shows the NetApp StorageGRID Appliance Installer interface. The top navigation bar includes 'Home', 'Configure Networking', 'Configure Hardware', 'Monitor Installation', and 'Advanced'. The 'Configure Hardware' section is active, showing the 'Node Encryption' configuration page. The page title is 'Node Encryption'. Below the title, there is a descriptive paragraph: 'Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.' Below this is the 'Encryption Status' section, which contains a yellow warning box: 'You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.' Underneath the warning box, there is a checkbox labeled 'Enable node encryption' which is checked. To the right of the checkbox is a blue 'Save' button. Below the checkbox and button is the 'Key Management Server Details' section, which is currently empty.

## 3. Select **Enable node encryption**.

You can unselect **Enable node encryption** without risk of data loss until you select **Save** and the appliance node accesses the KMS encryption keys in your StorageGRID system and begins disk encryption. You are not able to disable node encryption after the appliance is installed.



After you add an appliance that has node encryption enabled to a StorageGRID site that has a KMS, you cannot stop using KMS encryption for the node.

## 4. Select **Save**.

## 5. Deploy the appliance as a node in your StorageGRID system.

KMS-controlled encryption begins when the appliance accesses the KMS keys configured for your StorageGRID site. The installer displays progress messages during the KMS encryption process, which might take a few minutes depending on the number of disk volumes in the appliance.



Appliances are initially configured with a random non-KMS encryption key assigned to each disk volume. The disks are encrypted using this temporary encryption key, that is not secure, until the appliance that has node encryption enabled accesses the KMS keys configured for your StorageGRID site.

### After you finish

You can view node-encryption status, KMS details, and the certificates in use when the appliance node is in maintenance mode.

### Related information

[Administer StorageGRID](#)

[Monitoring node encryption in maintenance mode](#)

## Deploying a services appliance node

You can deploy a services appliance as a primary Admin Node, a non-primary Admin Node, or a Gateway Node. Both the SG100 and the SG1000 appliances can operate as Gateway Nodes and Admin Nodes (primary or non-primary) at the same time.

### Deploying a services appliance as a primary Admin Node

When you deploy a services appliance as a primary Admin Node, you use the StorageGRID Appliance Installer included on the appliance to install the StorageGRID software, or you upload the software version you want to install. You must install and configure the primary Admin Node before you install any other appliance node types. A primary Admin Node can connect to the Grid Network, and to the optional Admin Network and Client Network, if one or both are configured.

#### What you'll need

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.



If you have remapped any ports, you cannot use the same ports to configure load balancer endpoints. You can create endpoints using remapped ports, but those endpoints will be remapped to the original CLB ports and service, not the Load Balancer service. Follow the steps in the recovery and maintenance instructions for removing port remaps.



The CLB service is deprecated.

- You have a service laptop with a supported web browser.
- You know one of the IP addresses assigned to the appliance. You can use the IP address for any attached StorageGRID network.

#### About this task

To install StorageGRID on an appliance primary Admin Node:

- You use the StorageGRID Appliance Installer to install the StorageGRID software. If you want to install a different version of the software, you first upload it using the StorageGRID Appliance Installer.
- You wait as the software is installed.
- When the software has been installed, the appliance is rebooted automatically.

#### Steps

1. Open a browser, and enter the IP address for the appliance.  
**`https://services_appliance_IP:8443`**

The StorageGRID Appliance Installer Home page appears.

2. In the **This Node** section, select **Primary Admin**.
3. In the **Node name** field, enter the name you want to use for this appliance node, and click **Save**.

The node name is assigned to this appliance node in the StorageGRID system. It is shown on the Grid Nodes page in the Grid Manager.

4. Optionally, to install a different version of the StorageGRID software, follow these steps:
  - a. Download the installation archive from the NetApp Downloads page for StorageGRID.

[NetApp Downloads: StorageGRID](#)

- b. Extract the archive.
- c. From the StorageGRID Appliance Installer, select **Advanced** > **Upload StorageGRID Software**.
- d. Click **Remove** to remove the current software package.

NetApp® StorageGRID® Appliance Installer

Home | Configure Networking ▾ | Configure Hardware ▾ | Monitor Installation | **Advanced ▾**

### 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	11.3.0
Package Name	storagegrid-webscale-images-11-3-0_11.3.0-20190806.1731.4064510_amd64.deb

- e. Click **Browse** for the software package you downloaded and extracted, and then click **Browse** for the checksum file.

NetApp® StorageGRID® Appliance Installer

Home | Configure Networking ▾ | Configure Hardware ▾ | Monitor Installation | **Advanced ▾**

### 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	<input type="button" value="Browse"/>
Checksum File	<input type="button" value="Browse"/>

- f. Select **Home** to return to the Home page.

5. Confirm that the current state is “Ready to start installation of primary Admin Node name with software version x.y” and that the **Start Installation** button is enabled.



If you are deploying the Admin Node appliance as a node cloning target, stop the deployment process here and continue the node cloning procedure in recovery and maintenance.

#### Maintain & recover

6. From the StorageGRID Appliance Installer home page, click **Start Installation**.

#### Home

The installation is ready to be started. Review the settings below, and then click Start Installation.

#### This Node

Node type	Primary Admin (with Load Balancer)
Node name	xlr8r-8

Cancel Save

#### Installation

Current state	Ready to start installation of xlr8r-8 as primary Admin Node of a new grid running StorageGRID 11.3.0.
---------------	--

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.

#### Related information

[Deploying a services appliance as a Gateway or non-primary Admin Node](#)

#### Deploying a services appliance as a Gateway or non-primary Admin Node

When you deploy a services appliance as a Gateway Node or non-primary Admin Node, you use the StorageGRID Appliance Installer included on the appliance.

#### What you'll need

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.



If you have remapped any ports, you cannot use the same ports to configure load balancer endpoints. You can create endpoints using remapped ports, but those endpoints will be remapped to the original CLB ports and service, not the Load Balancer service. Follow the steps in the recovery and maintenance instructions for removing port remaps.



The CLB service is deprecated.

- The primary Admin Node for the StorageGRID system has been deployed.
- 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 a service laptop with a supported web browser.
- You know the IP address assigned to the appliance. You can use the IP address for any attached StorageGRID network.

### About this task

To install StorageGRID on a services appliance node:

- 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 the software is installed.

Partway through the appliance Gateway Node 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. The installation of a non-primary Admin Node does not require your approval.



Do not deploy the SG100 and SG1000 service appliances in the same site. Unpredictable performance might result.



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. You can also use the Appliance Installer to upload a JSON file that contains configuration information. See [Automating appliance installation and configuration](#).

### Steps

1. Open a browser, and enter the IP address for 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.

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	<ol style="list-style-type: none"> <li>a. Unselect the <b>Enable Admin Node discovery</b> check box.</li> <li>b. Enter the IP address manually.</li> <li>c. Click <b>Save</b>.</li> <li>d. Wait for the connection state for the new IP address to become ready.</li> </ol>
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"> <li>a. Select the <b>Enable Admin Node discovery</b> check box.</li> <li>b. Wait for the list of discovered IP addresses to be displayed.</li> <li>c. Select the primary Admin Node for the grid where this appliance Storage Node will be deployed.</li> <li>d. Click <b>Save</b>.</li> <li>e. Wait for the connection state for the new IP address to become ready.</li> </ol>

4. In the **Node name** field, enter the name you want to use for this appliance node, and click **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. Optionally, to install a different version of the StorageGRID software, follow these steps:
  - a. Download the installation archive from the NetApp Downloads page for StorageGRID.

[NetApp Downloads: StorageGRID](#)

- b. Extract the archive.
- c. From the StorageGRID Appliance Installer, select **Advanced > Upload StorageGRID Software**.
- d. Click **Remove** to remove the current software package.

NetApp® StorageGRID® Appliance Installer

Home | Configure Networking ▾ | Configure Hardware ▾ | Monitor Installation | Advanced ▾

### 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	11.3.0
Package Name	storagegrid-webscale-images-11-3-0_11.3.0-20190806.1731.4064510_amd64.deb

- e. Click **Browse** for the software package you downloaded and extracted, and then click **Browse** for the checksum file.

NetApp® StorageGRID® Appliance Installer

Home | Configure Networking ▾ | Configure Hardware ▾ | Monitor Installation | Advanced ▾

### 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	<input type="button" value="Browse"/>
Checksum File	<input type="button" value="Browse"/>

- f. Select **Home** to return to the Home page.
6. 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 installation and maintenance instructions for your appliance.

7. From the StorageGRID Appliance Installer home page, click **Start Installation**.

 The installation is ready to be started. Review the settings below, and then click Start Installation.

### This Node

Node type  

Node name

### Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

Connection state **Connection to 172.16.6.32 ready**

### Installation

Current state **Ready to start installation of GW-SG1000-003-074 into grid with Admin Node 172.16.6.32 running StorageGRID 11.3.0, using StorageGRID software downloaded from the Admin Node.**

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.

8. If your grid includes multiple appliance nodes, repeat the previous steps for each appliance.

#### Related information

[Deploying a services appliance as a primary Admin Node](#)

## Monitoring the services 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		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 are not re-run. If you are re-running an installation, any tasks that do not 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 in the appliance, 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 one of the following processes occurs:

- For all appliance nodes except the primary Admin Node, 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. Go to the next step.

- For appliance primary Admin Node installation, you do not need to approve the node. The appliance is rebooted. You can skip the next step.



During installation of an appliance primary Admin Node, a fifth phase appears (see the example screen shot showing four phases). If the fifth phase is in progress for more than 10 minutes, refresh the web 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	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...

```

4. Go to the Grid Manager, approve the pending grid node, and complete the StorageGRID installation process.

When you click **Install** from the Grid Manager, Stage 3 completes and stage 4, **Finalize Installation**, begins. When stage 4 completes, the appliance is rebooted.

## Automating appliance installation and configuration

You can automate the installation and configuration of your appliances and configuration of the whole StorageGRID system.

### About this task

Automating installation and configuration can be useful for deploying multiple StorageGRID instances or one large, complex StorageGRID instance.

To automate installation and configuration, use one or more of the following options:

- Create a JSON file that specifies the configuration settings for your appliances. Upload the JSON file using the StorageGRID Appliance Installer.



You can use the same file to configure more than one appliance.

- Use the `StorageGRIDconfigure-sga.py` Python script to automate the configuration of your appliances.
- Use additional Python scripts to configure other components of the whole StorageGRID system (the "grid").



You can use StorageGRID automation Python scripts directly, or you can use them as examples of how to use the StorageGRID Installation REST API in grid deployment and configuration tools you develop yourself. See the information about downloading and extracting the StorageGRID installation files in the Recovery and Maintenance instructions.

### Related information

[Maintain & recover](#)

### Automating appliance configuration using the StorageGRID Appliance Installer

You can automate the configuration of an appliance by using a JSON file that contains the configuration information. You upload the file using the StorageGRID Appliance Installer.

### What you'll need

- Your appliance must be on the latest firmware compatible with StorageGRID 11.5 or higher.
- You must be connected to the StorageGRID Appliance Installer on the appliance you are configuring using a supported browser.

### About this task

You can automate appliance configuration tasks such as configuring the following:

- Grid Network, Admin Network, and Client Network IP addresses
- BMC interface
- Network links
  - Port bond mode
  - Network bond mode
  - Link speed

Configuring your appliance using an uploaded JSON file is often more efficient than performing the

configuration manually using multiple pages in the StorageGRID Appliance Installer, especially if you have to configure many nodes. You must apply the configuration file for each node one at a time.



Experienced users who want to automate both the installation and configuration of their appliances can use the `configure-sga.py` script. [Automating installation and configuration of appliance nodes using the `configure-sga.py` script](#)

## Steps

1. Generate the JSON file using one of the following methods:

- The ConfigBuilder application

[ConfigBuilder.netapp.com](#)

- The `configure-sga.py` appliance configuration script. You can download the script from StorageGRID Appliance Installer (**Help > Appliance Configuration Script**). See the instructions on automating the configuration using the `configure-sga.py` script.

[Automating installation and configuration of appliance nodes using the `configure-sga.py` script](#)

The node names in the JSON file must follow these requirements:

- Must be a valid hostname containing at least 1 and no more than 32 characters
- Can use letters, numbers, and hyphens are allowed
- Cannot start or end with a hyphen or contain only numbers



Ensure that the node names (the top-level names) in the JSON file are unique, or you will not be able to configure more than one node using the JSON file.

2. Select **Advanced > Update Appliance Configuration**.

The Update Appliance Configuration page appears.

### Update Appliance Configuration

Use a JSON file to update this appliance's configuration. You can generate the JSON file from the [ConfigBuilder](#) application or from the [appliance configuration script](#).

**⚠** You might lose your connection if the applied configuration from the JSON file includes "link\_config" and/or "networks" sections. If you are not reconnected within 1 minute, re-enter the URL using one of the other IP addresses assigned to the appliance.

### Upload JSON

JSON  
configuration

Browse

Node name

-- Upload a file ▾

Apply JSON configuration

3. Select the JSON file with the configuration you want to upload.
  - a. Select **Browse**.
  - b. Locate and select the file.
  - c. Select **Open**.

The file is uploaded and validated. When the validation process is complete, the file name is shown next to a green check mark.



You might lose connection to the appliance if the configuration from the JSON file includes sections for "link\_config", "networks", or both. If you are not reconnected within 1 minute, re-enter the appliance URL using one of the other IP addresses assigned to the appliance.

#### Upload JSON

JSON configuration	<input type="button" value="Browse"/>	✓ appliances.orig.json
Node name	-- Select a node ▼	
<input type="button" value="Apply JSON configuration"/>		

The **Node name** drop down is populated with the top-level node names defined in the JSON file.



If the file is not valid, the file name is shown in red and an error message is displayed in a yellow banner. The invalid file is not applied to the appliance. You can use ConfigBuilder to ensure you have a valid JSON file.

4. Select a node from the list in the **Node name** drop down.

The **Apply JSON configuration** button is enabled.

#### Upload JSON

JSON configuration	<input type="button" value="Browse"/>	✓ appliances.orig.json
Node name	Lab-80-1000 ▼	
<input type="button" value="Apply JSON configuration"/>		

5. Select **Apply JSON configuration**.

The configuration is applied to the selected node.

## Automating installation and configuration of appliance nodes using the `configure-sga.py` script

You can use the `configure-sga.py` script to automate many of the installation and configuration tasks for StorageGRID appliance nodes, including installing and configuring a primary Admin Node. This script can be useful if you have a large number of appliances to configure. You can also use the script to generate a JSON file that contains appliance configuration information.

### What you'll need

- 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 primary Admin Node using the StorageGRID Appliance Installer.
- If you are installing the primary Admin Node, you know its IP address.
- If you are installing and configuring other nodes, the primary Admin Node has been deployed, and you know its IP address.
- For all nodes other than the primary Admin Node, 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 downloaded the `configure-sga.py` file. The file is included in the installation archive, or you can access it by clicking **Help > Appliance Installation Script** in the StorageGRID Appliance Installer.



This procedure is for advanced users with some experience using command-line interfaces. Alternatively, you can also use the StorageGRID Appliance Installer to automate the configuration.

[Automating appliance configuration using the StorageGRID Appliance Installer](#)

### Steps

1. Log in to the Linux machine you are using to run the Python script.
2. For general help with the script syntax and to see a list of the available parameters, enter the following:

```
configure-sga.py --help
```

The `configure-sga.py` script uses five subcommands:

- `advanced` for advanced StorageGRID appliance interactions, including BMC configuration and creating a JSON file containing the current configuration of the appliance
- `configure` for configuring the RAID mode, node name, and networking parameters
- `install` for starting a StorageGRID installation
- `monitor` for monitoring a StorageGRID installation
- `reboot` for rebooting the appliance

If you enter a subcommand (`advanced`, `configure`, `install`, `monitor`, or `reboot`) argument followed by the `--help` option you will get a different help text providing more detail on the options available within that subcommand:

```
configure-sga.py subcommand --help
```

3. To confirm the current configuration of the appliance node, enter the following where *SGA-install-ip* is any one of the IP addresses for the appliance node:

```
configure-sga.py configure SGA-INSTALL-IP
```

The results show current IP information for the appliance, including the IP address of the primary Admin Node and information about the Admin, Grid, and Client Networks.

```
Connecting to +https://10.224.2.30:8443+ (Checking version and
connectivity.)
2021/02/25 16:25:11: Performing GET on /api/versions... Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/system-info... Received
200
2021/02/25 16:25:11: Performing GET on /api/v2/admin-connection...
Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/link-config... Received
200
2021/02/25 16:25:11: Performing GET on /api/v2/networks... Received 200
2021/02/25 16:25:11: Performing GET on /api/v2/system-config... Received
200
```

#### StorageGRID Appliance

```
Name:          LAB-SGA-2-30
Node type:     storage
```

#### StorageGRID primary Admin Node

```
IP:           172.16.1.170
State:        unknown
Message:      Initializing...
Version:      Unknown
```

#### Network Link Configuration

##### Link Status

Link	State	Speed (Gbps)
----	-----	-----
1	Up	10
2	Up	10
3	Up	10
4	Up	10
5	Up	1
6	Down	N/A

##### Link Settings

```
Port bond mode:  FIXED
Link speed:      10GBE
```

```
Grid Network:    ENABLED
Bonding mode:    active-backup
```

```

VLAN:          novlan
MAC Addresses: 00:a0:98:59:8e:8a 00:a0:98:59:8e:82

Admin Network: ENABLED
Bonding mode:  no-bond
MAC Addresses: 00:80:e5:29:70:f4

Client Network: ENABLED
Bonding mode:  active-backup
VLAN:          novlan
MAC Addresses: 00:a0:98:59:8e:89 00:a0:98:59:8e:81

```

Grid Network

```

CIDR:          172.16.2.30/21 (Static)
MAC:           00:A0:98:59:8E:8A
Gateway:       172.16.0.1
Subnets:      172.17.0.0/21
                172.18.0.0/21
                192.168.0.0/21
MTU:           1500

```

Admin Network

```

CIDR:          10.224.2.30/21 (Static)
MAC:           00:80:E5:29:70:F4
Gateway:       10.224.0.1
Subnets:      10.0.0.0/8
                172.19.0.0/16
                172.21.0.0/16
MTU:           1500

```

Client Network

```

CIDR:          47.47.2.30/21 (Static)
MAC:           00:A0:98:59:8E:89
Gateway:       47.47.0.1
MTU:           2000

```

```

#####
##### If you are satisfied with this configuration, #####
##### execute the script with the "install" sub-command. #####
#####

```

4. If you need to change any of the values in the current configuration, use the `configure` subcommand to update them. For example, if you want to change the IP address that the appliance uses for connection to the primary Admin Node to `172.16.2.99`, enter the following:  
`configure-sga.py configure --admin-ip 172.16.2.99 SGA-INSTALL-IP`
5. If you want to back up the appliance configuration to a JSON file, use the `advanced` and `backup-file`

subcommands. For example, if you want to back up the configuration of an appliance with IP address `SGA-INSTALL-IP` to a file named `appliance-SG1000.json`, enter the following:

```
configure-sga.py advanced --backup-file appliance-SG1000.json SGA-INSTALL-IP
```

The JSON file containing the configuration information is written to the same directory you executed the script from.



Check that the top-level node name in the generated JSON file matches the appliance name. Do not make any changes to this file unless you are an experienced user and have a thorough understanding of StorageGRID APIs.

6. When you are satisfied with the appliance configuration, use the `install` and `monitor` subcommands to install the appliance:

```
configure-sga.py install --monitor SGA-INSTALL-IP
```

7. If you want to reboot the appliance, enter the following:

```
configure-sga.py reboot SGA-INSTALL-IP
```

### Automating the configuration of StorageGRID

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

#### What you'll need

- 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>	Sample 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 sample 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
```

### After you finish

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, you need to extract 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.
```

### Overview of installation REST APIs

StorageGRID provides two REST APIs for performing installation tasks: the StorageGRID Installation API and the StorageGRID Appliance Installer API.

Both APIs use 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 (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs 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 in the event that 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.

## StorageGRID Appliance Installer API

The StorageGRID Appliance Installer API can be accessed over HTTPS from *Controller\_IP:8443*.

To access the API documentation, go to the StorageGRID Appliance Installer on the appliance and select **Help > API Docs** from the menu bar.

The StorageGRID Appliance Installer API includes the following sections:

- **clone** — Operations to configure and control node cloning.
- **encryption** — Operations to manage encryption and view encryption status.
- **hardware configuration** — Operations to configure system settings on attached hardware.
- **installation** — Operations for starting the appliance installation and for monitoring installation status.
- **networking** — Operations related to the Grid, Admin, and Client Network configuration for a StorageGRID appliance and appliance port settings.
- **setup** — Operations to help with initial appliance installation setup including requests to get information about the system and update the primary Admin Node IP.
- **support** — Operations for rebooting the controller and getting logs.
- **upgrade** — Operations related to upgrading appliance firmware.
- **uploadsg** — Operations for uploading StorageGRID installation files.

## Troubleshooting the hardware installation

If you encounter issues during the installation, you might find it helpful to review

troubleshooting information related to hardware setup and connectivity issues.

### Related information

[Hardware setup appears to hang](#)

[Troubleshooting connection issues](#)

### Viewing boot-up codes for the appliance

When you apply power to the appliance, the BMC logs a series of boot-up codes. You can view these codes on a graphical console that is connected to the BMC management port.

### What you'll need

- You know how to access the BMC dashboard.
- If you want to use a kernel-based virtual machine (KVM), you have experience deploying and using KVM applications.
- If you want to use serial-over-LAN (SOL), you have experience using IPMI SOL console applications.

### Steps

1. Select one of the following methods for viewing the boot-up codes for the appliance controller, and gather the required equipment.

Method	Required equipment
VGA console	<ul style="list-style-type: none"><li>• VGA-capable monitor</li><li>• VGA cable</li></ul>
KVM	<ul style="list-style-type: none"><li>• KVM application</li><li>• RJ-45 cable</li></ul>
Serial port	<ul style="list-style-type: none"><li>• DB-9 serial cable</li><li>• Virtual serial terminal</li></ul>
SOL	<ul style="list-style-type: none"><li>• Virtual serial terminal</li></ul>

2. If you are using a VGA console, perform these steps:
  - a. Connect a VGA-capable monitor to the VGA port on the back of the appliance.
  - b. View the codes displayed on the monitor.
3. If you are using BMC KVM, perform these steps:
  - a. Connect to the BMC management port and log into the BMC web interface.
  - b. Select **Remote Control**.
  - c. Launch the KVM.
  - d. View the codes on the virtual monitor.
4. If you are using a serial port and terminal, perform these steps:

- a. Connect to the DB-9 serial port on the back of the appliance.
  - b. Use settings 115200 8-N-1.
  - c. View the codes printed over the serial terminal.
5. If you are using SOL, perform these steps:
- a. Connect to the IPMI SOL using the BMC IP address and login credentials.

```
ipmitool -I lanplus -H 10.224.3.91 -U root -P calvin sol activate
```

- b. View the codes on the virtual serial terminal.
6. Use the table to look up the codes for your appliance.

Code	Indicates
HI	The master boot script has started.
HP	The system is checking to see if the network interface card (NIC) firmware needs to be updated.
RB	The system is rebooting after applying firmware updates.
FP	The hardware subsystem firmware update checks have been completed. Inter-controller communication services are starting.
HC	The system is checking for existing StorageGRID installation data.
HO	The StorageGRID appliance is running.
HA	StorageGRID is running.

## Related information

[Accessing the BMC interface](#)

## Viewing error codes for the appliance

If a hardware error occurs when the appliance is booting up, the BMC logs an error code. As required, you can view these error codes using the BMC interface, and then work with technical support to resolve the issue.

## What you'll need

- You know how to access the BMC dashboard.

## Steps

1. From the BMC dashboard, select **BIOS POST Code**.

2. Review the information displayed for Current Code and the Previous Code.

If any of the following error codes are shown, work with technical support to resolve the issue.

<b>Code</b>	<b>Indicates</b>
0x0E	Microcode not found
0x0F	Microcode not loaded
0x50	Memory initialization error. Invalid memory type or incompatible memory speed.
0x51	Memory initialization error. SPD reading has failed.
0x52	Memory initialization error. Invalid memory size or memory modules do not match.
0x53	Memory initialization error. No usable memory detected.
0x54	Unspecified memory initialization error
0x55	Memory not installed
0x56	Invalid CPU type or speed
0x57	CPU mismatch
0x58	CPU self-test failed, or possible CPU cache error
0x59	CPU micro-code is not found, or micro-code update failed
0x5A	Internal CPU error
0x5B	Reset PPI is not available
0x5C	PEI phase BMC self-test failure
0xD0	CPU initialization error
0xD1	North bridge initialization error
0xD2	South bridge initialization error

<b>Code</b>	<b>Indicates</b>
0xD3	Some architectural protocols are not available
0xD4	PCI resource allocation error. Out of resources.
0xD5	No space for legacy option ROM
0xD6	No console output devices are found
0xD7	No console input devices are found
0xD8	Invalid password
0xD9	Error loading boot option (LoadImage returned error)
0xDA	Boot option failed (StartImage returned error)
0xDB	Flash update failed
0xDC	Reset protocol is not available
0xDD	DXE phase BMC self-test failure
0xE8	MRC: ERR_NO_MEMORY
0xE9	MRC: ERR_LT_LOCK
0xEA	MRC: ERR_DDR_INIT
0xEB	MRC: ERR_MEM_TEST
0xEC	MRC: ERR_VENDOR_SPECIFIC
0xED	MRC: ERR_DIMM_COMPAT
0xEE	MRC: ERR_MRC_COMPATIBILITY
0xEF	MRC: ERR_MRC_STRUCT
0xF0	MRC: ERR_SET_VDD
0xF1	MRC: ERR_IOT_MEM_BUFFER

<b>Code</b>	<b>Indicates</b>
0xF2	MRC: ERR_RC_INTERNAL
0xF3	MRC: ERR_INVALID_REG_ACCESS
0xF4	MRC: ERR_SET_MC_FREQ
0xF5	MRC: ERR_READ_MC_FREQ
0x70	MRC: ERR_DIMM_CHANNEL
0x74	MRC: ERR_BIST_CHECK
0xF6	MRC: ERR_SMBUS
0xF7	MRC: ERR_PCU
0xF8	MRC: ERR_NGN
0xF9	MRC: ERR_INTERLEAVE_FAILURE

### **Hardware setup appears to hang**

The StorageGRID Appliance Installer might not be available if hardware faults or cabling errors prevent the appliance from completing its boot-up processing.

#### **Steps**

1. Review the LEDs on the appliance and the boot-up and error codes displayed in the BMC.
2. If you need help resolving an issue, contact technical support.

#### **Related information**

[Viewing boot-up codes for the appliance](#)

[Viewing error codes for the appliance](#)

### **Troubleshooting connection issues**

If you encounter connection issues during the StorageGRID appliance installation, you should perform the corrective action steps listed.

#### **Unable to connect to the appliance**

If you cannot connect to the services appliance, there might be a network issue, or the hardware installation might not have been completed successfully.

#### **Steps**

1. Try to ping the appliance using the IP address for the appliance :  
`ping services_appliance_IP`
2. If you receive no response from the ping, confirm you are using the correct IP address.

You can use the IP address of the appliance on the Grid Network, the Admin Network, or the Client Network.

3. If the IP address is correct, check appliance cabling, QSFP or SFP transceivers, and the network setup.

If that does not resolve the issue, contact technical support.

4. If the ping was successful, open a web browser.
5. Enter the URL for the StorageGRID Appliance Installer:  
`https://appliances_controller_IP:8443`

The Home page appears.

### Rebooting the services appliance while the StorageGRID Appliance Installer is running

You might need to reboot the services appliance while the StorageGRID Appliance Installer is running. For example, you might need to reboot the services appliance if the installation fails.

#### About this task


This procedure only applies when the services appliance is running the StorageGRID Appliance Installer. Once the installation is completed, this step no longer works because the StorageGRID Appliance Installer is no longer available.

#### Steps

1. From the menu bar of the StorageGRID Appliance Installer, click **Advanced > Reboot Controller**.

The Reboot Controller page appears.

2. From the StorageGRID Appliance Installer, click **Advanced > Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
  - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



The services appliance is rebooted.

## Maintaining the appliance

You might need to perform maintenance procedures on the appliance. The procedures in this section assume that the appliance has already been deployed as a Gateway Node or an Admin Node in a StorageGRID system.

### Steps

- [Placing an appliance into maintenance mode](#)
- [Turning the controller identify LED on and off](#)
- [Locating the controller in a data center](#)
- [Replacing the services appliance](#)
- [Replacing a power supply in the services appliance](#)
- [Replacing a fan in the services appliance](#)
- [Replacing a drive in the services appliance](#)
- [Changing the link configuration of the services appliance](#)
- [Changing the MTU setting](#)
- [Checking the DNS server configuration](#)
- [Monitoring node encryption in maintenance mode](#)

### Placing an appliance into maintenance mode

You must place the appliance into maintenance mode before performing specific maintenance procedures.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission. For details, see the instructions for administering StorageGRID.

### About this task

Placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.



The password and host key for a StorageGRID appliance in maintenance mode remain the same as they were when the appliance was in service.

### Steps

1. From the Grid Manager, select **Nodes**.
2. From the tree view of the Nodes page, select the appliance Storage Node.
3. Select **Tasks**.

Overview Hardware Network Storage Objects ILM Events **Tasks**

### Reboot

Shuts down and restarts the node.

Reboot

### Maintenance Mode

Places the appliance's compute controller into maintenance mode.

Maintenance Mode

4. Select **Maintenance Mode**.

A confirmation dialog box appears.

**⚠ Enter Maintenance Mode on SGA-106-15**

You must place the appliance's compute controller into maintenance mode to perform certain maintenance procedures on the appliance.

Attention: All StorageGRID services on this node will be shut down. Wait a few minutes for the node to reboot into maintenance mode.

If you are ready to start, enter the provisioning passphrase and click OK.

Provisioning Passphrase

Cancel OK

5. Enter the provisioning passphrase, and select **OK**.

A progress bar and a series of messages, including "Request Sent," "Stopping StorageGRID," and "Rebooting," indicate that the appliance is completing the steps for entering maintenance mode.


## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

**Attention:** Your request has been sent, but the appliance might take 10-15 minutes to enter maintenance mode. Do not perform maintenance procedures until this tab indicates maintenance mode is ready, or data could become corrupted.

 Request Sent

When the appliance is in maintenance mode, a confirmation message lists the URLs you can use to access the StorageGRID Appliance Installer.

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

This node is currently in maintenance mode. Navigate to one of the URLs listed below and perform any necessary maintenance procedures.

- <https://172.16.2.106:8443>
- <https://10.224.2.106:8443>
- <https://47.47.2.106:8443>
- <https://169.254.0.1:8443>

When you are done with any required maintenance procedures, you must exit maintenance mode by clicking Reboot Controller from the StorageGRID Appliance Installer.

6. To access the StorageGRID Appliance Installer, browse to any of the URLs displayed.

If possible, use the URL containing the IP address of the appliance's Admin Network port.




Accessing <https://169.254.0.1:8443> requires a direct connection to the local management port.


7. From the StorageGRID Appliance Installer, confirm that the appliance is in maintenance mode.

⚠ This node is in maintenance mode. Perform any required maintenance procedures. If you want to exit maintenance mode manually to resume normal operation, go to **Advanced > Reboot Controller** to **reboot** the controller.

8. Perform any required maintenance tasks.
9. After completing maintenance tasks, exit maintenance mode and resume normal node operation. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



### Turning the controller identify LED on and off


The blue identify LED on the front and back of the controller can be turned on to help locate the appliance in a data center.

#### What you'll need

You must have the BMC IP address of the controller you want to identify.

#### Steps

1. Access the controller BMC interface.
2. Select **Server Identify**.
3. Select **ON** and then select **Perform Action**.



## Result

The blue identify LEDs light on the front (shown) and rear of the controller.



If a bezel is installed on the controller, it might be difficult to see the front identify LED.

## After you finish

To turn off the controller identify LED:

- Press the identify LED switch on the controller front panel.
- From the controller BMC interface, select **Server Identify**, select **OFF** and then select **Perform Action**.

The blue identify LEDs on the front and rear of the controller go off.



### Related information

[Locating the controller in a data center](#)

[Accessing the BMC interface](#)

### Locating the controller in a data center

Locate the controller so that you can perform hardware maintenance or upgrades.

### What you'll need

- You have determined which controller requires maintenance.

(Optional) To help locate the controller in your data center, turn on the blue identify LED.

[Turning the controller identify LED on and off](#)

### Steps

1. Find the controller requiring maintenance in the data center.
  - Look for a lit blue identify LED on the front or rear of the controller.

The front identify LED is behind the controller front bezel and might be difficult to see if the bezel is installed.



- Check the tags attached to the front of each controller for a matching part number.
2. Remove the controller front bezel, if one is installed, to access the front panel controls and indicators.

3. Optional: Turn off the blue identify LED if you used it to locate the controller.

- Press the identify LED switch on the controller front panel.
- Use the controller BMC interface.

#### Turning the controller identify LED on and off

## Replacing the services appliance

You might need to replace the appliance if it is not functioning optimally or if it has failed.

### What you'll need

- You have a replacement appliance with the same part number as the appliance you are replacing.
- You have labels to identify each cable that is connected to the appliance.
- You have physically located the appliance that you are replacing in the data center. See [Locating the controller in a data center](#).
- The appliance has been placed maintenance mode. See [Placing an appliance into maintenance mode](#).

### About this task

The StorageGRID node will not be accessible while you replace the appliance. If the appliance is functioning sufficiently, you can perform a controlled shutdown at the start of this procedure.



If you are replacing the appliance before installing StorageGRID software, you might not be able to access the StorageGRID Appliance Installer immediately after completing this procedure. While you can access the StorageGRID Appliance Installer from other hosts on the same subnet as the appliance, you cannot access it from hosts on other subnets. This condition should resolve itself within 15 minutes (when any ARP cache entries for the original appliance time out), or you can clear the condition immediately by purging any old ARP cache entries manually from the local router or gateway.

### Steps

1. When the appliance has been placed maintenance mode, shut down the appliance.

a. Log in to the grid node:

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

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

b. Shut down the appliance:

**shutdown -h now**

2. Use one of two methods to verify that the power for the appliance is off:

- The power indicator LED on the front of the appliance is off.
- The Power Control page of the BMC interface indicates that the appliance is off.

3. If the StorageGRID networks attached to the appliance use DHCP servers, update DNS/network and IP

address settings.

- a. Locate the MAC address label on the front of the appliance, and determine the MAC address for the Admin Network port.



The MAC address label lists the MAC address for the BMC management port.

To determine the MAC address for the Admin Network port, you must add **2** to the hexadecimal number on the label. For example, if the MAC address on the label ends in **09**, the MAC address for the Admin Port would end in **0B**. If the MAC address on the label ends in **(y)FF**, the MAC address for the Admin Port would end in **(y+1)01**. You can easily make this calculation by opening Calculator in Windows, setting it to Programmer mode, selecting Hex, typing the MAC address, then typing **+ 2 =**.

- b. Ask your network administrator to associate the DNS/network and IP address for the appliance you removed with the MAC address for the replacement appliance.



You must ensure that all IP addresses for the original appliance have been updated before you apply power to the replacement appliance. Otherwise, the appliance will obtain new DHCP IP addresses when it boots up and might not be able to reconnect to StorageGRID. This step applies to all StorageGRID networks that are attached to the appliance.



If the original appliance used static IP address, the new appliance will automatically adopt the IP addresses of the appliance you removed.

#### 4. Remove and replace the appliance:

- a. Label the cables and then disconnect the cables and any network transceivers.



To prevent degraded performance, do not twist, fold, pinch, or step on the cables.

- b. Remove the failed appliance from the cabinet or rack.
- c. Transfer the two power supplies, eight cooling fans, and two SSDs from the failed appliance to the replacement appliance.

Follow the instructions provided for replacing these components.

- d. Install the replacement appliance into the cabinet or rack.
- e. Replace the cables and any optical transceivers.
- f. Power on the appliance and monitor the appliance LEDs and boot-up codes.

Use the BMC interface to monitor boot-up status.

5. Confirm that the appliance node appears in the Grid Manager and that no alerts appear.

#### Related information

[Installing the appliance into a cabinet or rack \(SG100 and SG1000\)](#)

[Viewing status indicators on the SG100 and SG1000 appliances](#)

[Viewing boot-up codes for the appliance](#)

## Replacing a power supply in the services appliance

The services appliance has two power supplies for redundancy. If one of the power supplies fails, you must replace it as soon as possible to ensure that the appliance has redundant power.

### What you'll need

- You have unpacked the replacement power supply unit.
- You have physically located the appliance where you are replacing the power supply in the data center.

#### [Locating the controller in a data center](#)

- You can confirmed that the other power supply is installed and running.

### About this task


The figure shows the two power supply units for the SG100, which are accessible from the back of the appliance.




The power supplies for the SG1000 are identical.

### Steps

1. Unplug the power cord from the power supply.
2. Lift the cam handle.




3. Press the blue latch and pull the power supply out.



4. Slide the replacement power supply into the chassis.

Ensure that the blue latch is on the right side when you slide the unit in.



5. Push the cam handle down to secure the power supply.
6. Attach the power cord to the power supply, and ensure that the green LED comes on.

### Replacing a fan in the services appliance

The services appliance has eight cooling fans. If one of the fans fails, you must replace it as soon as possible to ensure that the appliance has proper cooling.

#### What you'll need

- You have unpacked the replacement fan.
- You have physically located the appliance where you are replacing the fan in the data center.

#### [Locating the controller in a data center](#)

- You have confirmed that the other fans are installed and running.

- The appliance has been placed maintenance mode.

### Placing an appliance into maintenance mode

#### About this task

The appliance node will not be accessible while you replace the fan.

The photograph shows a fan for the services appliance. The cooling fans are accessible after you take the top cover off of the appliance.



Each of the two power supply units also contain a fan. Those fans are not included in this procedure.




#### Steps


1. When the appliance has been placed maintenance mode, shut down the appliance.
  - 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. Shut down the services appliance:  
**`shutdown -h now`**
2. Use one of two methods to verify that the power for the services appliance is off:
    - The power indicator LED on the front of the appliance is off.
    - The Power Control page of the BMC interface indicates that the appliance is off.
  3. Lift the latch on the top cover and remove the cover from the appliance.
  4. Locate the fan that failed.




5. Lift the failed fan out of the chassis.




6. Slide the replacement fan into the open slot in the chassis.

Line up the edge of the fan with the guide pin. The pin is circled in the photograph.



7. Press the fan's connector firmly into the circuit board.



8. Put the top cover back on the appliance, and press the latch down to secure the cover in place.
9. Power on the appliance and monitor the controller LEDs and boot-up codes.

Use the BMC interface to monitor boot-up status.

10. Confirm that the appliance node appears in the Grid Manager and that no alerts appear.

### Replacing a drive in the services appliance

The SSDs in the services appliance contain the StorageGRID operating system. Additionally, when the appliance is configured as an Admin Node, the SSDs also contain audit logs, metrics, and database tables. The drives are mirrored using RAID1 for redundancy. If one of the drives fails, you must replace it as soon as possible to ensure redundancy.

#### What you'll need

- You have physically located the appliance where you are replacing the drive in the data center.

#### [Locating the controller in a data center](#)

- You have verified which drive has failed by noting that its left LED is blinking amber.



If you remove the working drive, you will bring down the appliance node. See the information about viewing status indicators to verify the failure.

- You have obtained the replacement drive.
- You have obtained proper ESD protection.

#### Steps


1. Verify that the drive's left LED is blinking amber.

You can also use the Grid Manager to monitor the status of the SSDs. Select **Nodes**. Then select **Appliance Node > Hardware**. If a drive has failed, the Storage RAID Mode field contains a message about which drive has failed.

2. Wrap the strap end of the ESD wristband around your wrist, and secure the clip end to a metal ground to prevent static discharge.
3. Unpack the replacement drive, and set it on a static-free, level surface near the appliance.

Save all packing materials.


4. Press the release button on the failed drive.



The handle on the drive springs open partially, and the drive releases from the slot.

5. Open the handle, slide the drive out, and place it on a static-free, level surface.
6. Press the release button on the replacement drive before you insert it into the drive slot.

The latch springs open.



7. Insert the replacement drive in the slot, and then close the drive handle.



Do not use excessive force while closing the handle.

When the drive is fully inserted, you hear a click.

The drive is automatically rebuilt with mirrored data from the working drive. You can check the status of the rebuild by using the Grid Manager. Select **Nodes**. Then select **Appliance Node > Hardware**. The Storage RAID Mode field contains a “rebuilding” message until the drive is completely rebuilt.

8. Contact technical support about the drive replacement.

Technical support will provide instructions for returning the failed drive.

## Changing the link configuration of the services appliance

You can change the Ethernet link configuration of the services appliance. You can change the port bond mode, the network bond mode, and the link speed.

### What you'll need

- You must place the appliance in maintenance mode. Putting a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.

[Placing an appliance into maintenance mode](#)


## About this task

Options for changing the Ethernet link configuration of the services appliance include:

- Changing **Port bond mode** from Fixed to Aggregate, or from Aggregate to Fixed
- Changing **Network bond mode** from Active-Backup to LACP, or from LACP to Active-Backup
- Enabling or disabling VLAN tagging, or changing the value of a VLAN tag
- Changing the link speed

## Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > Link Configuration**.



2. Make the desired changes to the link configuration.

For more information on the options, see “Configuring network links.”

3. When you are satisfied with your selections, click **Save**.



You might lose your connection if you made changes to the network or link you are connected through. If you are not reconnected within 1 minute, re-enter the URL for the StorageGRID Appliance Installer using one of the other IP addresses assigned to the appliance:

**`https://services_appliance_IP:8443`**

4. Make any necessary changes to the IP addresses for the appliance.

If you made changes to the VLAN settings, the subnet for the appliance might have changed. If you need to change the IP addresses for the appliance, follow the instructions for configuring IP addresses.

### [Configuring StorageGRID IP addresses](#)


5. Select **Configure Networking > Ping Test** from the menu.
6. Use the Ping Test tool to check connectivity to IP addresses on any networks that might have been affected by the link configuration changes you made when configuring the appliance.

In addition to any other tests you choose to perform, confirm that you can ping the Grid Network IP address of the primary Admin Node, and the Grid Network IP address of at least one other node. If necessary, return to the instructions for configuring network links, and correct any issues.


7. Once you are satisfied that your link configuration changes are working, reboot the node. From the

StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Changing the MTU setting

You can change the MTU setting that you assigned when you configured IP addresses for the appliance node.

### What you'll need

The appliance has been placed maintenance mode.

### Placing an appliance into maintenance mode

#### Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > IP Configuration**.
2. Make the desired changes to the MTU settings for the Grid Network, Admin Network, and Client Network.

### Grid Network

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

IP Assignment  Static  DHCP

IPv4 Address (CIDR)

Gateway

**⚠** All required Grid Network subnets must also be defined in the Grid Network Subnet List on the Primary Admin Node before starting installation.

Subnets (CIDR)  **×**

**×**

**+ ×**

MTU  **⬆**



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.



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


3. When you are satisfied with the settings, select **Save**.

4. Reboot the node. From the StorageGRID Appliance Installer, select **Advanced** > **Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Administer StorageGRID](#)

## Checking the DNS server configuration

You can check and temporarily change the domain name system (DNS) servers that are currently in use by this appliance node.

## What you'll need

The appliance has been placed maintenance mode.

## Placing an appliance into maintenance mode

### About this task

You might need to change the DNS server settings if an encrypted appliance cannot connect to the key management server (KMS) or KMS cluster because the hostname for the KMS was specified as a domain name instead of an IP address. Any changes that you make to the DNS settings for the appliance are temporary and are lost when you exit maintenance mode. To make these changes permanent, specify the DNS servers in Grid Manager (**Maintenance > Network > DNS Servers**).

- Temporary changes to the DNS configuration are necessary only for node-encrypted appliances where the KMS server is defined using a fully qualified domain name, instead of an IP address, for the hostname.
- When a node-encrypted appliance connects to a KMS using a domain name, it must connect to one of the DNS servers defined for the grid. One of these DNS servers then translates the domain name into an IP address.
- If the node cannot reach a DNS server for the grid, or if you changed the grid-wide DNS settings when a node-encrypted appliance node was offline, the node is unable to connect to the KMS. Encrypted data on the appliance cannot be decrypted until the DNS issue is resolved.

To resolve a DNS issue preventing KMS connection, specify the IP address of one or more DNS servers in the StorageGRID Appliance Installer. These temporary DNS settings allow the appliance to connect to the KMS and decrypt data on the node.

For example, if the DNS server for the grid changes while an encrypted node was offline, the node will not be able to reach the KMS when it comes back online, since it is still using the previous DNS values. Entering the new DNS server IP address in the StorageGRID Appliance Installer allows a temporary KMS connection to decrypt the node data.

### Steps

1. From the StorageGRID Appliance Installer, select **Configure Networking > DNS Configuration**.
2. Verify that the DNS servers specified are correct.

#### DNS Servers

 Configuration changes made on this page will not be passed to the StorageGRID software after appliance installation.

#### Servers

Server 1	<input type="text" value="10.224.223.135"/>	
Server 2	<input type="text" value="10.224.223.136"/>	 
<input type="button" value="Cancel"/>		<input type="button" value="Save"/>

3. If required, change the DNS servers.



Changes made to the DNS settings are temporary and are lost when you exit maintenance mode.

4. When you are satisfied with the temporary DNS settings, select **Save**.


The node uses the DNS server settings specified on this page to reconnect to the KMS, allowing data on the node to be decrypted.

5. After node data is decrypted, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:
  - Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
  - Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



When the node reboots and rejoins the grid, it uses the system-wide DNS servers listed in the Grid Manager. After rejoining the grid, the appliance will no longer use the temporary DNS servers specified in the StorageGRID Appliance Installer while the appliance was in maintenance mode.

It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



### Monitoring node encryption in maintenance mode

If you enabled node encryption for the appliance during installation, you can monitor the

node-encryption status of each appliance node, including the node-encryption state and key management server (KMS) details.

### What you'll need

- Node encryption must have been enabled for the appliance during installation. You cannot enable node encryption after the appliance is installed.
- The appliance has been placed into maintenance mode.

### Placing an appliance into maintenance mode

### Steps

1. From the StorageGRID Appliance Installer, select **Configure Hardware > Node Encryption**.

#### Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

#### Encryption Status

 You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

Save

#### Key Management Server Details

View the status and configuration details for the KMS that manages the encryption key for this appliance. You must use the Grid Manager to make configuration changes.

KMS display name	thales
External key UID	41b0306abcce451facfe01b1b4870ae1c1ec6bd5e3849d790223766baf35c57
Hostnames	10.96.99.164 10.96.99.165
Port	5696

Server certificate



Client certificate



#### Clear KMS Key

 Do not clear the KMS key if you need to access or preserve any data on this appliance.

If you want to reinstall this appliance node (for example, in another grid), you must clear the KMS key. When the KMS key is cleared, all data on this appliance is deleted.

Clear KMS Key and Delete Data

The Node Encryption page includes these three sections:

- Encryption Status shows whether node encryption is enabled or disabled for the appliance.

- Key Management Server Details shows information about the KMS being used to encrypt the appliance. You can expand the server and client certificate sections to view certificate details and status.
  - To address issues with the certificates themselves, such as renewing expired certificates, see the information about KMS in the instructions for administering StorageGRID.
  - If there are unexpected problems connecting to KMS hosts, verify that the domain name system (DNS) servers are correct and that appliance networking is correctly configured.

#### Checking the DNS server configuration

- If you are unable to resolve your certificate issues, contact technical support.
- Clear KMS Key disables node encryption for the appliance, removes the association between the appliance and the key management server that was configured for the StorageGRID site, and deletes all data from the appliance. You must clear the KMS key before you can install the appliance into another StorageGRID system.

#### Clearing the key management server configuration




Clearing the KMS configuration deletes data from the appliance, rendering it permanently inaccessible. This data is not recoverable.

2. When you are done checking node-encryption status, reboot the node. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select one of these options:

- Select **Reboot into StorageGRID** to reboot the controller with the node rejoining the grid. Select this option if you are done working in maintenance mode and are ready to return the node to normal operation.
- Select **Reboot into Maintenance Mode** to reboot the controller with the node remaining in maintenance mode. Select this option if there are additional maintenance operations you need to perform on the node before rejoining the grid.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Related information

[Administer StorageGRID](#)

### Clearing the key management server configuration

Clearing the key management server (KMS) configuration disables node encryption on your appliance. After clearing the KMS configuration, the data on your appliance is permanently deleted and is no longer accessible. This data is not recoverable.

### What you'll need

If you need to preserve data on the appliance, you must perform a node decommission procedure before you clear the KMS configuration.



When KMS is cleared, data on the appliance will be permanently deleted and no longer accessible. This data is not recoverable.

Decommission the node to move any data it contains to other nodes in StorageGRID. See the recovery and maintenance instructions for grid node decommissioning.

### About this task

Clearing the appliance KMS configuration disables node encryption, removing the association between the appliance node and the KMS configuration for the StorageGRID site. Data on the appliance is then deleted and the appliance is left in a pre-install state. This process cannot be reversed.

You must clear the KMS configuration:

- Before you can install the appliance into another StorageGRID system, that does not use a KMS or that uses a different KMS.



Do not clear the KMS configuration if you plan to reinstall an appliance node in a StorageGRID system that uses the same KMS key.

- Before you can recover and reinstall a node where the KMS configuration was lost and the KMS key is not recoverable.
- Before returning any appliance that was previously in use at your site.
- After decommissioning a appliance that had node encryption enabled.



Decommission the appliance before clearing KMS to move its data to other nodes in your StorageGRID system. Clearing KMS before decommissioning the appliance will result in data loss and might render the appliance inoperable.

## Steps

1. Open a browser, and enter one of the IP addresses for the appliance's compute controller.

**`https://Controller_IP:8443`**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

2. Select **Configure Hardware > Node Encryption**.

### Node Encryption

Node encryption allows you to use an external key management server (KMS) to encrypt all StorageGRID data on this appliance. If node encryption is enabled for the appliance and a KMS is configured for the site, you cannot access any data on the appliance unless the appliance can communicate with the KMS.

### Encryption Status

You can only enable node encryption for an appliance during installation. You cannot enable or disable the node encryption setting after the appliance is installed.

Enable node encryption

Save

### Key Management Server Details

View the status and configuration details for the KMS that manages the encryption key for this appliance. You must use the Grid Manager to make configuration changes.

KMS display name	thales
External key UID	41b0306abcce451facfe01b1b4870ae1c1ec6bd5e3849d790223766baf35c57
Hostnames	10.96.99.164 10.96.99.165
Port	5696

Server certificate >

Client certificate >

### Clear KMS Key

Do not clear the KMS key if you need to access or preserve any data on this appliance.


If you want to reinstall this appliance node (for example, in another grid), you must clear the KMS key. When the KMS key is cleared, all data on this appliance is deleted.

Clear KMS Key and Delete Data



If the KMS configuration is cleared, data on the appliance will be permanently deleted. This data is not recoverable.

3. At the bottom of the window, select **Clear KMS Key and Delete Data**.
4. If you are sure that you want to clear the KMS configuration, type **clear** and select **Clear KMS Key and Delete Data**.



The KMS encryption key and all data are deleted from the node, and the appliance reboots. This can take up to 20 minutes.

5. Open a browser, and enter one of the IP addresses for the appliance's compute controller.  
**https://Controller\_IP:8443**

*Controller\_IP* is the IP address of the compute controller (not the storage controller) on any of the three StorageGRID networks.

The StorageGRID Appliance Installer Home page appears.

6. Select **Configure Hardware > Node Encryption**.
7. Verify that node encryption is disabled and that the key and certificate information in **Key Management Server Details** and the **Clear KMS Key and Delete Data** control are removed from the window.

Node encryption cannot be reenabled on the appliance until it is reinstalled in a grid.

### After you finish

After the appliance reboots and you have verified that KMS has been cleared and that the appliance is in a pre-install state, you can physically remove the appliance from your StorageGRID system. See the recovery and maintenance instructions for information about preparing an appliance for reinstallation.

### Related information

[Administer StorageGRID](#)



# Configure and manage

## Administer StorageGRID

Learn how to configure the StorageGRID system.

- [Administering a StorageGRID system](#)
- [Controlling administrator access to StorageGRID](#)
- [Configuring key management servers](#)
- [Managing tenants](#)
- [Configuring S3 and Swift client connections](#)
- [Managing StorageGRID networks and connections](#)
- [Configuring AutoSupport](#)
- [Managing Storage Nodes](#)
- [Managing Admin Nodes](#)
- [Managing Archive Nodes](#)
- [Migrating data into StorageGRID](#)

### Administering a StorageGRID system

Use these instructions to configure and administer a StorageGRID system.

These instructions describe how to use the Grid Manager to set up groups and users, create tenant accounts to allow S3 and Swift client applications to store and retrieve objects, configure and manage StorageGRID networks, configure AutoSupport, manage node settings, and more.



The instructions for managing objects with information lifecycle management (ILM) rules and policies have been moved to [Manage objects with ILM](#).

These instructions are for technical personnel who will be configuring, administering, and supporting a StorageGRID system after it has been installed.

#### What you'll need

- 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.

#### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87

Web browser	Minimum supported version
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

### Signing 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.

#### What you'll need

- You must have your login credentials.
- You must have the URL for the Grid Manager.
- You must be using a supported web browser.
- Cookies must be enabled in your web browser.
- You must have specific access permissions.

#### About this task

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, the Admin Nodes are not exactly the same:

- Alarm acknowledgments (legacy system) made on one Admin Node are not copied to other Admin Nodes. For this reason, the information displayed for alarms might not look the same on each Admin Node.
- Some maintenance procedures can only be performed from the primary Admin Node.

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 preferred Master, so that when you access the Grid Manager, you access it on the primary Admin Node unless the primary Admin Node is not available.

#### Steps

1. Launch a supported web browser.
2. In the browser's address bar, enter the URL for the Grid Manager:

```
https://FQDN_or_Admin_Node_IP/
```

where *FQDN\_or\_Admin\_Node\_IP* is a fully qualified domain name or the IP address of an Admin Node or

the virtual IP address of an HA group of Admin Nodes.

If you must access the Grid Manager on a port other than the standard port for HTTPS (443), enter the following, where *FQDN\_or\_Admin\_Node\_IP* is a fully qualified domain name or IP address, and *port* is the port number:


```
https://FQDN_or_Admin_Node_IP:port/
```

3. If you are prompted with a security alert, install the certificate using the browser's installation wizard.
4. Sign in to the Grid Manager:
  - If single sign-on (SSO) is not being used for your StorageGRID system:
    - i. Enter your username and password for the Grid Manager.
    - ii. Click **Sign In**.



The image shows the StorageGRID Grid Manager login page. On the left is the NetApp logo. The main heading is "StorageGRID® Grid Manager". Below the heading are two input fields: "Username" and "Password". A "Sign in" button is located at the bottom right of the form.

- If SSO is enabled for your StorageGRID system and this is the first time you have accessed the URL on this browser:
  - i. Click **Sign in**. You can leave the Account ID field blank.



The image shows the StorageGRID Sign in page. On the left is the NetApp logo. The main heading is "StorageGRID® Sign in". Below the heading is an "Account ID" input field containing a long string of zeros. Below the input field is the text "For Grid Manager, leave this field blank." A "Sign in" button is located at the bottom right of the form.

- ii. Enter your standard SSO credentials on your organization's SSO sign-in page. For example:

Sign in with your organizational account

someone@example.com

Password

Sign in

- If SSO is enabled for your StorageGRID system and you have previously accessed the Grid Manager or a tenant account:
  - i. Do either of the following:
    - Enter **0** (the account ID for the Grid Manager), and click **Sign in**.
    - Select **Grid Manager** if it appears in the list of recent accounts, and click **Sign in**.

StorageGRID® Sign in

Recent Grid Manager

Account ID 0

Sign in

- ii. 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 "Viewing the Dashboard" in the instructions for monitoring and troubleshooting StorageGRID.

Dashboard

**Health**

No current alerts. All grid nodes are connected.

**Information Lifecycle Management (ILM)**

Awaiting - Client 0 objects


Awaiting - Evaluation Rate 0 objects / second

Scan Period - Estimated 0 seconds

**Protocol Operations**

S3 rate 0 operations / second

Swift rate 0 operations / second



5. If you want to sign in to another Admin Node:

Option	Steps
SSO not enabled	<ol style="list-style-type: none"> <li>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.</li> <li>Enter your username and password for the Grid Manager.</li> <li>Click <b>Sign In</b>.</li> </ol>
SSO enabled	<p>In the browser's address bar, enter the fully qualified domain name or IP address of the other Admin Node.</p> <p>If you have signed in to one Admin Node, you can access other Admin Nodes without having to sign in again. However, if your SSO session expires, you are prompted for your credentials again.</p> <p><b>Note:</b> SSO is not available on the restricted Grid Manager port. You must use the default HTTPS port (443) if you want users to authenticate with single sign-on.</p>

**Related information**

[Web browser requirements](#)

[Controlling access through firewalls](#)

[Configuring server certificates](#)

[Configuring single sign-on](#)

[Managing admin groups](#)

[Managing high availability groups](#)

[Use a tenant account](#)


[Monitor & troubleshoot](#)

## Signing out of the Grid Manager

When you are done working with the Grid Manager, you must sign out to ensure that unauthorized users cannot access the StorageGRID system. Closing your browser might not sign you out of the system, based on browser cookie settings.

### Steps

1. Locate the **Sign Out** link in the top-right corner of the user interface.



2. Click **Sign Out**.

Option	Description
SSO not in use	<p>You are signed out of the Admin Node.</p> <p>The Grid Manager sign in page is displayed.</p> <p><b>Note:</b> If you signed into more than one Admin Node, you must sign out of each node.</p>
SSO enabled	<p>You are signed out of all Admin Nodes you were accessing. The StorageGRID sign in page is displayed. <b>Grid Manager</b> is listed as the default in the <b>Recent Accounts</b> drop-down, and the <b>Account ID</b> field shows 0.</p> <p><b>Note:</b> 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.</p>

### Related information

[Configuring single sign-on](#)

## Changing your password

If you are a local user of the Grid Manager, you can change your own password.

### What you'll need

You must be signed in to the Grid Manager using a supported browser.

### About this task

If you sign in to StorageGRID as a federated user or if single sign-on (SSO) is enabled, you cannot 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. Click **Save**.

## Changing 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 and encryption keys for the StorageGRID system.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have Maintenance or Root Access permissions.
- You must 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**.

NetApp® StorageGRID® Help ▾ | Root ▾ | Sign Out

Dashboard Alerts ▾ Nodes Tenants ILM ▾ Configuration ▾ Maintenance ▾ Support ▾

Grid Passwords  
Change the provisioning passphrase and other passwords for your StorageGRID system.

**Change Provisioning Passphrase**

The provisioning passphrase is required for any installation, expansion, or maintenance procedure that makes changes to the grid topology. This passphrase is also required to download backups of the grid topology information and encryption keys for the StorageGRID system. After changing the provisioning passphrase, you must download a new Recovery Package.

Current Provisioning Passphrase

New Provisioning Passphrase

Confirm New Provisioning Passphrase

2. Enter your current provisioning passphrase.
3. Enter the new passphrase. The passphrase must contain at least 8 and no more than 32 characters. Passphrases are case-sensitive.



Store the new provisioning passphrase in a secure location. It is required for installation, expansion, and maintenance procedures.

4. Re-enter the new passphrase, and click **Save**.

The system displays a green success banner when the provisioning passphrase change is complete. The change should take less than a minute.

NetApp® StorageGRID® Help ▾ | Root ▾ | Sign Out

Dashboard Alerts ▾ Nodes Tenants ILM ▾ Configuration ▾ Maintenance ▾ Support ▾

Grid Passwords  
Change the provisioning passphrase and other passwords for your StorageGRID system.

Provisioning passphrase successfully changed. Go to the [Recovery Package page](#) to download a new Recovery Package.

**Change Provisioning Passphrase**

The provisioning passphrase is required for any installation, expansion, or maintenance procedure that makes changes to the grid topology. This passphrase is also required to download backups of the grid topology information and encryption keys for the StorageGRID system. After changing the provisioning passphrase, you must download a new Recovery Package.

Current Provisioning Passphrase

New Provisioning Passphrase

Confirm New Provisioning Passphrase

5. Select the **Recovery Package page** link inside the success banner.
6. Download the new Recovery Package from the Grid Manager. Select **Maintenance > Recovery Package**

and enter the new provisioning passphrase.



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.

## Changing the browser session timeout

You can control whether Grid Manager and Tenant Manager users are signed out if they are inactive for more than a certain amount of time.

### What you'll need

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

### About this task

The GUI Inactivity Timeout defaults to 900 seconds (15 minutes). If a user's browser session is not active for this amount of time, the session times out.

As required, you can increase or decrease the timeout period by setting the GUI Inactivity Timeout display option.

If single sign-on (SSO) is enabled and a user's browser session times out, the system behaves as if the user clicked **Sign Out** manually. The user must reenter their SSO credentials to access StorageGRID again.

User session timeout can also be controlled by the following:



- A separate, non-configurable StorageGRID timer, which is included for system security. By default, 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 the value for the GUI Inactivity Timeout has not been reached. To renew the token, the user must sign back in.
- Timeout settings for the identity provider, assuming SSO is enabled for StorageGRID.

### Steps

1. Select **Configuration > System Settings > Display Options**.
2. For **GUI Inactivity Timeout**, enter a timeout period of 60 seconds or more.

Set this field to 0 if you do not want to use this functionality. Users are signed out 16 hours after they sign in, when their authentication tokens expire.



## Display Options

Updated: 2017-03-09 20:36:53 MST

Current Sender	ADMIN-DC1-ADM1
Preferred Sender	ADMIN-DC1-ADM1
GUI Inactivity Timeout	900
Notification Suppress All	<input type="checkbox"/>

Apply Changes

### 3. Click **Apply Changes**.

The new setting does not affect currently signed in users. Users must sign in again or refresh their browsers for the new timeout setting to take effect.

#### Related information

[How single sign-on works](#)

[Use a tenant account](#)

#### Viewing StorageGRID license information

You can view the license information for your StorageGRID system, such as the maximum storage capacity of your grid, whenever necessary.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

#### About this task

If there is an issue with the software license for this StorageGRID system, the Health panel on the Dashboard includes a License Status icon and a **License** link. The number indicates how many license-related issues there are.

#### Dashboard



## Step

To view the license, do one of the following:

- From the Health panel on the Dashboard, click the License status icon or the **License** link. This link appears only if there is an issue with the license.
- Select **Maintenance > System > License**.

The License Page appears and provides the following, read-only information about the current license:

- StorageGRID system ID, which is the unique identification number for this StorageGRID installation
- License serial number
- Licensed storage capacity of the grid
- Software license end date
- Support service contract end date
- Contents of the license text file



For licenses issued before StorageGRID 10.3, the licensed storage capacity is not included in the license file, and a "See License Agreement" message is displayed instead of a value.

## Updating 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.

### What you'll need

- You must have a new license file to apply to your StorageGRID system.
- You must have specific access permissions.
- You must have the provisioning passphrase.

### Steps

1. Select **Maintenance > System > License**.
2. Enter the provisioning passphrase for your StorageGRID system in the **Provisioning Passphrase** text box.
3. Click **Browse**.
4. In the Open dialog box, locate and select the new license file (.txt), and click **Open**.

The new license file is validated and displayed.

5. Click **Save**.

## Using 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.

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.

### 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 the information about using tenant accounts.
- `/private`: Access is restricted to Grid Manager users and is based on the configured group permissions. These APIs are intended for internal use only and are not publicly documented. These APIs are also subject to change without notice.

### Related information

[Use a tenant account](#)

[Prometheus: Query basics](#)

### Grid Management API operations

The Grid Management API organizes the available API operations into the following sections.

- **accounts** — Operations to manage storage tenant accounts, including creating new accounts and retrieving storage usage for a given account.
- **alarms** — Operations to list current alarms (legacy system), and return information about the health of the grid, including the current alerts and a summary of node connection states.
- **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*").



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 on 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.
- **endpoint-domain-names** — Operations to list and change 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).
- **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-health** — Operations on node health 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.
- **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.

## Issuing API requests

The Swagger user interface provides complete details and documentation for each API operation.

### What you'll need

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



Any API operations you perform using the API Docs webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

### Steps

1. Select **Help > API Documentation** from the Grid Manager header.
2. Select the desired operation.

When you expand an API operation, you can see the available HTTP actions, such as GET, PUT, UPDATE, and DELETE.

3. 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.

GET
/grid/groups Lists Grid Administrator Groups
🔒

Try it out

Name	Description
type string <small>(query)</small>	filter by group type  Available values : local, federated  <input type="text" value="--"/>
limit integer <small>(query)</small>	maximum number of results  Default value : 25  <input type="text" value="25"/>
marker string <small>(query)</small>	marker-style pagination offset (value is Group's URN)  <input type="text" value="marker - marker-style pagination offset (value"/>
includeMarker boolean <small>(query)</small>	if set, the marker element is also returned  <input type="text" value="--"/>
order string <small>(query)</small>	pagination order (desc requires marker)  Available values : asc, desc  <input type="text" value="--"/>

Responses
Response content type

Code	Description
200	successfully retrieved  Example Value   Model  <pre style="background-color: #2e3436; color: #eeeeec; padding: 10px; border: 1px solid #2e3436;">{   "responseTime": "2021-03-29T14:22:19.673Z",   "status": "success",   "apiVersion": "3.3",   "deprecated": false,   "data": [     {       "displayName": "Developers",</pre>

4. 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.
5. Determine if you need to modify the example request body. If so, you can click **Model** to learn the requirements for each field.
6. Click **Try it out**.
7. Provide any required parameters, or modify the request body as required.
8. Click **Execute**.
9. Review the response code to determine if the request was successful.

## Grid Management API versioning

The Grid Management API uses versioning to support non-disruptive upgrades.

For example, this Request URL specifies version 3 of the API.

```
https://hostname_or_ip_address/api/v3/authorize
```

The major version of the Tenant Management API is bumped when changes are made that are **not compatible** with older versions. The minor version of the Tenant Management 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

When you install StorageGRID software for the first time, only the most recent version of the Grid Management 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 use the Grid Management API to configure the supported versions. See the “config” section of the Swagger API documentation for more information. You should deactivate support for the older version after updating all Grid Management 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 v1 API at POST "/api/v1/authorize"
```

## Determining which API versions are supported in the current release

Use the following API request to return a list of the supported API major versions:

```
GET https://{{IP-Address}}/api/versions
{
  "responseTime": "2019-01-10T20:41:00.845Z",
  "status": "success",
  "apiVersion": "3.0",
  "data": [
    2,
    3
  ]
}
```

### Specifying an API version for a request

You can specify the API version using a path parameter (`/api/v3`) or a header (`Api-Version: 3`). If you provide both values, the header value overrides the path value.

```
curl https://[IP-Address]/api/v3/grid/accounts

curl -H "Api-Version: 3" https://[IP-Address]/api/grid/accounts
```

### Protecting 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.

### Using the API if single sign-on is enabled

If single sign-on (SSO) has been enabled for your StorageGRID system, you cannot use the standard Authenticate API requests to sign in to and sign out of the Grid Management API or the Tenant Management API.

### Signing in to 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 obtain an authentication token from AD FS that is valid for the Grid Management API or the Tenant Management API.

#### What you'll need

- 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 or CentOS, `./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 username
- The domain where StorageGRID is installed
- The address for StorageGRID
- If you want to access the Tenant Management API, enter tenant account ID.

```
python3 /tmp/storagegrid-ssoauth.py
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=fZHRT0MwFIZfhh...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=For
msAuthentication" --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=fZHRTomwFIZfhb...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=MjAxOC0xMS0wNzoxNjoxMjMjOjVpcMQ==;
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.

### Signing 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.

#### About this task

If required, you can sign out of the StorageGRID API simply 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://ads.example.com/ads/ls/?SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3
  D",
  "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
```

## Using StorageGRID 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 do not 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).

An external certificate authority (CA) can issue custom certificates that are fully compliant with your organization's information security policies. StorageGRID also includes a built-in certificate authority (CA) that generates internal CA certificates during system installation. These internal CA certificates are used, by default, to secure internal StorageGRID traffic. Although you can use the internal CA certificates 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 are not recommended.

- Custom CA certificates do not 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.

### System hardening

- 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.

Certificate	Certificate type	Description	Navigation location	Details
Administrator client certificate	Client	<p>Installed on each client, allowing StorageGRID to authenticate external client access.</p> <ul style="list-style-type: none"> <li>• Allows authorized external clients to access the StorageGRID Prometheus database.</li> <li>• Allows secure monitoring of StorageGRID using external tools.</li> </ul>	<b>Configuration &gt; Access Control &gt; Client Certificates</b>	<a href="#">Configuring administrator client certificates</a>
Identity federation certificate	Server	Authenticates the connection between StorageGRID and an external Active Directory, OpenLDAP, or Oracle Directory Server. Used for identity federation, which allows admin groups and users to be managed by an external system.	<b>Configuration &gt; Access Control &gt; Identity Federation</b>	<a href="#">Using identity federation</a>
Single sign-on (SSO) certificate	Server	Authenticates the connection between Active Directory Federation Services (AD FS) and StorageGRID that is used for single sign-on (SSO) requests.	<b>Configuration &gt; Access Control &gt; Single Sign-on</b>	<a href="#">Configuring single sign-on</a>

Certificate	Certificate type	Description	Navigation location	Details
Key management server (KMS) certificate	Server and client	<p>Authenticates the connection between StorageGRID and an external key management server (KMS), which provides encryption keys to StorageGRID appliance nodes.</p>	<p><b>Configuration &gt; System Settings &gt; Key Management Server</b></p>	<p><a href="#">Adding a key management server (KMS)</a></p>
Email alert notification certificate	Server and client	<p>Authenticates the connection between an SMTP email server and StorageGRID that is used for alert notifications.</p> <ul style="list-style-type: none"> <li>• If communications with the SMTP server requires Transport Layer Security (TLS), you must specify the email server CA certificate.</li> <li>• Specify a client certificate only if the SMTP email server requires client certificates for authentication.</li> </ul>	<p><b>Alerts &gt; Email Setup</b></p>	<p><a href="#">Monitor &amp; troubleshoot</a></p>

Certificate	Certificate type	Description	Navigation location	Details
Load balancer endpoint certificate	Server	<p>Authenticates the connection between S3 or Swift clients and the StorageGRID Load Balancer service on Gateway Nodes or Admin Nodes. You 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.</p> <p><b>Note:</b> The load balancer certificate is the most used certificate during normal StorageGRID operation.</p>	<p><b>Configuration &gt; Network Settings &gt; Load Balancer Endpoints</b></p>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring load balancer endpoints</a></li> <li>• Creating a load balancer endpoint for FabricPool</li> </ul> <p><a href="#">Configure StorageGRID for FabricPool</a></p>

Certificate	Certificate type	Description	Navigation location	Details
Management Interface Server Certificate	Server	<p>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.</p> <p>This certificate also authenticates Grid Management API and Tenant Management API connections.</p> <p>You can use the internal CA certificate or upload a custom certificate.</p>	<b>Configuration &gt; Network Settings &gt; Server Certificates</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring server certificates</a></li> <li>• <a href="#">Configuring a custom server certificate for the Grid Manager and the Tenant Manager</a></li> </ul>
Cloud Storage Pool endpoint certificate	Server	Authenticates the connection from the 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.	<b>ILM &gt; Storage Pools</b>	<a href="#">Manage objects with ILM</a>
Platform services endpoint certificate	Server	Authenticates the connection from the StorageGRID platform service to an S3 storage resource.	<b>Tenant Manager &gt; STORAGE (S3) &gt; Platform services endpoints</b>	<a href="#">Use a tenant account</a>

Certificate	Certificate type	Description	Navigation location	Details
Object Storage API Service Endpoint Server Certificate	Server	Authenticates secure S3 or Swift client connections to the Local Distribution Router (LDR) service on a Storage Node or to the deprecated Connection Load Balancer (CLB) service on a Gateway Node.	<b>Configuration &gt; Network Settings &gt; Load Balancer Endpoints</b>	<a href="#">Configuring a custom server certificate for connections to the Storage Node or the CLB service</a>

### 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 or Swift 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.

## Controlling administrator access to StorageGRID

You can control administrator access to the StorageGRID system by opening or closing firewall ports, managing admin groups and users, configuring single sign-on (SSO), and providing client certificates to allow secure external access to StorageGRID metrics.

- [Controlling access through firewalls](#)
- [Using identity federation](#)
- [Managing admin groups](#)
- [Managing local users](#)
- [Using single sign-on \(SSO\) for StorageGRID](#)
- [Configuring administrator client certificates](#)

## Controlling access through firewalls

When you want to control access through firewalls, you open or close specific ports at the external firewall.

### Controlling access 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.

Port	Description	If port is open...
443	Default HTTPS port for Admin Nodes	<p>Web browsers and management API clients can access the Grid Manager, the Grid Management API, the Tenant Manager, and the Tenant Management API.</p> <p><b>Note:</b> Port 443 is also used for some internal traffic.</p>
8443	Restricted Grid Manager port on Admin Nodes	<ul style="list-style-type: none"> <li>• Web browsers and management API clients can access the Grid Manager and the Grid Management API using HTTPS.</li> <li>• Web browsers and management API clients cannot access the Tenant Manager or the Tenant Management API.</li> <li>• Requests for internal content will be rejected.</li> </ul>
9443	Restricted Tenant Manager port on Admin Nodes	<ul style="list-style-type: none"> <li>• Web browsers and management API clients can access the Tenant Manager and the Tenant Management API using HTTPS.</li> <li>• Web browsers and management API clients cannot access the Grid Manager or the Grid Management API.</li> <li>• Requests for internal content will be rejected.</li> </ul>



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

[Signing in to the Grid Manager](#)

[Creating a tenant account if StorageGRID is not using SSO](#)

[Summary: IP addresses and ports for client connections](#)

[Managing untrusted Client Networks](#)

[Install Ubuntu or Debian](#)

[Install VMware](#)

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

## Using identity federation

Using identity federation makes setting up groups and users faster, and it allows users to sign in to StorageGRID using familiar credentials.

### Configuring identity federation

You can configure identity federation if you want admin groups and users to be managed in another system such as Active Directory, OpenLDAP, or Oracle Directory Server.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- If you plan to enable single sign-on (SSO), you must use Active Directory as the federated identity source and AD FS as the identity provider. See “Requirements for using single sign-on.”
- You must be using Active Directory, OpenLDAP, or Oracle Directory Server as the identity provider.



If you want to use an LDAP v3 service that is not listed, you must contact technical support.

- 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.

#### About this task

You must configure an identity source for the Grid Manager if you want to import the following types of federated groups:

- Administration 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 do not 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.

#### Steps

1. Select **Configuration > Access Control > Identity Federation**.
2. Select **Enable identity federation**.

The fields for configuring the LDAP server appear.

3. In the LDAP service type section, select the type of LDAP service you want to configure.

You can select **Active Directory**, **OpenLDAP**, or **Other**.



If you select **OpenLDAP**, you must configure the OpenLDAP server. See the guidelines for configuring an OpenLDAP server.



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.

- **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. In the Configure LDAP server section, enter the required LDAP server and network connection information.

- **Hostname:** The server hostname 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`

- **Password:** The password associated with the username.
- **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.

6. In the **Transport Layer Security (TLS)** section, select a security setting.

- **Use STARTTLS (recommended):** Use STARTTLS to secure communications with the LDAP server. This is the recommended option.
- **Use LDAPS:** The LDAPS (LDAP over SSL) option uses TLS to establish a connection to the LDAP server. This option is supported for compatibility reasons.
- **Do not use TLS:** The network traffic between the StorageGRID system and the LDAP server will not be secured.



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 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.

8. Optionally, select **Test connection** to validate your connection settings for the LDAP server.

A confirmation message appears in the upper right corner of the page if the connection is valid.

9. If the connection is valid, select **Save**.

The following screenshot shows example configuration values for an LDAP server that uses Active Directory.

## LDAP service type

Select the type of LDAP service you want to configure.

Active Directory

OpenLDAP

Other

## Configure LDAP server (All fields are required)

Hostname

my-active-directory.example.com

Port

389

Username

MyDomain\Administrator

Password

••••••••

Group Base DN

DC=storagegrid,DC=example,DC=com

User Base DN

DC=storagegrid,DC=example,DC=com

### Related information

[Supported ciphers for outgoing TLS connections](#)

[Requirements for using single sign-on](#)

[Creating a tenant account](#)

[Use a tenant account](#)

### 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.

## 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 Administrator's Guide for OpenLDAP.

## 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 Administrator's Guide for OpenLDAP.

## Related information

[OpenLDAP documentation: Version 2.4 Administrator's Guide](#)

## Forcing 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.

## What you'll need

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

## Steps

1. Select **Configuration > Access Control > Identity Federation**.

The Identity Federation page appears. The **Synchronize** button is at the bottom of the page.

### Synchronize

StorageGRID periodically synchronizes federated groups and users from the configured LDAP server. Clicking the button below will immediately start the synchronization process against the saved LDAP server.

Synchronize

2. Click **Synchronize**.

A confirmation message indicates that synchronization started successfully. 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.

## Disabling 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.

### What you'll need

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

### About this task

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 or alarms will not be raised for accounts that have not been synchronized.
- The **Enable Identity Federation** check box 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.

### Steps

1. Select **Configuration > Access Control > Identity Federation**.
2. Uncheck the **Enable Identity Federation** check box.
3. Click **Save**.

### Related information

[Disabling single sign-on](#)

## Managing 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.

### Creating admin groups

Admin groups allow you to determine which users can access which features and operations in the Grid Manager and the Grid Management API.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- If you plan to import a federated group, you must have configured identity federation and the federated group must already exist in the configured identity source.

### Steps

1. Select **Configuration > Access Control > Admin Groups**.

The Admin Groups page appears and lists any existing admin groups.

## Admin Groups

Add and manage local and federated user groups, allowing member users to sign in to the Grid Manager. Set group permissions to control access to specific pages and features.

<input type="button" value="+ Add"/> <input type="button" value="Clone"/> <input type="button" value="Edit"/> <input type="button" value="Remove"/>			
Name	ID	Group Type ?	Access Mode ?
<input checked="" type="radio"/> Flintstone	264083d0-23b5-3046-9bd4-88b7097731ab	Federated	Read-write
<input type="radio"/> Simpson	cc8ad11f-68d0-f84a-af29-e7a6fc63a2	Federated	Read-only
<input type="radio"/> ILM (read-only_group)	88446141-9599-4543-b183-9c227ce7767a	Local	Read-only
<input type="radio"/> API Developers	974b2faa-f9a1-4cfc-b364-914cdba2905f	Local	Read-write
<input type="radio"/> ILM Admins (read-write)	a528c0c2-2417-4559-86ed-f0d2e31da820	Local	Read-write
<input type="radio"/> Maintenance Users	7e3400ec-de8c-45a7-8bb8-e1496b362a8d	Local	Read-write

Group Type  Show  rows per page

### 2. Select **Add**.

The Add Group dialog box appears.

### Add Group

Create a new local group or import a group from the external identity source.

Group Type ?  Local  Federated

Display Name

Unique Name ?

Access Mode ?  Read-write  Read-only

#### Management Permissions

<input type="checkbox"/> Root Access ?	<input type="checkbox"/> Manage Alerts ?
<input type="checkbox"/> Acknowledge Alarms ?	<input type="checkbox"/> Grid Topology Page Configuration ?
<input type="checkbox"/> Other Grid Configuration ?	<input type="checkbox"/> Tenant Accounts ?
<input type="checkbox"/> Change Tenant Root Password ?	<input type="checkbox"/> Maintenance ?
<input type="checkbox"/> Metrics Query ?	<input type="checkbox"/> ILM ?
<input type="checkbox"/> Object Metadata Lookup ?	<input type="checkbox"/> Storage Appliance Administrator ?

3. For Group Type, select **Local** if you want to create a group that will be used only within StorageGRID, or select **Federated** if you want to import a group from the identity source.
4. If you selected **Local**, enter a display name for the group. The display name is the name that appears in the Grid Manager. For example, "Maintenance Users" or "ILM Administrators."
5. Enter a unique name for the group.
  - **Local**: Enter whatever unique name you want. For example, "ILM Administrators."
  - **Federated**: Enter the group's name exactly as it appears in the configured identity source.
6. 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 cannot 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.

7. Select one or more management 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.

8. Select **Save**.

The new group is created. If this is a local group, you can now add one or more users. If this is a federated group, the identity source manages which users belong to the group.

## Related information

[Managing local users](#)

### 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.

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
- Monitor grid topology
- View current and resolved alerts
- View current and historical alarms (legacy system)

- Change their own password (local users only)
- View certain information on the Configuration and Maintenance pages

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.

### Manage Alerts

This permission provides access to options for managing alerts. Users must have this permission to manage silences, alert notifications, and alert rules.

### Acknowledge Alarms (legacy system)

This permission provides access to acknowledge and respond to alarms (legacy system). All signed-in users can view current and historical alarms.

If you want a user to monitor grid topology and acknowledge alarms only, you should assign this permission.

### Grid Topology Page Configuration

This permission provides access to the following menu options:

- Configuration tabs available from the pages in **Support > Tools > Grid Topology**.
- **Reset event counts** link on the **Nodes > Events** tab.

### 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.

- **Alarms** (legacy system):
  - Global Alarms
  - Legacy Email Setup
- **ILM**:
  - Storage Pools
  - Storage Grades
- **Configuration > Network Settings**
  - Link Cost
- **Configuration > System Settings**:
  - Display Options
  - Grid Options
  - Storage Options

- **Configuration > Monitoring:**

- Events

- **Support:**

- AutoSupport

## Tenant Accounts

This permission provides access to the **Tenants > Tenant Accounts** page.



Version 1 of the Grid Management API (which has been deprecated) uses this permission to manage tenant group policies, reset Swift admin passwords, and manage root user S3 access keys.

## Change Tenant Root Password

This permission provides access to the **Change Root Password** option on the Tenant Accounts page, allowing you to control who can change the password for the tenant's local root user. Users who do not have this permission cannot see the **Change Root Password** option.



You must assign the Tenant Accounts permission to the group before you can assign this permission.

## Maintenance

This permission provides access to the following menu options:

- **Configuration > System Settings:**

- Domain Names\*
- Server Certificates\*

- **Configuration > Monitoring:**

- Audit\*

- **Configuration > Access Control:**

- Grid Passwords

- **Maintenance > Maintenance Tasks**

- Decommission
- Expansion
- Recovery

- **Maintenance > Network:**

- DNS Servers\*
- Grid Network\*
- NTP Servers\*

- **Maintenance > System:**

- License\*
- Recovery Package

- Software Update
- **Support > Tools:**
  - Logs
- Users who do not have the Maintenance permission can view, but not edit, the pages marked with an asterisk.

## Metrics Query

This permission provides access to the **Support > Tools > Metrics** page. This permission also provides access to custom Prometheus metrics queries using the **Metrics** section of the Grid Management API.

## ILM

This permission provides access to the following **ILM** menu options:

- **Erasure Coding**
- **Rules**
- **Policies**
- **Regions**



Access to the **ILM > Storage Pools** and **ILM > Storage Grades** menu options is controlled by the Other Grid Configuration and Grid Topology Page Configuration permissions.

## Object Metadata Lookup

This permission provides access to the **ILM > Object Metadata Lookup** menu option.

## Storage Appliance Administrator

This permission provides access to the E-Series SANtricity System Manager on storage appliances through the Grid Manager.

## 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.

## Deactivating features from the Grid Management 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 the 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 can ensure that no Admin user—including the root user and users belonging to groups with the Root Access permission—can change the password for any tenant account's root user.

## Reactivating 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 cannot 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.

For details, see the instructions for implementing S3 or Swift client applications.

### Steps

1. Access the Swagger documentation for 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.”

4. 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 Alarm Acknowledgment feature, send this PUT request:

```
{ "grid": { "alarmAcknowledgment": true } }
```

## Related information

[Using the Grid Management API](#)

### Modifying an admin group

You can modify an admin group to change the permissions associated with the group. For local admin groups, you can also update the display name.

### What you'll need

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

### Steps

1. Select **Configuration > Access Control > Admin Groups**.
2. Select the group.

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. You can then use your browser's find feature to search for a specific item in the currently displayed rows.

3. Click **Edit**.
4. Optionally, for local groups, enter the group's name that will appear to users, for example, "Maintenance Users."

You cannot change the unique name, which is the internal group name.

5. Optionally, change the group's Access Mode.
  - **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 cannot 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.

6. Optionally, add or remove group permissions.

See information about admin group permissions.

7. Select **Save**.

## Related information

[Admin group permissions](#)

### Deleting an admin 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 admin users from the group, but does not delete the admin users.

## What you'll need

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

## About this task

When you delete a group, users assigned to that group will lose all access privileges to the Grid Manager, unless they are granted privileges by a different group.

## Steps

1. Select **Configuration > Access Control > Admin Groups**.
2. Select the name of the group.

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. You can then use your browser's find feature to search for a specific item in the currently displayed rows.

3. Select **Remove**.
4. Select **OK**.

## Managing local users

You can create local users and assign them to local admin groups to determine which Grid Manager features these users can access.

The Grid Manager includes one predefined local user, named "root." Although you can add and remove local users, you cannot remove the root user.



If single sign-on (SSO) has been enabled, local users cannot sign in to StorageGRID.

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

## Creating a local user

If you have created local admin groups, you can create one or more local users and assign each user to one or more groups. The group's permissions control which Grid Manager features the user can access.

## About this task

You can only create local users, and you can only assign these users to local admin groups. Federated users and federated groups are managed using the external identity source.

## Steps

1. Select **Configuration > Access Control > Admin Users**.
2. Click **Create**.
3. Enter the user's display name, unique name, and password.
4. Assign the user to one or more groups that govern the access permissions.

The list of group names is generated from the Groups table.

5. Click **Save**.

## Related information

[Managing admin groups](#)

### Modifying a local user's account

You can modify a local admin user's account to update the user's display name or group membership. You can also temporarily prevent a user from accessing the system.

### About this task

You can edit local users only. Federated user details are automatically synchronized with the external identity source.

### Steps

1. Select **Configuration > Access Control > Admin Users**.
2. Select the user you want to edit.

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. You can then use your browser's find feature to search for a specific item in the currently displayed rows.

3. Click **Edit**.
4. Optionally, make changes to the name or group membership.
5. Optionally, to prevent the user from accessing the system temporarily, check **Deny Access**.
6. Click **Save**.

The new settings are applied the next time the user signs out and then signs back in to the Grid Manager.

### Deleting a local user's account

You can delete accounts for local users that no longer require access to the Grid Manager.

### Steps

1. Select **Configuration > Access Control > Admin Users**.
2. Select the local user you want to delete.



You cannot delete the predefined root local user.

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. You can then use your browser's find feature to search for a specific item in the currently displayed rows.

3. Click **Remove**.
4. Click **OK**.

### Changing a local user's password

Local users can change their own passwords using the **Change Password** option in the Grid Manager banner. In addition, users who have access to the Admin Users page can change passwords for other local users.

## About this task

You can change passwords for local users only. Federated users must change their own passwords in the external identity source.

## Steps

1. Select **Configuration > Access Control > Admin Users**.
2. From the Users page, select the user.

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. You can then use your browser's find feature to search for a specific item in the currently displayed rows.

3. Click **Change Password**.
4. Enter and confirm the password, and click **Save**.

## Using single sign-on (SSO) for StorageGRID

The StorageGRID system supports single sign-on (SSO) using the Security Assertion Markup Language 2.0 (SAML 2.0) standard. When SSO is enabled, 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 cannot sign in to StorageGRID.

- [How single sign-on works](#)
- [Requirements for using single sign-on](#)
- [Configuring single sign-on](#)

## How single sign-on works

Before enabling single sign-on (SSO), review how the StorageGRID sign-in and sign-out processes are affected when SSO is enabled.

## Signing 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:



- 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. Click **Sign in**

StorageGRID redirects you to your organization's SSO sign-in page. For example:

Sign in with your organizational account



[Sign in](#)

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 that has adequate access permission, you are signed in to the Grid Manager or the Tenant Manager, depending on which account you selected.
5. Optionally, access other Admin Nodes, or access the Grid Manager or the Tenant Manager, if you have adequate permissions.

You do not need to reenter your SSO credentials.

### Signing 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. Click **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
Tenant Manager on one or more Admin Nodes	Tenant Manager on any Admin Node	Tenant Manager on all Admin Nodes

If you are signed in to...	And you sign out from...	You are signed out of...
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 for using single sign-on

Before enabling single sign-on (SSO) for a StorageGRID system, review the requirements in this section.



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.

### Identity provider requirements

The identity provider (IdP) for SSO must meet the following requirements:

- Either of the following versions of Active Directory Federation Service (AD FS):
  - AD FS 4.0, included with Windows Server 2016



Windows Server 2016 should be using the [KB3201845 update](#), or higher.

- AD FS 3.0, included with Windows Server 2012 R2 update, or higher.
- Transport Layer Security (TLS) 1.2 or 1.3
- Microsoft .NET Framework, version 3.5.1 or higher

### Server certificate requirements

StorageGRID uses a Management Interface Server 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 SSO relying party trusts for StorageGRID in AD FS, you use the server certificate as the signature certificate for StorageGRID requests to AD FS.

If you have not already installed a custom server 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.



Using an Admin Node's default server certificate in the AD FS relying party trust 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 in AD FS 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`.

### Related information

[Controlling access through firewalls](#)

[Configuring a custom server certificate for the Grid Manager and the Tenant Manager](#)

### Configuring 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.

- [Confirming federated users can sign in](#)
- [Using sandbox mode](#)
- [Creating relying party trusts in AD FS](#)
- [Testing relying party trusts](#)
- [Enabling single sign-on](#)
- [Disabling single sign-on](#)
- [Temporarily disabling and reenabling single sign-on for one Admin Node](#)

### Confirming 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You are using Active Directory as the federated identity source and AD FS as the identity provider.

[Requirements for using single sign-on](#)

### 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 Control > Identity Federation**.
  - c. Confirm that the **Enable Identity Federation** check box 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, unselect the check box, and click **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 click **Edit Account**.
  - c. If the **Uses Own Identity Source** check box is selected, uncheck the box and click **Save**.

### Edit Tenant Account

#### Tenant Details

Display Name

**Uses Own Identity Source**

Allow Platform Services

Storage Quota (optional)

The Tenant Accounts page appears.

- d. Select the tenant account, click **Sign In**, and sign in to the tenant account as the local root user.
- e. From the Tenant Manager, click **Access Control > Groups**.
- f. Ensure that at least one federated group from the Grid Manager has been assigned the Root Access permission for this tenant.
- g. Sign out.

- h. Confirm you can sign back in to the tenant as a user in the federated group.

## Related information

[Requirements for using single sign-on](#)

[Managing admin groups](#)

[Use a tenant account](#)

## Using sandbox mode

You can use sandbox mode to configure and test Active Directory Federation Services (AD FS) relying party trusts before you enforce single sign-on (SSO) for StorageGRID users. After SSO is enabled, you can reenabling sandbox mode to configure or test new and existing relying party trusts. Reenabling sandbox mode temporarily disables SSO for StorageGRID users.

## What you'll need

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

## About this task

When SSO is enabled and a user attempts to sign in to an Admin Node, StorageGRID sends an authentication request to AD FS. In turn, AD FS sends an authentication response back to StorageGRID, indicating whether the authorization request was successful. For successful requests, the response includes a universally unique identifier (UUID) for the user.

To allow StorageGRID (the service provider) and AD FS (the identity provider) to communicate securely about user authentication requests, you must configure certain settings in StorageGRID. Next, you must use AD FS to create a relying party trust for every 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.



Using sandbox mode is highly recommended, but not strictly required. If you are prepared to create AD FS relying party trusts immediately after you configure SSO in StorageGRID, and you do not need to test the SSO and single logout (SLO) processes for each Admin Node, click **Enabled**, enter the StorageGRID settings, create a relying party trust for each Admin Node in AD FS, and then click **Save** to enable SSO.

## 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 do not appear, confirm you have configured Active Directory as the federated identity source. See “Requirements for using single sign-on.”

### 2. Select the **Sandbox Mode** option.

The Identity Provider and Relying Party settings appear. In the Identity Provider section, the **Service Type** field is read only. It shows the type of identity federation service you are using (for example, Active Directory).

### 3. In the Identity Provider section:

- a. Enter the Federation Service name, exactly as it appears in 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 whether you want to use Transport Layer Security (TLS) 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 and paste the certificate in the **CA Certificate** text box.

- **Do not use TLS:** Do not use a TLS certificate to secure the connection.

### 4. In the Relying Party section, specify the relying party identifier you will use for StorageGRID Admin Nodes when you configure relying party trusts.

- For example, if your grid has only one Admin Node and you do not 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 includes a relying party identifier for each Admin Node, based on the node’s hostname.

**NOTE:** 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.

### 5. Click **Save**.

- A green check mark appears on the **Save** button for a few seconds.

Save ✓

- The Sandbox mode confirmation notice appears, confirming that sandbox mode is now enabled. You can use this mode while you use AD FS to configure a relying party trust for each Admin Node and test the single sign-in (SSO) and single logout (SLO) processes.

### 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

**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.

### Related information

[Requirements for using single sign-on](#)

### Creating 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.

### Creating a relying party trust using Windows PowerShell

You can use Windows PowerShell to quickly create one or more relying party trusts.

### What you'll need

- You have configured SSO in StorageGRID, and you know the fully qualified domain name (or the IP address) and the relying party identifier for each Admin Node in your system.



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.

### About this task

These instructions apply to AD FS 4.0, which is included with Windows Server 2016. If you are using AD FS 3.0, which is included with Windows 2012 R2, you will notice slight differences in the procedure. See the Microsoft AD FS documentation if you have questions.

### Steps

1. From the Windows start menu, right-click the PowerShell icon, and select **Run as Administrator**.
2. At the PowerShell command prompt, enter the following command:

```
Add-AdfsRelyingPartyTrust -Name "Admin_Node_Identifier" -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. Click **Apply**, and click **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. Click **Add rule**.
- d. On the Select Rule Template page, select **Send LDAP Attributes as Claims** from the list, and click **Next**.
- e. On the Configure Rule page, enter a display name for this rule.

For example, **ObjectGUID to Name ID**.

- f. For the Attribute Store, select **Active Directory**.
- g. In the LDAP Attribute column of the Mapping table, type **objectGUID**.
- h. In the Outgoing Claim Type column of the Mapping table, select **Name ID** from the drop-down list.

- i. Click **Finish**, and click **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 simply 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.

## Creating 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.

### What you'll need

- You have configured SSO in StorageGRID, and you know the fully qualified domain name (or the IP address) and the relying party identifier for each Admin Node in your system.



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.

### About this task

These instructions apply to AD FS 4.0, which is included with Windows Server 2016. If you are using AD FS 3.0, which is included with Windows 2012 R2, you will notice slight differences in the procedure. See the Microsoft AD FS documentation if you have questions.

### Steps

1. In Windows Server Manager, click **Tools**, and then select **AD FS Management**.
2. Under Actions, click **Add Relying Party Trust**.
3. On the Welcome page, choose **Claims aware**, and click **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. Click **Add rule**:
- c. On the Select Rule Template page, select **Send LDAP Attributes as Claims** from the list, and click **Next**.
- d. On the Configure Rule page, enter a display name for this rule.

For example, **ObjectGUID to Name ID**.

- e. For the Attribute Store, select **Active Directory**.
- f. In the LDAP Attribute column of the Mapping table, type **objectGUID**.
- g. In the Outgoing Claim Type column of the Mapping table, select **Name ID** from the drop-down list.
- h. Click **Finish**, and click **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 simply 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.

## Creating a relying party trust manually

If you choose not to import the data for the relying party trusts, you can enter the values manually.

### What you'll need

- You have configured SSO in StorageGRID, and you know the fully qualified domain name (or the IP address) and the relying party identifier for each Admin Node in your system.



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 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.
- 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.

## About this task

These instructions apply to AD FS 4.0, which is included with Windows Server 2016. If you are using AD FS 3.0, which is included with Windows 2012 R2, you will notice slight differences in the procedure. See the Microsoft AD FS documentation if you have questions.

## Steps

1. In Windows Server Manager, click **Tools**, and then select **AD FS Management**.
2. Under Actions, click **Add Relying Party Trust**.
3. On the Welcome page, choose **Claims aware**, and click **Start**.
4. Select **Enter data about the relying party manually**, and click **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** check box.
- 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, click **Add rule**:
  - a. On the Select Rule Template page, select **Send LDAP Attributes as Claims** from the list, and click **Next**.
  - b. On the Configure Rule page, enter a display name for this rule.  
  
For example, **ObjectGUID to Name ID**.
  - c. For the Attribute Store, select **Active Directory**.

- d. In the LDAP Attribute column of the Mapping table, type **objectGUID**.
  - e. In the Outgoing Claim Type column of the Mapping table, select **Name ID** from the drop-down list.
  - f. Click **Finish**, and click **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. Click **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. Click **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 do not 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.

**Note:** 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. Click **Apply**, and click **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.

## Testing relying party trusts

Before you enforce the use of single sign-on (SSO) for StorageGRID, confirm that single sign-on and single logout (SLO) are correctly configured. If you created a relying party trust for each Admin Node, confirm you can use SSO and SLO for each Admin Node.

### What you'll need

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

- You have configured one or more relying party trusts in AD FS.

## Steps

1. Select **Configuration > Access Control > Single Sign-on**.

The Single Sign-on page appears, with the **Sandbox Mode** option selected.

2. In the instructions for sandbox mode, locate the link to your identity provider's sign-on page.

The URL is derived from the value you entered in the **Federated 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.

3. Click the link, or copy and paste the URL into a browser, to access your identity provider's sign-on page.
4. 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 click **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**

You are prompted to enter your username and password.

5. 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.

6. Repeat the previous steps to confirm you can sign in to any other Admin Nodes.

If all SSO sign-in and logout operations are successful, you are ready to enable SSO.

## Enabling single sign-on

After using sandbox mode to test all of your StorageGRID relying party trusts, you are ready to enable single sign-on (SSO).

### What you'll need

- You must have imported at least one federated group from the identity source and assigned Root Access management permissions to the group. You must 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.
- You must have tested all relying party trusts using sandbox mode.

### Steps

1. Select **Configuration > Access Control > Single Sign-on**.

The Single Sign-on page appears with **Sandbox Mode** selected.

2. Change the SSO Status to **Enabled**.
3. Click **Save**.

A warning message appears.

### Warning

#### Enable single sign-on

After you enable SSO, no local users—including the root user—will be able to sign in to the Grid Manager, the Tenant Manager, the Grid Management API, or the Tenant Management API.

Before proceeding, confirm the following:

- You have imported at least one federated group from the identity source and assigned Root Access management permissions to the group. You must 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.
- You have tested all relying party trusts using sandbox mode.

Are you sure you want to enable single sign-on?

Cancel

OK

4. Review the warning, and click **OK**.

Single sign-on is now 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.

## Disabling 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must 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. Click **Save**.

A warning message appears indicating that local users will now be able to sign in.



4. Click **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 disabling and reenabling 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 reenabling SSO for one Admin Node. To disable and then reenabling SSO, you must access the node's command shell.

### What you'll need

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

- You must 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 reenables 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** check box 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. Click **Save**.

Clicking **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. Click **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.

## Related information

[Configuring single sign-on](#)

## Configuring administrator client certificates

You can use client certificates to allow authorized external clients to access the StorageGRID Prometheus database. Client certificates provide a secure way to use 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.

## Adding administrator client certificates

To add a client certificate, you can provide your own certificate or generate one using the Grid Manager.

## What you'll need

- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.
- You must know the IP address or domain name of the Admin Node.
- You must have configured the StorageGRID Management Interface Server Certificate and have the corresponding CA bundle
- If you want to upload your own certificate, the public key and private key for the certificate must be available on your local computer.

## Steps

1. In the Grid Manager, select **Configuration > Access Control > Client Certificates**.

The Client Certificates page appears.

### Client Certificates

You can upload or generate one or more client certificates to allow StorageGRID to authenticate external client access.



2. Select **Add**.

The Upload Certificate page appears.

## Upload Certificate

Name 

Allow Prometheus 

### Certificate Details

Upload the public key for the client certificate.

Upload Client Certificate

Generate Client Certificate

Cancel

Save

3. Type a name between 1 and 32 characters for the certificate.
4. To access Prometheus metrics using your external monitoring tool, select the **Allow Prometheus** check box.
5. Upload or generate a certificate:
  - a. To upload a certificate, go [here](#).
  - b. To generate a certificate, go [here](#).
6. To upload a certificate:
  - a. Select **Upload Client Certificate**.
  - b. Browse for the public key for the certificate.

After you upload the public key for the certificate, the **Certificate metadata** and **Certificate PEM** fields are populated.

## Upload Certificate

Name  test-certificate-upload

Allow Prometheus 

### Certificate Details

Upload the public key for the client certificate.

Upload Client Certificate

Generate Client Certificate

Uploaded file name: client (1).crt

Certificate metadata 

```
Subject DN: /C=US/ST=California/L=Sunnyvale/O=Example Co./OU=IT/CN=*.s3.example.com
Serial Number: 0D:0E:FC:16:75:B8:BE:3E:7D:47:4D:05:49:08:F3:7B:E8:4A:71:90
Issuer DN: /C=US/ST=California/L=Sunnyvale/O=Example Co./OU=IT/CN=*.s3.example.com
Issued On: 2020-06-19T22:11:56.000Z
Expires On: 2021-06-19T22:11:56.000Z
SHA-1 Fingerprint: 13:AA:D6:06:2B:90:FE:B7:7B:EB:1A:83:BE:C3:62:39:B7:A6:E7:F0
SHA-256 Fingerprint: 5C:29:06:6B:CF:81:50:B8:4F:A9:56:F7:A7:AB:3C:36:FA:3D:B7:32:A4:C9:74:85:2C:8D:E6:67:37:C3:AC:60
```

Certificate PEM 

```
-----BEGIN CERTIFICATE-----
MIIDmzCCAoOgAwIBAgIUdQ78FnH4vj59R00FSQjze+hKcZAwDQYJKoZIhvcNAQEL
BQAwDELMAkGA1UEBhMCVVMxEzARBgNVBAgMCkNhbgG1mb3JuaWEwEjAQBgNVBAcM
CVN1bm55dmFeZTEUMBIGA1UECgwLRXhhbXBsZSBDbj4xCzAJBgNVBAsMAk1UMRkw
FwYDQDDBAQLnMzLmV4YW1wbGUuY29tMB4XDTIwMDYxOTIyMTE1cKDIxMDYx
OTIyMTE1NlowdDELMAkGA1UEBhMCVVMxEzARBgNVBAgMCkNhbgG1mb3JuaWEwEjAQ
BgNVBAcMCVN1bm55dmFeZTEUMBIGA1UECgwLRXhhbXBsZSBDbj4xCzAJBgNVBAsM
Ak1UMRkwFwYDQDDBAQLnMzLmV4YW1wbGUuY29tMIIBIjANBgkqhkiG9w0BAQEF
AAOCAQ8AMIIBCgKCAQEA5Vgq2MjVvotLeGtq1Co4coJmsQ2ygRhuwSza0bgMnjf
cwUgHNVFXGuG1zY/Tl37r3Dk5buZfyGYAeJ6mqbQA6cE3yp0p5SHx7Cm/ANJknFw6
```

Copy certificate to clipboard

Cancel

Save

- c. Select **Copy certificate to clipboard** and paste the certificate to your external monitoring tool.
  - d. Use an editing tool to copy and paste the private key to your external monitoring tool.
  - e. Select **Save** to save the certificate in the Grid Manager.
7. To generate a certificate:
- a. Select **Generate Client Certificate**.
  - b. Enter the domain name or IP address of the Admin Node.
  - c. Optionally, enter an X.509 subject, also referred to as the Distinguished Name (DN), to identify the administrator who owns the certificate.
  - d. Optionally, select the number of days the certificate is valid. The default is 730 days.
  - e. Select **Generate**.

The **Certificate metadata**, **Certificate PEM**, and **Certificate private key** fields are populated.

## Upload Certificate

Name  test-certificate-generate

Allow Prometheus 

### Certificate Details

Upload the public key for the client certificate.

Upload Client Certificate

Generate Client Certificate

Certificate metadata 

```
Subject DN: /CN=test.com
Serial Number: 08:F8:FB:76:B2:13:E4:DF:54:93:3D:35:56:6F:2A:03:53:B0:E2:0
A
Issuer DN: /CN=test.com
Issued On: 2020-11-20T22:44:46.000Z
Expires On: 2022-11-20T22:44:46.000Z
SHA-1 Fingerprint: 6E:DB:8C:F8:3E:20:68:E4:C6:42:52:5F:32:7E:E7:93:66:69:F3:3
D
SHA-256 Fingerprint: 73:D3:51:83:ED:D3:89:AD:7B:89:4C:AF:AE:34:76:B6:42:FE:0D:
EF:78:C0:A4:66:C2:EB:65:64:C3:D4:7A:B0
```

Certificate PEM 

```
-----BEGIN CERTIFICATE-----
MIIECyCCAbOgAwIBAgIUUCPj7dxIT8N9Uge01Vm8qA1Ov4gowDQYJKoZIhvcNAQEL
BQAwEwERMA8GA1UEAwwIdGVudC5jb20wHhcNMjA0MTIwMjI0NDQ2WWhcNMjIwNDQ2
WjATMREwDwYDVQQDDAh0ZXN0LmNvbnRCCASAwDQYJKoZIhvcNAQEBBQADggEPADCCAQcC
ggEBBAK02dB9mx2jFrGuBb2ZMjcidf/+TeKxLcBgm+4vIwt1gvrR
XgH231B9YIQn/Vo729R2mNKKyBwkyQTkGCO2Ixxv08TBLcIWf88TgcIcMyt1V1F
OasEWFYs4O2xxjnR3/X+AX+6s2WZLcVa+3CDjGu4ic0V/uVQax4yA1T9SoRnjBmOa
LCVjL6iVnkUGB88GkyUFcOaMjeL6TN1QocFv9VEB0xBKCP4D7FDbaIy2F9Ng5s9
EEOQoLN=NxXCasLO4D7jZqFqOYUpFJSMOoh1x0nSpQ78Z5KfYwVtDKg5v52P9UBM
1o6GeuoFaW+dbpLZMPO9N1VtFhghXe9RaxN8e-jkCAwEAAaMXXMBUwEwYDVDR0RBww
-----
```

Copy certificate to clipboard

Certificate private key 

```
-----BEGIN RSA PRIVATE KEY-----
MIIEpQIBAAKCAQEAsT2OH2bHaM+sa4Fv2kyNyJ1/+1NwxEu0Eab7i8jC2KWC/BFe
AdneUH1ghCf9WjvblHaY0oxIHCTJBOQYI5kjG+/RjMEt4h29sKx0BwigzK2VWU7
OwF2jPg7bFG0cxf9s4Bf7xN12kixV75IIOMa7iJzRX+5VDFHjIDVf1KqgeMgYSoe
JWmvgJWeRQYFI2uTJQ945ggyOwvPM2VDOgW/1UQHTEEoKngFeUNtojLL/02DmtJS
Q8Cge202xcJxMe7gPuNmoW05h8kUnow6iHXHSfm1Dvxnkp9jBw0MqDm/nY/xQEaW
jw266h9pb5lukt2k703VW0WGCfD70DPE3yyOQIDAQABoIBAQCfEUfV4pE0Hqov
2uEL6De4yXMTvg/8Gn+W8mvtDgQB4xWEGQrkk1kiEUG+HThYzFJcn6XX0vACDYAC/
Hh1Q67xDPvRjdpuR0cr1W8evveEmpBx99MqH9Y2UGx6Yub3USJaqfDvjA4Nvaon
MxaYJRFBIvAR7E2x2xXVY8b0sRPAjrn0YCa1Lct5Y0K79e0G8naTmwIdm2YM6EE
-----
```

Copy private key to clipboard

 You will not be able to view the certificate private key after you close this dialog. To save the keys for future reference, copy and paste the values to another location.

Cancel

Save

- f. Select **Copy certificate to clipboard** and paste the certificate to your external monitoring tool.
- g. Select **Copy private key to clipboard** and paste the key to your external monitoring tool.



You will not be able to view the private key after you close the dialog box. Copy the key to a safe location.

- h. Select **Save** to save the certificate in the Grid Manager.

8. Configure the following settings on your external monitoring tool, such as Grafana.

A Grafana example is shown in the following screenshot:

Name ⓘ  Default

## HTTP

URL ⓘ

Access  Help >

Whitelisted Cookies ⓘ

## Auth

Basic auth  With Credentials ⓘ

TLS Client Auth  With CA Cert ⓘ

Skip TLS Verify

Forward OAuth Identity ⓘ

### TLS/SSL Auth Details ⓘ

CA Cert

ServerName

Client Cert

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 Authorization** and **With CA Cert**.
- d. Copy and paste the Management Interface Server Certificate or CA bundle to **CA Cert** under TLS/SSL Auth Details.
- e. **ServerName**: Enter the domain name of the Admin Node.

ServerName must match the domain name as it appears in the Management Interface Server Certificate.

- f. 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 and troubleshooting StorageGRID.

### Related information

[Using StorageGRID security certificates](#)

[Configuring a custom server certificate for the Grid Manager and the Tenant Manager](#)

[Monitor & troubleshoot](#)

### Editing administrator client certificates

You can edit a certificate to change its name, enable or disable Prometheus access, or upload a new certificate when the current one has expired.

### What you'll need

- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.
- You must know the IP address or domain name of the Admin Node.
- If you want to upload a new certificate and private key, they must be available on your local computer.

### Steps

1. Select **Configuration > Access Control > Client Certificates**.

The Client Certificates page appears. The existing certificates are listed.

Certificate expiration dates 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.

<input type="button" value="+ Add"/> <input type="button" value="✎ Edit"/> <input type="button" value="✕ Remove"/>			
	Name	Allow Prometheus	Expiration Date
<input type="radio"/>	test-certificate-upload	✓	2021-06-19 16:11:56 MDT
<input checked="" type="radio"/>	test-certificate-generate	✓	2022-08-20 09:42:00 MDT

Displaying 2 certificates.

2. Select the radio button to the left of the certificate you want to edit.
3. Select **Edit**.

The Edit Certificate dialog box appears.

Edit Certificate test-certificate-generate

Name

Allow Prometheus

### Certificate Details

Upload the public key for the client certificate.

Certificate metadata

```
Subject DN: /CN=test.com
Serial Number: 0C:11:87:6C:1E:FD:13:16:F3:F2:06:D9:DA:6D:BC:CE:2A:A9:C3:53
Issuer DN: /CN=test.com
Issued On: 2020-11-23T15:53:33.000Z
Expires On: 2022-11-23T15:53:33.000Z
SHA-1 Fingerprint: AE:E6:70:A7:D3:C3:39:7A:09:F9:62:9B:81:8A:87:CD:43:16:89:A7
SHA-256 Fingerprint: 63:07:BF:FF:08:1E:84:F1:D4:67:C6:16:B0:35:26:00:C6:A3:13:11:7E:5E:9
0:EC:7A:7B:EF:23:14:55:3D:56
```

Certificate PEM

```
-----BEGIN CERTIFICATE-----
MIICyzCCRbCgAwIBAgIUDBGHbE79Exbz8gbZ2m28ziqpwlMwDQYJKoZIhvcNAQEL
BQAwEzERMASGA1UEAwIdGVzdC5jb20wHhcNMjAxMTIzMTU1MzE1MzE1MzE1MzE1
MTU1MzE1MzE1MzE1MzE1MzE1MzE1MzE1MzE1MzE1MzE1MzE1MzE1MzE1MzE1MzE1
ggEPADCCAQcCggEBAKdgEcneCDFDsLjvLnX9ow6oPrdU7m2EN6SS6xdVI155sCH+
hkwO5a2Mym7EhbNrfwOt2nMjQkcaKlrk8OAmutRgG6N1N12FIW0qYQuzFQ0QddLq
n7ymFw6wSa9zYSu7bLp84Yn0/LSDPk+h3Jio7Mrt2X70It5ZDRwFmbLNvEvYEtTS
h+FbNh885AIRO2eLxvC0IRij1bySe76wK+Wmc97HdxRSQyxIWk6BD47XC+dOrv55
wvtjc/41qc5xsE6XmJs2yJg4VARr10y8Icwa9fr00+xPwIdCQNWxkpWJXeBnCoXx
YcQxbWzjz+iVLJqLTMxU8xTTI30rUqN00M82GJUCAwEAAAMXMBUwEwYDVR0RBAww
-----END CERTIFICATE-----
```

4. Make the desired changes to the certificate.
5. Select **Save** to save the certificate in the Grid Manager.
6. If you uploaded a new certificate:
  - a. Select **Copy certificate to clipboard** to paste the certificate to your external monitoring tool.
  - b. Use an editing tool to copy and paste the new private key to your external monitoring tool.
  - c. Save and test the certificate and private key in your external monitoring tool.
7. If you generated a new certificate:
  - a. Select **Copy certificate to clipboard** to paste the certificate to your external monitoring tool.
  - b. Select **Copy private key to clipboard** to paste the certificate to your external monitoring tool.



You will not be able to view or copy the private key after you close the dialog box. Copy the key to a safe location.

- c. Save and test the certificate and private key in your external monitoring tool.

## Removing administrator client certificates

If you no longer need a certificate, you can remove it.


### What you'll need

- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.

### Steps

1. Select **Configuration > Access Control > Client Certificates**.

The Client Certificates page appears. The existing certificates are listed.




	Name	Allow Prometheus	Expiration Date
<input type="radio"/>	test-certificate-upload	✓	2021-06-19 16:11:56 MDT
<input checked="" type="radio"/>	test-certificate-generate	✓	2022-08-20 09:42:00 MDT

Displaying 2 certificates.

2. Select the radio button to the left of the certificate you want to remove.
3. Select **Remove**.

A confirmation dialog box appears.



4. Select **OK**.

The certificate is removed.

## Configuring key management servers

You can configure one or more external key management servers (KMS) to protect the data on specially configured appliance nodes.

### 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).

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 cannot 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.

### Reviewing StorageGRID encryption methods

StorageGRID provides a number of 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 ( <b>Configuration &gt; System Settings &gt; Key Management Server</b> ) 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 <b>Node Encryption</b> enabled during installation. All data on the appliance is protected against physical loss or removal from the data center. Can be used with some StorageGRID storage and services appliances.
Drive security in SANtricity System Manager	If the Drive Security feature is enabled for a storage 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 Federal Information Processing Standard (FIPS) drives. All data on the secured drives is protected against physical loss or removal from the data center. Cannot be used with some storage appliances or with any service appliances.  <a href="#">SG6000 storage appliances</a>  <a href="#">SG5700 storage appliances</a>  <a href="#">SG5600 storage appliances</a>

Encryption option	How it works	Applies to
Stored Object Encryption grid option	<p>The <b>Stored Object Encryption</b> option can be enabled in the Grid Manager (<b>Configuration &gt; System Settings &gt; Grid Options</b>). When enabled, any new objects that are not encrypted at the bucket level or at the object level are encrypted during ingest.</p>	<p>Newly ingested S3 and Swift object data. Existing stored objects are not encrypted. Object metadata and other sensitive data are not encrypted.</p> <p><a href="#">Configuring stored object encryption</a></p>
S3 bucket encryption	<p>You issue a PUT Bucket encryption request to enable encryption for the bucket. Any new objects that are not encrypted at the object level are encrypted during ingest.</p>	<p>Newly ingested S3 object data only. Encryption must be specified for the bucket. Existing bucket objects are not encrypted. Object metadata and other sensitive data are not encrypted.</p> <p><a href="#">Use S3</a></p>
S3 object server-side encryption (SSE)	<p>You issue an S3 request to store an object and include the <code>x-amz-server-side-encryption</code> request header.</p>	<p>Newly ingested S3 object data only. Encryption must be specified for the object. Object metadata and other sensitive data are not encrypted.</p> <p>StorageGRID manages the keys.</p> <p><a href="#">Use S3</a></p>
S3 object server-side encryption with customer-provided keys (SSE-C)	<p>You issue an S3 request to store an object and include three request headers.</p> <ul style="list-style-type: none"> <li>• <code>x-amz-server-side-encryption-customer-algorithm</code></li> <li>• <code>x-amz-server-side-encryption-customer-key</code></li> <li>• <code>x-amz-server-side-encryption-customer-key-MD5</code></li> </ul>	<p>Newly ingested S3 object data only. Encryption must be specified for the object. Object metadata and other sensitive data are not encrypted.</p> <p>Keys are managed outside of StorageGRID.</p> <p><a href="#">Use S3</a></p>

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.	<p>All object data, metadata, and system configuration data, assuming every volume or datastore is encrypted.</p> <p>An external encryption method provides tighter control over encryption algorithms and keys. Can be combined with the other methods listed.</p>
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.	<p>Object data and metadata only (system configuration data is not encrypted).</p> <p>An external encryption method provides tighter control over encryption algorithms and keys. Can be combined with the other methods listed.</p> <p><a href="#">Amazon Simple Storage Service - Developer Guide: Protecting data using client-side encryption</a></p>

### Using 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 grid 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.

### Overview of 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.

### Setting 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.	<a href="#">Configuring StorageGRID as a client in the KMS</a>
Obtain the required information for the StorageGRID client on the KMS.	<a href="#">Configuring StorageGRID as a client in the KMS</a>
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.	<a href="#">Adding a key management server (KMS)</a>

### Setting 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 cannot enable the **Node Encryption** setting after an appliance is added to the grid, and you cannot use external key management for appliances that do not 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 cannot 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.

For details, refer to the following:

- [SG100 & SG1000 services appliances](#)
- [SG6000 storage appliances](#)
- [SG5700 storage appliances](#)
- [SG5600 storage appliances](#)

### 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, cannot 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.

### What are the KMIP requirements?

StorageGRID supports KMIP version 1.4.

### [Key Management Interoperability Protocol Specification Version 1.4](#)

Communications between the appliance nodes and the configured KMS use secure TLS connections. StorageGRID supports the following TLS v1.2 ciphers for KMIP:

- TLS\_ECDHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384
- TLS\_ECDHE\_ECDSA\_WITH\_AES\_256\_GCM\_SHA384

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.

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.

### 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 cannot enable node encryption after an appliance is added to the grid, and you cannot use external key management for appliances that do not have node encryption enabled.

You can use the configured KMS for the following StorageGRID appliances and appliance nodes:

Appliance	Node type
SG1000 services appliance	Admin Node or Gateway Node
SG100 services appliance	Admin Node or Gateway Node
SG6000 storage appliance	Storage Node
SG5700 storage appliance	Storage Node
SG5600 storage appliance	Storage Node

You cannot use the configured KMS for software-based (non-appliance) nodes, including the following:

- Nodes deployed as virtual machines (VMs)
- Nodes deployed within Docker containers 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 cannot 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 rotating the encryption key, use the KMS software to rotate from the last used version of the key to a new version of the same key. Do not rotate to an entirely different key.



Never attempt to rotate a key by changing the key name (alias) for the KMS in the Grid Manager. Instead, rotate the key by updating the key version in the KMS software. Use the same key alias for new keys as was used for previous keys. If you change the key alias for a configured KMS, StorageGRID might not be able to decrypt your data.

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 cannot 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.

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)


## 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 do not 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.



4. 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
<p>You have one or more site-specific KMS entries, and you want to use one of them as the default KMS.</p>	<p>Edit the site-specific KMS. In the <b>Manages keys for</b> field, select <b>Sites not managed by another KMS (default KMS)</b>. The site-specific KMS will now be used as the default KMS. It will apply to any sites that do not have a dedicated KMS.</p> <p><a href="#">Editing a key management server (KMS)</a></p>
<p>You have a default KMS and you add a new site in an expansion. You do not want to use the default KMS for the new site.</p>	<ol style="list-style-type: none"> <li>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.</li> <li>2. Using the Grid Manager, add the new KMS and select the site.</li> </ol> <p><a href="#">Adding a key management server (KMS)</a></p>
<p>You want the KMS for a site to use a different server.</p>	<ol style="list-style-type: none"> <li>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.</li> <li>2. Using the Grid Manager, edit the existing KMS configuration and enter the new host name or IP address.</li> </ol> <p><a href="#">Adding a key management server (KMS)</a></p>

## Configuring 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.

### About this task

These instructions apply to Thales CipherTrust Manager k170v, versions 2.0, 2.1, and 2.2. If you have questions about using a different key management server with StorageGRID, contact technical support.

### [Thales CipherTrust Manager](#)

#### 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. From the KMS software, create an AES encryption key for each KMS or KMS cluster.

The encryption key needs to be exportable.

### 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.



The encryption key must already exist in the KMS. StorageGRID does not create or manage KMS keys.

### 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.

## Adding a key management server (KMS)

You use the StorageGRID Key Management Server wizard to add each KMS or KMS cluster.

### What you'll need

- You must have reviewed the [considerations and requirements for using a key management server](#).
- You must have [configured StorageGRID as a client in the KMS](#), and you must have the required information for each KMS or KMS cluster
- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.

### 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.

## Considerations for changing the KMS for a site

### Steps

1. [Step 1: Enter KMS Details](#)
2. [Step 2: Upload Server Certificate](#)
3. [Step 3: Upload Client Certificates](#)

### Step 1: Enter KMS Details

In Step 1 (Enter KMS Details) of the Add a Key Management Server wizard, you provide details about the KMS or KMS cluster.

### Steps

1. Select **Configuration > System Settings > Key Management Server**.

The Key Management Server page appears with the Configuration Details tab selected.

Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details [Encrypted Nodes](#)

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

[+ Create](#) [Edit](#) [Remove](#)

KMS Display Name	Key Name	Manages keys for	Hostname	Certificate Status
No key management servers have been configured. Select <a href="#">Create</a> .				

2. Select **Create**.

Step 1 (Enter KMS Details) of the Add a Key Management Server wizard appears.

## Add a Key Management Server



Enter information about the external key management server (KMS) and the StorageGRID client you configured in that KMS. If you are configuring a KMS cluster, select + to add a hostname for each server in the cluster.

KMS Display Name	<input type="text"/>
Key Name	<input type="text"/>
Manages keys for	<input type="text" value="-- Choose One --"/>
Port	<input type="text" value="5696"/>
Hostname	<input type="text"/>

+

Cancel

Next

3. Enter the following information for the KMS and the StorageGRID client you configured in that KMS.

Field	Description
KMS Display 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.

Field	Description
Manages keys for	<p>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.</p> <ul style="list-style-type: none"> <li>• Select a site if this KMS will manage encryption keys for the appliance nodes at a specific site.</li> <li>• Select <b>Sites not managed by another KMS (default KMS)</b> to configure a default KMS that will apply to any sites that do not have a dedicated KMS and to any sites you add in subsequent expansions.</li> </ul> <p><b>Note:</b> 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.</p>
Port	<p>The port the KMS server uses for Key Management Interoperability Protocol (KMIP) communications. Defaults to 5696, which is the KMIP standard port.</p>
Hostname	<p>The fully qualified domain name or IP address for the KMS.</p> <p><b>Note:</b> The 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.</p>

4. If you are using a KMS cluster, select the plus sign **+** to add a hostname for each server in the cluster.
5. Select **Next**.

Step 2 (Upload Server Certificate) of the Add a Key Management Server wizard appears.

### 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.

## Add a Key Management Server



Upload a server certificate signed by the certificate authority (CA) on the external key management server (KMS) or a certificate bundle. The server certificate allows the KMS to authenticate itself to StorageGRID.

Server Certificate 

Cancel

Back

Next

2. Upload the certificate file.

The server certificate metadata appears.

## Add a Key Management Server



Upload a server certificate signed by the certificate authority (CA) on the external key management server (KMS) or a certificate bundle. The server certificate allows the KMS to authenticate itself to StorageGRID.

Server Certificate ⓘ  k170vCA.pem

### Server Certificate Metadata

```
Server DN: /C=US/ST=MD/L=Belcamp/O=Gemalto/CN=KeySecure Root CA
Serial Number: 71:CD:6D:72:53:B5:6D:0A:8C:69:13:0D:4D:D7:81:0E
Issue DN: /C=US/ST=MD/L=Belcamp/O=Gemalto/CN=KeySecure Root CA
Issued On: 2020-10-15T21:12:45.000Z
Expires On: 2030-10-13T21:12:45.000Z
SHA-1 Fingerprint: EE:E4:6E:17:86:DF:56:B4:F5:AF:A2:3C:BD:56:6B:10:DB:B2:5A:79
```

Cancel

Back

Next



If you uploaded a certificate bundle, the metadata for each certificate appears on its own tab.

### 3. Select **Next**.

Step 3 (Upload Client Certificates) of the Add a Key Management Server wizard appears.

#### 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.

## Add a Key Management Server



Upload the client certificate and the client certificate private key. The client certificate is issued to StorageGRID by the external key management server (KMS), and it allows StorageGRID to authenticate itself to the KMS.

Client Certificate 

Client Certificate Private Key 

Cancel

Back

Save

### 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.

The metadata for the client certificate and the client certificate private key appear.

## Add a Key Management Server



Upload the client certificate and the client certificate private key. The client certificate is issued to StorageGRID by the external key management server (KMS), and it allows StorageGRID to authenticate itself to the KMS.

Client Certificate   k170vClientCert.pem

```
Server DN: /CN=admin/UID=  
Serial Number: 7D:5A:8A:27:02:40:C8:F5:19:A1:28:22:E7:D6:E2:EB  
Issue DN: /C=US/ST=MD/L=Belcamp/O=Gemalto/CN=KeySecure Root CA  
Issued On: 2020-10-15T23:31:49.000Z  
Expires On: 2022-10-15T23:31:49.000Z  
SHA-1 Fingerprint: A7:10:AC:39:85:42:80:8F:FF:62:AD:A1:BD:CF:4C:90:F3:E9:36:69
```

Client Certificate Private Key   k170vClientKey.pem

Cancel

Back

Save

### 5. Select **Save**.

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 **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**.

## Add a Key Management Server



Upload the client certificate and the client certificate private key. The client certificate is issued to StorageGRID by the external key management server (KMS), and it allows StorageGRID to authenticate itself to the KMS.

Client Certificate ⓘ  k170vClientCert.pem

Server DN: /CN=admin/UID=  
Serial Number: 7D:5A:8A:27:02:40:C8:F5:19:A1:28:22:E7:D6:E2:EB  
Issue DN: /C=US/ST=MD/L=Belcamp/O=Gemalto/CN=KeySecure Root CA  
Issued On: 2020-10-15T23:31:49.000Z  
Expires On: 2022-10-15T23:31:49.000Z  
SHA-1 Fingerprint: A7:10:AC:39:85:42:80:8F:FF:62:AD:A1:BD:CF:4C:90:F3:E9:36:69

Client Certificate Private Key ⓘ  k170vClientKey.pem

Select **Force Save** to save this KMS without testing the external connections. If there is an issue with the configuration, you might not be able to reboot any FDE-enabled appliance nodes at the affected site, and you might lose access to your data.

Cancel

Back

Force Save

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.

### Warning

Confirm force-saving the KMS configuration

Are you sure you want to save this KMS without testing the external connections?

If there is an issue with the configuration, you might not be able to reboot any appliance nodes with node encryption enabled at the affected site, and you might lose access to your data.

Cancel

OK

The KMS configuration is saved but the connection to the KMS is not tested.

## Viewing KMS details

You can view information about each key management server (KMS) in your StorageGRID system, including the current status of the server and client certificates.

### Steps

1. Select **Configuration > System Settings > Key Management Server**.

The Key Management Server page appears. The Configuration Details tab shows any key management servers that are configured.

#### Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details Encrypted Nodes

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

+ Create Edit Remove

KMS Display Name	Key Name	Manages keys for	Hostname	Certificate Status
Default KMS	test	Sites not managed by another KMS (default KMS)	10.96.99.164	✓ All certificates are valid

2. Review the information in the table for each KMS.

Field	Description
KMS Display Name	The descriptive name of the KMS.
Key Name	The key alias for the StorageGRID client in the KMS.
Manages keys for	The StorageGRID site associated with the KMS.  This field displays the name of a specific StorageGRID site or <b>Sites not managed by another KMS (default KMS)</b> .

Field	Description
Hostname	<p>The fully qualified domain name or IP address of the KMS.</p> <p>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.</p> <p>For example: 10.10.10.10 and 10.10.10.11 or 10.10.10.10 and 2 others.</p> <p>To view all hostnames in a cluster, select a KMS and then select <b>Edit</b>.</p>
Certificate Status	<p>Current state of the server certificate, optional CA certificate, and the client certificate: valid, expired, nearing expiration, or unknown.</p> <p><b>Note:</b> It might take StorageGRID as long as 30 minutes to get updates to the certificate status. You must refresh your web browser to see the current values.</p>

- If the Certificate Status is Unknown, wait up to 30 minutes and then refresh your web browser.



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 actual status.

- If the Certificate Status column indicates that a certificate has expired or is nearing expiration, address the issue as soon as possible.

See the recommended actions for the **KMS CA certificate expiration**, **KMS client certificate expiration**, and **KMS server certificate expiration** alerts in the instructions for monitoring and troubleshooting StorageGRID.



You must address any certificate issues as soon as possible to maintain data access.

## Related information

[Monitor & troubleshoot](#)

## Viewing 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 > System Settings > Key Management Server**.

The Key Management Server page appears. The Configuration Details tab shows any key management servers that have been configured.

#### Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details **Encrypted Nodes**

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

+ Create Edit Remove

KMS Display Name	Key Name	Manages keys for	Hostname	Certificate Status
Default KMS	test	Sites not managed by another KMS (default KMS)	10.96.99.164	✓ All certificates are valid

### 2. From the top of the page, select the **Encrypted Nodes** tab.

#### Key Management Server

If your StorageGRID system includes appliance nodes with Full Disk Encryption (FDE) enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID data at rest.

Configuration Details **Encrypted Nodes**

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

The Encrypted Nodes tab lists the appliance nodes in your StorageGRID system that have the **Node Encryption** setting enabled.

Configuration Details Encrypted Nodes

Review the KMS status for all appliance nodes that have node encryption enabled. Address any issues immediately to ensure your data is fully protected. If no KMS exists for a site, select Configuration Details and add a KMS.

**Nodes with Encryption Enabled**

Node Name	Node Type	Site	KMS Display Name	Key UID	Status
SGA-010-096-104-67	Storage Node	Data Center 1	Default KMS	41b0...5c57	✓ Connected to KMS (2021-03-12 10:59:32 MST)

### 3. Review the information in the table for each appliance node.

Column	Description
Node Name	The name of the appliance node.

Column	Description
Node Type	The type of node: Storage, Admin, or Gateway.
Site	The name of the StorageGRID site where the node is installed.
KMS Display 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.  <a href="#">Adding a key management server (KMS)</a>
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, hover your cursor over the cell.  A dash (--) indicates the key UID is unknown, possibly because of a connection issue between the appliance node and the KMS.
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.  <b>Note:</b> You must refresh your web browser to see the new values.

4. 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
- See the recommended actions for these alerts in the instructions for monitoring and troubleshooting StorageGRID.



You must address any issues immediately to ensure that your data is fully protected.

## Related information

[Monitor & troubleshoot](#)

## Editing a key management server (KMS)

You might need to edit the configuration of a key management server, for example, if a certificate is about to expire.

### What you'll need

- You must have reviewed the [considerations and requirements for using a key management server](#).
- If you plan to update the site selected for a KMS, you must have reviewed the [considerations for changing the KMS for a site](#).
- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.

### Steps

1. Select **Configuration > System Settings > Key Management Server**.

The Key Management Server page appears and shows all key management servers that have been configured.

#### Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details    Encrypted Nodes

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

+ Create    Edit    Remove				
KMS Display Name	Key Name	Manages keys for	Hostname	Certificate Status
<input checked="" type="radio"/> Default KMS	test	Sites not managed by another KMS (default KMS)	10.96.99.164	✓ All certificates are valid

2. Select the KMS you want to edit, and select **Edit**.
3. Optionally, update the details in **Step 1 (Enter KMS Details)** of the Edit a Key Management Server wizard.

Field	Description
KMS Display Name	A descriptive name to help you identify this KMS. Must be between 1 and 64 characters.

Field	Description
Key Name	<p>The exact key alias for the StorageGRID client in the KMS. Must be between 1 and 255 characters.</p> <p>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.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  <p>Never attempt to rotate a key by changing the key name (alias) for the KMS. Instead, rotate the key by updating the key version in the KMS software. 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.</p> <p><a href="#">Considerations and requirements for using a key management server</a></p> </div>
Manages keys for	<p>If you are editing a site-specific KMS and you do not already have a default KMS, optionally select <b>Sites not managed by another KMS (default KMS)</b>. This selection converts a site-specific KMS to the default KMS, which will apply to all sites that do not have a dedicated KMS and to any sites added in an expansion.</p> <p><b>Note:</b> If you are editing a site-specific KMS, you cannot select another site. If you are editing the default KMS, you cannot select a specific site.</p>
Port	<p>The port the KMS server uses for Key Management Interoperability Protocol (KMIP) communications. Defaults to 5696, which is the KMIP standard port.</p>
Hostname	<p>The fully qualified domain name or IP address for the KMS.</p> <p><b>Note:</b> The 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.</p>

4. If you are configuring a KMS cluster, select the plus sign **+** to add a hostname for each server in the cluster.
5. Select **Next**.
  - Step 2 (Upload Server Certificate) of the Edit a Key Management Server wizard appears.
6. If you need to replace the server certificate, select **Browse** and upload the new file.
7. Select **Next**.

Step 3 (Upload Client Certificates) of the Edit a Key Management Server wizard appears.

8. If you need to replace the client certificate and the client certificate private key, select **Browse** and upload the new files.
9. Select **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.

10. 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.

A warning dialog box with a yellow header bar containing a warning triangle icon and the word "Warning". The main text area contains the following text: "Confirm force-saving the KMS configuration", "Are you sure you want to save this KMS without testing the external connections?", and "If there is an issue with the configuration, you might not be able to reboot any appliance nodes with node encryption enabled at the affected site, and you might lose access to your data." At the bottom right, there are two buttons: "Cancel" (grey) and "OK" (blue).

The KMS configuration is saved but the connection to the KMS is not tested.

## Removing 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.

### What you'll need

- You must have reviewed the [considerations and requirements for using a key management server](#).
- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.

## 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 > System Settings > Key Management Server**.

The Key Management Server page appears and shows all key management servers that have been configured.

### Key Management Server

If your StorageGRID system includes appliance nodes with node encryption enabled, you can use an external key management server (KMS) to manage the encryption keys that protect your StorageGRID at rest.

Configuration Details    Encrypted Nodes

You can configure more than one KMS (or KMS cluster) to manage the encryption keys for appliance nodes. For example, you can configure one default KMS to manage the keys for all appliance nodes within a group of sites and a second KMS to manage the keys for the appliance nodes at a particular site.

Before adding a KMS:

- Ensure that the KMS is KMIP-compliant.
- Configure StorageGRID as a client in the KMS.
- Enable node encryption for each appliance during appliance installation. You cannot enable node encryption after an appliance is added to the grid and you cannot use a KMS for appliances that do not have node encryption enabled.

For complete instructions, see [administering StorageGRID](#).

KMS Display Name	Key Name	Manages keys for	Hostname	Certificate Status
<input checked="" type="radio"/> Default KMS	test	Sites not managed by another KMS (default KMS)	10.96.99.164	✓ All certificates are valid

2. Select the radio button for the KMS you want to remove, and select **Remove**.
3. Review the considerations in the warning dialog.

## Warning

### Delete KMS Configuration

You can only remove a KMS in these cases:

- You are removing a site-specific KMS for a site that has no appliance nodes with node encryption enabled.
- You are removing the default KMS, but a site-specific KMS already exists for each site with node encryption.

Are you sure you want to delete the Default KMS KMS configuration?

Cancel

OK

#### 4. Select **OK**.

The KMS configuration is removed.

## Managing tenants

As a grid administrator, you create and manage the tenant accounts that S3 and Swift clients use to store and retrieve objects, monitor storage usage, and manage the actions that clients are able to perform using your StorageGRID system.

### What tenant accounts are

Tenant accounts allow client applications that use the Simple Storage Service (S3) REST API or the Swift REST API to store and retrieve objects on StorageGRID.

Each tenant account supports the use of a single protocol, which you specify when you create the account. To store and retrieve objects to a StorageGRID system with both protocols, you must create two tenant accounts: one for S3 buckets and objects, and one for Swift containers and objects. Each tenant account has its own account ID, authorized groups and users, buckets or containers, and objects.

Optionally, you can create additional tenant accounts if you want to segregate the objects stored on your system by different entities. For example, you might set up multiple tenant accounts in 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 simply use S3 buckets and bucket policies to segregate objects between the departments in an enterprise. You do not need to use tenant accounts. See the instructions for implementing S3 client applications 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.

## Creating and configuring tenant accounts

When you create a tenant account, you specify the following information:

- Display name for the tenant account.
- Which client protocol will be used by the tenant account (S3 or Swift).
- For S3 tenant accounts: Whether the tenant account has permission to use platform services with S3 buckets. If you permit tenant accounts to use platform services, you must ensure that the grid is configured to support their use. See “Managing platform services.”
- Optionally, a storage quota for the tenant account—the maximum number of gigabytes, terabytes, or petabytes available for the tenant's objects. If the quota is exceeded, the tenant cannot create new objects.



A tenant's storage quota represents a logical amount (object size), not a physical amount (size on disk).

- If identity federation is enabled for the StorageGRID system, which federated group has Root Access permission to configure the tenant account.
- If single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid's identity source, and the initial password for the tenant's local root user.

After a tenant account is created, you can perform the following tasks:

- **Manage platform services for the grid:** If you enable platform services for tenant accounts, ensure that you understand how platform services messages are delivered and the networking requirements that the use of platform services place on your StorageGRID deployment.
- **Monitor a tenant account's storage usage:** After tenants begin using their accounts, you can use Grid Manager to monitor how much storage each tenant consumes.

If you have set quotas for tenants, you can enable the **Tenant quota usage high** alert to determine if tenants are consuming their quotas. If enabled, this alert is triggered when a tenant has used 90% of its quota. For more information, see the alerts reference in the instructions for monitoring and troubleshooting StorageGRID.

- **Configure client operations:** You can configure if some types of client operations are forbidden.

## Configuring S3 tenants

After an S3 tenant account is created, tenant users can access the Tenant Manager to perform tasks such as the following:

- Setting up identity federation (unless the identity source is shared with the grid) and creating local groups and users
- Managing S3 access keys
- Creating and managing S3 buckets
- Monitoring storage usage
- Using platform services (if enabled)



S3 tenant users can create and manage S3 access key and buckets with the Tenant Manager, but they must use an S3 client application to ingest and manage objects.

## Configuring Swift tenants

After a Swift tenant account is created, the tenant's root user can access the Tenant Manager to perform tasks such as the following:

- Setting up identity federation (unless the identity source is shared with the grid), and creating local groups and users
- Monitoring storage usage



Swift users must have the Root Access permission to access the Tenant Manager. However, the Root Access permission does not allow users to authenticate into the Swift REST API to create containers and ingest objects. Users must have the Swift Administrator permission to authenticate into the Swift REST API.

## Related information

## Creating a tenant account

You must create at least one tenant account to control access to the storage in your StorageGRID system.

### What you'll need

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

### Steps

1. Select **Tenants**.

The Tenant Accounts page appears and lists any existing tenant accounts.

#### Tenant Accounts

View information for each tenant account.

**Note:** 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 and select **View Details**.

The screenshot shows the 'Tenant Accounts' page interface. At the top, there is a toolbar with buttons for '+ Create', 'View details', 'Edit', 'Actions', and 'Export to CSV'. To the right of these buttons is a search box labeled 'Search by Name/ID'. Below the toolbar is a table header with columns: 'Display Name', 'Space Used', 'Quota Utilization', 'Quota', 'Object Count', and 'Sign in'. The table body is empty, showing 'No results found.' At the bottom right, there is a 'Show 20 rows per page' control.

2. Select **Create**.

The Create Tenant Account page appears. The fields included on the page depend on whether single sign-on (SSO) has been enabled for the StorageGRID system.

- If SSO is not being used, the Create Tenant Account page looks like this.

## Create Tenant Account

### Tenant Details

Display Name

Protocol  S3  Swift

Storage Quota (optional)

### Authentication [?](#)

Configure how the tenant account will be accessed.

Uses Own Identity Source

Specify a password for the tenant's local root user.

Username root

Password

Confirm Password

Cancel

Save

- If SSO is enabled, the Create Tenant Account page looks like this.

## Create Tenant Account

### Tenant Details

Display Name

Protocol  S3  Swift

Allow Platform Services

Storage Quota (optional)

### Authentication

Because single sign-on is enabled, the tenant must use the Grid Manager's identity federation service, and no local users can sign in. You must select an existing federated group to have the initial Root Access permission for the tenant.

Uses Own Identity Source

Single sign-on is enabled. The tenant cannot use its own identity source.

Root Access Group

Cancel

Save

### Related information

[Using identity federation](#)

[Configuring single sign-on](#)

### Creating a tenant account if StorageGRID is not using SSO

When you create a tenant account, you specify a name, a client protocol, and optionally a storage quota. If StorageGRID is not using single sign-on (SSO), you must also specify whether the tenant account will use its own identity source and configure the initial password for the tenant's local root user.

### About this task

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 must have imported that federated group into the Grid Manager. You do not need to assign any Grid Manager permissions to this admin group. See the instructions for [managing admin groups](#).

### Steps

1. In the **Display Name** text box, enter a display name for this tenant account.

Display names do not need to be unique. When the tenant account is created, it receives a unique,

numeric Account ID.

2. Select the client protocol that will be used by this tenant account, either **S3** or **Swift**.
3. For S3 tenant accounts, keep the **Allow Platform Services** check box selected unless you do not want this tenant to use platform services for S3 buckets.

If platform services are enabled, a tenant can use features, such as CloudMirror replication, that access external services. You might want to disable the use of these features to limit the amount of network bandwidth or other resources a tenant consumes. See “Managing platform services.”

4. In the **Storage Quota** text box, optionally enter the maximum number of gigabytes, terabytes, or petabytes that you want to make available for this tenant’s objects. Then, select the units from the drop-down list.

Leave this field blank if you want this tenant to have an unlimited quota.



A tenant’s storage quota represents a logical amount (object size), not a physical amount (size on disk). ILM copies and erasure coding do not contribute to the amount of quota used. If the quota is exceeded, the tenant account cannot create new objects.



To monitor each tenant account’s storage usage, select **Usage**. Tenant accounts can also monitor their own storage usage from the Dashboard in the Tenant Manager or with the Tenant Management API. Note that a tenant’s storage usage values might become out of date if nodes are isolated from other nodes in the grid. The totals will be updated when network connectivity is restored.

5. If the tenant will manage its own groups and users, follow these steps.
  - a. Select the **Uses Own Identity Source** check box (default).



If this check box is selected and you want to use identity federation for tenant groups and users, the tenant must configure its own identity source. See the instructions for using tenant accounts.

- b. Specify a password for the tenant’s local root user.
6. If the tenant will use the groups and users configured for the Grid Manager, follow these steps.
  - a. Unselect the **Uses Own Identity Source** check box.
  - b. Do either or both of the following:

- In the Root Access Group field, select an existing federated group from the Grid Manager that should have the initial Root Access permission for the tenant.



If you have adequate permissions, the existing federated groups from the Grid Manager are listed when you click the field. Otherwise, enter the group’s unique name.

- Specify a password for the tenant’s local root user.
7. Click **Save**.

The tenant account is created.

8. Optionally, access the new tenant. Otherwise, go to the step for [accessing the tenant later](#).

If you are...	Do this...
Accessing the Grid Manager on a restricted port	Click <b>Restricted</b> 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>  <ul style="list-style-type: none"> <li>• <i>FQDN_or_Admin_Node_IP</i> is a fully qualified domain name or the IP address of an Admin Node</li> <li>• <i>port</i> is the tenant-only port</li> <li>• <i>20-digit-account-id</i> is the tenant's unique account ID</li> </ul>
Accessing the Grid Manager on port 443 but you did not set a password for the local root user	Click <b>Sign In</b> , and enter the credentials for a user in the Root Access federated group.
Accessing the Grid Manager on port 443 and you set a password for the local root user	Go to the next step to <a href="#">sign in as root</a> .

9. Sign in to the tenant as root:
  - a. From the Configure Tenant Account dialog box, click the **Sign in as root** button.

### Configure Tenant Account

✔ Account **S3 tenant** created successfully.

If you are ready to configure this tenant account, sign in as the tenant's root user. Then, click the links below.

Sign in as root

- [Buckets](#) - Create and manage buckets.
- [Groups](#) - Manage user groups, and assign group permissions.
- [Users](#) - Manage local users, and assign users to groups.

Finish

A green check mark appears on the button, indicating that you are now signed in to the tenant account as the root user.



b. Click 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.

c. Click **Finish**.

10. To access the tenant later:

If you are using...	Do one of these...
Port 443	<ul style="list-style-type: none"><li>• From the Grid Manager, select <b>Tenants</b>, and click <b>Sign in</b> to the right of the tenant name.</li><li>• Enter the tenant's URL in a web browser:  <code>https://FQDN_or_Admin_Node_IP/?accountId=20-digit-account-id/</code><ul style="list-style-type: none"><li>◦ <i>FQDN_or_Admin_Node_IP</i> is a fully qualified domain name or the IP address of an Admin Node</li><li>◦ <i>20-digit-account-id</i> is the tenant's unique account ID</li></ul></li></ul>
A restricted port	<ul style="list-style-type: none"><li>• From the Grid Manager, select <b>Tenants</b>, and click <b>Restricted</b>.</li><li>• Enter the tenant's URL in a web browser:  <code>https://FQDN_or_Admin_Node_IP:port/?accountId=20-digit-account-id</code><ul style="list-style-type: none"><li>◦ <i>FQDN_or_Admin_Node_IP</i> is a fully qualified domain name or the IP address of an Admin Node</li><li>◦ <i>port</i> is the tenant-only restricted port</li><li>◦ <i>20-digit-account-id</i> is the tenant's unique account ID</li></ul></li></ul>

## Related information

[Controlling access through firewalls](#)

[Managing platform services for S3 tenant accounts](#)

[Use a tenant account](#)

## Creating a tenant account if SSO is enabled

When you create a tenant account, you specify a name, a client protocol, and optionally a storage quota. If single sign-on (SSO) is enabled for StorageGRID, you also specify which federated group has Root Access permission to configure the tenant account.

## Steps

1. In the **Display Name** text box, enter a display name for this tenant account.

Display names do not need to be unique. When the tenant account is created, it receives a unique, numeric Account ID.

2. Select the client protocol that will be used by this tenant account, either **S3** or **Swift**.
3. For S3 tenant accounts, keep the **Allow Platform Services** check box selected unless you do not want this tenant to use platform services for S3 buckets.

If platform services are enabled, a tenant can use features, such as CloudMirror replication, that access external services. You might want to disable the use of these features to limit the amount of network bandwidth or other resources a tenant consumes. See “Managing platform services.”

4. In the **Storage Quota** text box, optionally enter the maximum number of gigabytes, terabytes, or petabytes that you want to make available for this tenant’s objects. Then, select the units from the drop-down list.

Leave this field blank if you want this tenant to have an unlimited quota.



A tenant’s storage quota represents a logical amount (object size), not a physical amount (size on disk). ILM copies and erasure coding do not contribute to the amount of quota used. If the quota is exceeded, the tenant account cannot create new objects.



To monitor each tenant account’s storage usage, select **Usage**. Tenant accounts can also monitor their own storage usage from the Dashboard in the Tenant Manager or with the Tenant Management API. Note that a tenant’s storage usage values might become out of date if nodes are isolated from other nodes in the grid. The totals will be updated when network connectivity is restored.

5. Notice that the **Uses Own Identity Source** check box is unchecked and disabled.

Because SSO is enabled, the tenant must use the identity source that was configured for the Grid Manager. No local users can sign in.

6. In the **Root Access Group** field, select an existing federated group from the Grid Manager to have the initial Root Access permission for the tenant.



If you have adequate permissions, the existing federated groups from the Grid Manager are listed when you click the field. Otherwise, enter the group’s unique name.

7. Click **Save**.

The tenant account is created. The Tenant Accounts page appears, and it includes a row for the new tenant.

8. If you are a user in the Root Access group, optionally click the **Sign in** link for the new tenant to immediately access the Tenant Manager, where you can configure the tenant. Otherwise, provide the URL for the **Sign in** link to the tenant account’s administrator. (The URL for a tenant is the fully qualified domain name or IP address of any Admin Node, followed by `/?accountId=20-digit-account-id`.)



An access denied message is displayed if you click **Sign in**, but you do not belong to the Root Access group for the tenant account.

## Related information

[Configuring single sign-on](#)

### Changing the password for a 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.

#### What you'll need

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

#### About this task

If single sign-on (SSO) is enabled for your StorageGRID system, the local root user cannot 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**.

The Tenant Accounts page appears and lists all existing tenant accounts.

#### Tenant Accounts

View information for each tenant account.

**Note:** 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 and select **View Details**.

	Display Name	Space Used	Quota Utilization	Quota	Object Count	Sign in
<input checked="" type="radio"/>	Account01	500.00 KB	0.00%	20.00 GB	100	
<input type="radio"/>	Account02	2.50 MB	0.01%	30.00 GB	500	
<input type="radio"/>	Account03	605.00 MB	4.03%	15.00 GB	31,000	
<input type="radio"/>	Account04	1.00 GB	10.00%	10.00 GB	200,000	
<input type="radio"/>	Account05	0 bytes	—	Unlimited	0	

Buttons: + Create, View details, Edit, Actions, Export to CSV. Search by Name/ID. Show 20 rows per page.

2. Select the tenant account you want to edit.

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. Use the search box to search for a tenant account by display name or tenant ID.

The View Details, Edit, and Actions buttons become enabled.

3. From the **Actions** drop-down, select **Change Root Password**.

## Change Root User Password - Account03

Username root

New Password

Confirm New Password

4. Enter the new password for the tenant account.
5. Select **Save**.

### Related information

[Controlling administrator access to StorageGRID](#)

### Editing a tenant account

You can edit a tenant account to change the display name, change the identity source setting, allow or disallow platform services, or enter a storage quota.

### What you'll need

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

### Steps

1. Select **Tenants**.

The Tenant Accounts page appears and lists all existing tenant accounts.

### Tenant Accounts

View information for each tenant account.

**Note:** 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 and select **View Details**.

	Display Name  	Space Used   	Quota Utilization   	Quota   	Object Count   	Sign in 
<input checked="" type="radio"/>	Account01	500.00 KB	0.00%	20.00 GB	100	
<input type="radio"/>	Account02	2.50 MB	0.01%	30.00 GB	500	
<input type="radio"/>	Account03	605.00 MB	4.03%	15.00 GB	31,000	
<input type="radio"/>	Account04	1.00 GB	10.00%	10.00 GB	200,000	
<input type="radio"/>	Account05	0 bytes	—	Unlimited	0	

Show  rows per page

2. Select the tenant account you want to edit.

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. Use the search box to search for a tenant account by display name or tenant ID.

3. Select **Edit**.

The Edit Tenant Account page appears. This example is for a grid that does not use single sign-on (SSO). This tenant account has not configured its own identity source.

## Edit Tenant Account

### Tenant Details

Display Name

Allow Platform Services

Storage Quota (optional)

Uses Own Identity Source

4. Change the values for the fields as required.

- a. Change the display name for this tenant account.
- b. Change the setting of the **Allow Platform Services** check box to determine whether the tenant account can use platform services for their S3 buckets.



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.

- c. For **Storage Quota**, change the number of maximum number of gigabytes, terabytes, or petabytes available for this tenant's objects, or leave the field blank if you want this tenant to have an unlimited quota.

A tenant's storage quota represents a logical amount (object size), not a physical amount (size on disk). ILM copies and erasure coding do not contribute to the amount of quota used.



To monitor each tenant account's storage usage, select **Usage**. Tenant accounts can also monitor their own usage from the Dashboard in the Tenant Manager or with the Tenant Management API. Note that a tenant's storage usage values might become out of date if nodes are isolated from other nodes in the grid. The totals will be updated when network connectivity is restored.

- d. Change the setting of the **Uses Own Identity Source** check box to determine whether the tenant account will use its own identity source or the identity source that was configured for the Grid Manager.



If the **Uses Own Identity Source** check box is:

- Disabled and checked, 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 unchecked, SSO is enabled for the StorageGRID system. The tenant must use the identity source that was configured for the Grid Manager.

5. Select **Save**.

#### Related information

[Managing platform services for S3 tenant accounts](#)

[Use a tenant account](#)

#### Deleting a tenant account

You can delete a tenant account if you want to permanently remove the tenant's access to the system.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have removed all buckets (S3), containers (Swift), and objects associated with the tenant account.

#### Steps

1. Select **Tenants**.
2. Select the tenant account you want to delete.

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. Use the search box to search for a tenant account by display name or tenant ID.

3. From the **Actions** drop-down, select **Remove**.
4. Select **OK**.

#### Related information

[Controlling administrator access to StorageGRID](#)

#### Managing platform services for S3 tenant accounts

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.

- [What platform services are](#)
- [Networking and ports for platform services](#)
- [Per-site delivery of platform services messages](#)
- [Troubleshooting platform services](#)

## What platform services are

Platform services include CloudMirror replication, event notifications, and the search integration service.

These services allow tenants to use the following functionality with their S3 buckets:

- **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 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 Amazon Simple Notification Service™ (SNS).

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, you must be aware of the following recommendations:

- 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.
- 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.

## Related information

[Use a tenant account](#)

[Configuring Storage proxy settings](#)

[Monitor & troubleshoot](#)

## Networking 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 Simple Notification Service (SNS) messages
- 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
- **443**: For endpoint URIs that begin with https

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

[Configuring Storage proxy settings](#)


[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.

### Troubleshooting 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, a Simple Notification Service 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 are not 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

<input type="checkbox"/>	Display name	Last error	Type	URI	URN
<input type="checkbox"/>	my-endpoint-2	2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es::mydomain/sveloso/_doc
<input type="checkbox"/>	my-endpoint-3	3 days ago	Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1
<input type="checkbox"/>	my-endpoint-5	12 days ago	Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example3
<input type="checkbox"/>	my-endpoint-4		Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example2
<input type="checkbox"/>	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 are not blocked.

## Determining 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 cannot 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, the Total Events (SMTT) alarm is triggered in the Grid Manager. To view the Total Events alarm:

1. Select **Nodes**.
2. Select **site > grid node > Events**.
3. View Last Event at the top of the table.

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

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

The tenant can resubmit the existing values to avoid making unwanted changes.

## Platform services messages cannot 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.

If platform services messages cannot be delivered because of an unrecoverable error, the Total Events (SMTT) alarm is triggered in the Grid Manager.

## 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 Failure Rate chart.



### 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 about troubleshooting platform services endpoints, see the instructions for using tenant accounts.

[Use a tenant account](#)

### Related information

[Monitor & troubleshoot](#)

[Configuring Storage proxy settings](#)

## Configuring S3 and Swift client connections

As a grid administrator, you manage the configuration options that control how S3 and Swift tenants can connect client applications to your StorageGRID system to store and retrieve data. There are a number of different options to meet different client and tenant requirements.

Client applications can store or retrieve objects by connecting to any of the following:

- The Load Balancer service on Admin Nodes or Gateway Nodes, or optionally, the virtual IP address of a high availability (HA) group of Admin Nodes or Gateway Nodes
- The CLB service on Gateway Nodes, or optionally, the virtual IP address of a high availability group of Gateway Nodes



The CLB service is deprecated. Clients configured before the StorageGRID 11.3 release can continue to use the CLB service on Gateway Nodes. All other client applications that depend on StorageGRID to provide load balancing should connect using the Load Balancer service.

- Storage Nodes, with or without an external load balancer

You can optionally configure the following features on your StorageGRID system:

- **Load Balancer service:** You enable clients to use the Load Balancer service by creating load balancer endpoints for client connections. When creating a load balancer endpoint, you specify a port number, whether the endpoint accepts HTTP or HTTPS connections, the type of client (S3 or Swift) that will use the endpoint, and the certificate to be used for HTTPS connections (if applicable).
- **Untrusted Client Network:** You can make the Client Network more secure by configuring it as untrusted. When the Client Network is untrusted, clients can only connect using load balancer endpoints.
- **High availability groups:** You can create an HA group of Gateway Nodes or Admin Nodes to create an active-backup configuration, or you can use round-robin DNS or a third-party load balancer and multiple HA groups to achieve an active-active configuration. Client connections are made using the virtual IP addresses of HA groups.

You can also enable the use of HTTP for clients that connect to StorageGRID either directly to Storage Nodes or using the CLB service (deprecated), and you can configure S3 API endpoint domain names for S3 clients.

### Summary: IP addresses and ports for client connections

Client applications can connect to StorageGRID using the IP address of a grid node and

the port number of a service on that node. If high availability (HA) groups are configured, client applications can connect using the virtual IP address of the HA group.

**About this task**

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. The instructions describe how to find this information in the Grid Manager if load balancer endpoints and high availability (HA) groups are already configured.

Where connection is made	Service that client connects to	IP address	Port
HA group	Load Balancer	Virtual IP address of an HA group	<ul style="list-style-type: none"> <li>• Load balancer endpoint port</li> </ul>
HA group	CLB <b>Note:</b> The CLB service is deprecated.	Virtual IP address of an HA group	Default S3 ports: <ul style="list-style-type: none"> <li>• HTTPS: 8082</li> <li>• HTTP: 8084</li> </ul> Default Swift ports: <ul style="list-style-type: none"> <li>• HTTPS:8083</li> <li>• HTTP:8085</li> </ul>
Admin Node	Load Balancer	IP address of the Admin Node	<ul style="list-style-type: none"> <li>• Load balancer endpoint port</li> </ul>
Gateway Node	Load Balancer	IP address of the Gateway Node	<ul style="list-style-type: none"> <li>• Load balancer endpoint port</li> </ul>
Gateway Node	CLB <b>Note:</b> The CLB service is deprecated.	IP address of the Gateway Node <b>Note:</b> By default, HTTP ports for CLB and LDR are not enabled.	Default S3 ports: <ul style="list-style-type: none"> <li>• HTTPS: 8082</li> <li>• HTTP: 8084</li> </ul> Default Swift ports: <ul style="list-style-type: none"> <li>• HTTPS:8083</li> <li>• HTTP:8085</li> </ul>
Storage Node	LDR	IP address of Storage Node	Default S3 ports: <ul style="list-style-type: none"> <li>• HTTPS: 18082</li> <li>• HTTP: 18084</li> </ul> Default Swift ports: <ul style="list-style-type: none"> <li>• HTTPS: 18083</li> <li>• HTTP:18085</li> </ul>

## Examples

To connect an S3 client 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 an S3 Load Balancer endpoint is 10443, then an S3 client could use the following URL to connect to StorageGRID:

- `https://192.0.2.5:10443`

To connect a Swift client 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 192.0.2.6 and the port number of a Swift Load Balancer endpoint is 10444, then a Swift client could use the following URL to connect to StorageGRID:

- `https://192.0.2.6:10444`

It is possible to configure a DNS name for the IP address that clients use to connect to StorageGRID. Contact your local network administrator.

## Steps

1. Sign in to the Grid Manager using a supported 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. Click **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 Settings > 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 Settings > Load Balancer Endpoints**.

The Load Balancer Endpoints page appears, showing the list of endpoints that have already been configured.

- b. Select an endpoint, and click **Edit endpoint**.

The Edit Endpoint window opens and displays additional details about the endpoint.

- c. Confirm that the endpoint you have selected is configured for use with the correct protocol (S3 or Swift), then click **Cancel**.
- d. Note the port number for the endpoint that you want to use for a client connection.



If the port number is 80 or 443, the endpoint is configured only on Gateway Nodes, since those ports are reserved on Admin Nodes. All other ports are configured on both Gateway Nodes and Admin Nodes.

## Managing load balancing

You can use the StorageGRID load balancing functions to handle ingest and retrieval workloads from S3 and Swift clients. Load balancing maximizes speed and connection capacity by distributing the workloads and connections across multiple Storage Nodes.

You can achieve load balancing in your StorageGRID system in the following ways:

- Use the Load Balancer service, which is installed on Admin Nodes and Gateway Nodes. The Load Balancer service provides Layer 7 load balancing and performs TLS termination of client requests, inspects the requests, and establishes new secure connections to the Storage Nodes. This is the recommended load balancing mechanism.
- Use the Connection Load Balancer (CLB) service, which is installed on Gateway Nodes only. The CLB service provides Layer 4 load balancing and supports link costs.



The CLB service is deprecated.

- Integrate a third-party load balancer. Contact your NetApp account representative for details.

### How load balancing works - Load Balancer service

The Load Balancer service distributes incoming network connections from client applications to Storage Nodes. To enable load balancing, you must configure load balancer endpoints using the Grid Manager.

You can configure load balancer endpoints only for Admin Nodes or Gateway Nodes, since these node types contain the Load Balancer service. You cannot configure endpoints for Storage Nodes or Archive Nodes.

Each load balancer endpoint specifies a port, a protocol (HTTP or HTTPS), a service type (S3 or Swift), and a binding mode. HTTPS endpoints require a server certificate. Binding modes allow you to restrict the accessibility of endpoint ports to:

- Specific high availability (HA) virtual IP addresses (VIPs)
- Specific network interfaces of specific nodes

## Port considerations

Clients can access any of the endpoints you configure on any node running the Load Balancer service, with two exceptions: ports 80 and 443 are reserved on Admin Nodes, so endpoints configured on these ports support load balancing operations only on Gateway Nodes.

If you have remapped any ports, you cannot use the same ports to configure load balancer endpoints. You can create endpoints using remapped ports, but those endpoints will be remapped to the original CLB ports and service, not the Load Balancer service. Follow the steps in the recovery and maintenance instructions for removing port remaps.



The CLB service is deprecated.

## CPU availability

The Load Balancer service on each Admin Node and Gateway Node operates independently when forwarding S3 or Swift 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.

## Related information

[Maintain & recover](#)

## Configuring load balancer endpoints

You can create, edit, and remove load balancer endpoints.

## Creating load balancer endpoints

Each load balancer endpoint specifies a port, a network protocol (HTTP or HTTPS), and a service type (S3 or Swift). If you create an HTTPS endpoint, you must upload or generate a server certificate.

## What you'll need

- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.
- If you have previously remapped ports you intend to use for the Load Balancer service, you must have removed the remaps.



If you have remapped any ports, you cannot use the same ports to configure load balancer endpoints. You can create endpoints using remapped ports, but those endpoints will be remapped to the original CLB ports and service, not the Load Balancer service. Follow the steps in the recovery and maintenance instructions for removing port remaps.



The CLB service is deprecated.

## Steps

1. Select **Configuration > Network Settings > Load Balancer Endpoints**.

The Load Balancer Endpoints page appears.

## Load Balancer Endpoints

Load balancer endpoints define Gateway Node and Admin Node ports that accept and load balance S3 and Swift requests to Storage Nodes. HTTPS endpoint certificates are configured per endpoint.

 Changes to endpoints can take up to 15 minutes to be applied to all nodes.

 Add endpoint port  Edit endpoint  Remove endpoint port

Display name	Port	Using HTTPS
--------------	------	-------------

*No endpoints configured.*

### 2. Select **Add endpoint**.

The Create Endpoint dialog box appears.

Create Endpoint

Display Name

Port

Protocol  HTTP  HTTPS

Endpoint Binding Mode  Global  HA Group VIPs  Node Interfaces

### 3. Enter a display name for the endpoint, which will appear in the list on the Load Balancer Endpoints page.

### 4. Enter a port number, or leave the pre-filled port number as is.

If you enter port number 80 or 443, the endpoint is configured only on Gateway Nodes, since these ports are reserved on Admin Nodes.



Ports used by other grid services are not permitted. See the networking guidelines for a list of ports used for internal and external communications.

### 5. Select **HTTP** or **HTTPS** to specify the network protocol for this endpoint.

### 6. Select an endpoint binding mode.

- **Global** (default): The endpoint is accessible on all Gateway Nodes and Admin Nodes on the specified port number.

## Create Endpoint

Display Name

Port

Protocol  HTTP  HTTPS

Endpoint Binding Mode  Global  HA Group VIPs  Node Interfaces

 This endpoint is currently bound globally. All nodes will use this endpoint unless an endpoint with an overriding binding mode exists for a specific port.

Cancel Save

- **HA Group VIPs:** The endpoint is accessible only through the virtual IP addresses defined for the selected HA groups. Endpoints defined in this mode can reuse the same port number, as long as the HA groups defined by those endpoints do not overlap with each other.

Select the HA groups with the virtual IP addresses where you want the endpoint to appear.

## Create Endpoint

Display Name

Port

Protocol  HTTP  HTTPS

Endpoint Binding Mode  Global  HA Group VIPs  Node Interfaces

Name	Description	Virtual IP Addresses	Interfaces
<input type="checkbox"/> Group1		192.168.5.163	CO-REF-DC1-ADM1:eth0 (preferred Master)
<input type="checkbox"/> Group2		47.47.5.162	CO-REF-DC1-ADM1:eth2 (preferred Master)

Displaying 2 HA groups.

 No HA groups selected. You must select one or more HA Groups; otherwise, this endpoint will act as a globally bound endpoint.

Cancel Save

- **Node Interfaces:** The endpoint is accessible only on the designated nodes and network interfaces. Endpoints defined in this mode can reuse the same port number as long as those interfaces do not overlap with each other.

Select the node interfaces where you want the endpoint to appear.

## Create Endpoint

Display Name

Port

Protocol  HTTP  HTTPS

Endpoint Binding Mode  Global  HA Group VIPs  Node Interfaces

Node	Interface
<input type="checkbox"/> CO-REF-DC1-ADM1	eth0
<input type="checkbox"/> CO-REF-DC1-ADM1	eth1
<input type="checkbox"/> CO-REF-DC1-ADM1	eth2
<input type="checkbox"/> CO-REF-DC1-GW1	eth0
<input type="checkbox"/> CO-REF-DC2-ADM1	eth0
<input type="checkbox"/> CO-REF-DC2-GW1	eth0

 No node interfaces selected. You must select one or more node interfaces; otherwise, this endpoint will act as a globally bound endpoint.

7. Select **Save**.

The Edit Endpoint dialog box appears.

8. Select **S3** or **Swift** to specify the type of traffic this endpoint will serve.

### Edit Endpoint Unsecured Port A (port 10449)

#### Endpoint Service Configuration

Endpoint service type  S3  Swift

9. If you selected **HTTP**, select **Save**.

The unsecured endpoint is created. The table on the Load Balancer Endpoints page lists the endpoint's display name, port number, protocol, and endpoint ID.

10. If you selected **HTTPS** and you want to upload a certificate, select **Upload Certificate**.

## Load Certificate

Upload the PEM-encoded custom certificate, private key, and CA bundle files.

Server Certificate

Certificate Private Key

CA Bundle

Cancel

Save

- a. Browse for the server certificate and the certificate private key.

To enable S3 clients to connect using an S3 API endpoint domain name, use a multi-domain or wildcard certificate that matches all domain names that the client might use to connect to the grid. For example, the server certificate might use the domain name `*.example.com`.

[Configuring S3 API endpoint domain names](#)

- b. Optionally browse for a CA bundle.
- c. Select **Save**.

The PEM-encoded certificate data for the endpoint appears.

11. If you selected **HTTPS** and you want to generate a certificate, select **Generate Certificate**.

## Generate Certificate

Domain 1

IP 1

Subject

Days valid

Cancel

Generate

- a. Enter a domain name or an IP address.

You can use wildcards to represent the fully qualified domain names of all Admin Nodes and Gateway Nodes running the Load Balancer service. For example, `*.sgws.foo.com` uses the `*` wildcard to represent `gn1.sgws.foo.com` and `gn2.sgws.foo.com`.

## Configuring S3 API endpoint domain names

- b. Select **+** to add any other domain names or IP addresses.

If you are using high availability (HA) groups, add the domain names and IP addresses of the HA virtual IPs.

- c. Optionally, enter an X.509 subject, also referred to as the Distinguished Name (DN), to identify who owns the certificate.
- d. Optionally, select the number of days the certificate is valid. The default is 730 days.
- e. Select **Generate**.

The certificate metadata and the PEM-encoded certificate data for the endpoint appear.

12. Click **Save**.

The endpoint is created. The table on the Load Balancer Endpoints page lists the endpoint's display name, port number, protocol, and endpoint ID.

### Related information

[Maintain & recover](#)

[Network guidelines](#)

[Managing high availability groups](#)

[Managing untrusted Client Networks](#)

### Editing load balancer endpoints

For an unsecured (HTTP) endpoint, you can change the endpoint service type between S3 and Swift. For a secured (HTTPS) endpoint, you can edit the endpoint service type and view or change the security certificate.

### What you'll need

- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.

### Steps

1. Select **Configuration > Network Settings > Load Balancer Endpoints**.

The Load Balancer Endpoints page appears. The existing endpoints are listed in the table.

Endpoints with certificates that will expire soon are identified in the table.



- Change the endpoint binding mode.  
For a secured (HTTPS) endpoint, you can:
- Change the endpoint service type between S3 and Swift.
- Change the endpoint binding mode.
- View the security certificate.
- Upload or generate a new security certificate when the current certificate is expired or about to expire.

Select a tab to display detailed information about the default StorageGRID server certificate or a CA signed certificate that was uploaded.



To change the protocol for an existing endpoint, for example from HTTP to HTTPS, you must create a new endpoint. Follow the instructions for creating load balancer endpoints, and select the desired protocol.

5. Click **Save**.

### Related information

[Creating load balancer endpoints](#)

### Removing load balancer endpoints

If you no longer need a load balancer endpoint, you can remove it.

#### What you'll need

- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.

#### Steps

1. Select **Configuration > Network Settings > Load Balancer Endpoints**.

The Load Balancer Endpoints page appears. The existing endpoints are listed in the table.

#### Load Balancer Endpoints

Load balancer endpoints define Gateway Node and Admin Node ports that accept and load balance S3 and Swift requests to Storage Nodes. HTTPS endpoint certificates are configured per endpoint.

<input type="button" value="+ Add endpoint"/> <input type="button" value="✎ Edit endpoint"/> <input type="button" value="✕ Remove endpoint"/>			
	Display name	Port	Using HTTPS
<input type="radio"/>	Unsecured Endpoint 5	10444	No
<input checked="" type="radio"/>	Secured Endpoint 1	10443	Yes

Displaying 2 endpoints.

2. Select the radio button to the left of the endpoint you want to remove.
3. Click **Remove endpoint**.

A confirmation dialog box appears.

## Warning

Remove Endpoint

Are you sure you want to remove endpoint 'Secured Endpoint 1'?

Cancel

OK

4. Click **OK**.

The endpoint is removed.

### How load balancing works - CLB service

The Connection Load Balancer (CLB) service on Gateway Nodes is deprecated. The Load Balancer service is now the recommended load balancing mechanism.

The CLB service uses Layer 4 load balancing to distribute incoming TCP network connections from client applications to the optimal Storage Node based on availability, system load, and the administrator-configured link cost. When the optimal Storage Node is chosen, the CLB service establishes a two-way network connection and forwards the traffic to and from the chosen node. The CLB does not consider the Grid Network configuration when directing incoming network connections.

To view information about the CLB service, select **Support > Tools > Grid Topology**, and then expand a Gateway Node until you can select **CLB** and the options below it.

The screenshot shows the Grid Topology interface. On the left, a tree view under 'StorageGRID Webscale Deployment' shows 'Data Center 1' expanded to 'DC1-G1-98-161', which contains 'SSM', 'CLB', 'HTTP', 'Events', and 'Resources'. The 'CLB' node is highlighted with a blue box. On the right, the 'Overview: Summary - DC1-G1-98-161' page is displayed, showing a 'Storage Capacity' table.

Storage Capacity	
Storage Nodes Installed:	N/A
Storage Nodes Readable:	N/A
Storage Nodes Writable:	N/A
Installed Storage Capacity:	N/A
Used Storage Capacity:	N/A
Used Storage Capacity for Data:	N/A
Used Storage Capacity for Metadata:	N/A
Usable Storage Capacity:	N/A

If you choose to use the CLB service, you should consider configuring link costs for your StorageGRID system.

### Related information

[What link costs are](#)

[Updating link costs](#)

## Managing untrusted Client Networks

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 (see the information about external communications in the network guidelines).

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.

### 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 Untrusted Client Networks page, 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 service 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 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 Amazon Web Services.

## Related information

[Network guidelines](#)

[Configuring load balancer endpoints](#)

### Specifying a node's Client Network is untrusted

If you are using a Client Network, you can specify whether each node's Client Network is trusted or untrusted. You can also specify the default setting for new nodes added in an expansion.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.
- 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.



Existing client connections might fail if load balancer endpoints have not been configured.

## Steps

1. Select **Configuration > Network Settings > Untrusted Client Network**.

The Untrusted Client Networks page appears.

This page lists all nodes in your StorageGRID system. The Unavailable Reason column includes an entry if the Client Network on the node must be trusted.

### Untrusted Client Networks

If you are using a Client Network, you can specify whether a node trusts inbound traffic from the Client Network. If the Client Network is untrusted, the node only accepts inbound traffic on ports configured as [load balancer endpoints](#).

#### Set New Node Default

This setting applies to new nodes expanded into the grid.

New Node Client Network Default  Trusted  Untrusted

#### Select Untrusted Client Network Nodes

Select nodes that should have untrusted Client Network enforcement.

<input type="checkbox"/>	Node Name	Unavailable Reason
<input type="checkbox"/>	DC1-ADM1	
<input type="checkbox"/>	DC1-G1	
<input type="checkbox"/>	DC1-S1	
<input type="checkbox"/>	DC1-S2	
<input type="checkbox"/>	DC1-S3	
<input type="checkbox"/>	DC1-S4	

Client Network untrusted on 0 nodes.

Save

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:** 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 page to change the setting for a specific new node.



This setting does not affect the existing nodes in your StorageGRID system.

3. In the **Select Untrusted Client Network Nodes** section, select the nodes that should allow client connections only on explicitly configured load balancer endpoints.

You can select or unselect the check box in the title to select or unselect all nodes.

#### 4. Click **Save**.

The new firewall rules are immediately added and enforced. Existing client connections might fail if load balancer endpoints have not been configured.

### Related information

[Configuring load balancer endpoints](#)

### Managing high availability groups

High availability (HA) groups can be used to provide highly available data connections for S3 and Swift clients. HA groups can also be used to provide highly available connections to the Grid Manager and the Tenant Manager.

- [What an HA group is](#)
- [How HA groups are used](#)
- [Configuration options for HA groups](#)
- [Creating a high availability group](#)
- [Editing a high availability group](#)
- [Removing a high availability group](#)

#### What an HA group is

High availability groups use virtual IP addresses (VIPs) to provide active-backup access to Gateway Node or Admin Node services.

An HA group consists of one or more network interfaces on Admin Nodes and Gateway Nodes. When creating an HA group, you select network interfaces belonging to the Grid Network (eth0) or the Client Network (eth2). All interfaces in an HA group must be within the same network subnet.

An HA group maintains one or more virtual IP addresses that are added to the active interface in the group. If the active interface becomes unavailable, the virtual IP addresses are moved to another interface. This 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.

The active interface in an HA group is designated as the Master. All other interfaces are designated as Backup. To view these designations, select **Nodes > node > Overview**.

Overview

Hardware

Network

Storage

Load Balancer

Events

Tasks

Node Information 

Name	DC1-ADM1
Type	Admin Node
ID	711b7b9b-8d24-4d9f-877a-be3fa3ac27e8
Connection State	 Connected
Software Version	11.4.0 (build 20200515.2346.8edcbbf)
HA Groups	Fabric Pools, Master
IP Addresses	192.168.2.208, 10.224.2.208, 47.47.2.208, 47.47.4.219 <a href="#">Show more</a> 

When creating an HA group, you specify one interface to be the preferred Master. The preferred Master is the active interface unless a failure occurs that causes the VIP addresses to be reassigned to a Backup interface. When the failure is resolved, the VIP addresses are automatically moved back to the preferred Master.

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 failure of the CLB service (deprecated) or services for the Grid Manager or the Tenant Manager.

If the HA group includes interfaces from more than two nodes, the active interface might move to any other node's interface during failover.

### How HA groups are used

You might want to use high availability (HA) groups for several reasons.

- 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 and Swift 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 and the CLB service (deprecated).

Purpose of HA group	Add nodes of this type to the HA group
Access to Grid Manager	<ul style="list-style-type: none"> <li>• Primary Admin Node (<b>preferred Master</b>)</li> <li>• Non-primary Admin Nodes</li> </ul> <p><b>Note:</b> The primary Admin Node must be the preferred Master. Some maintenance procedures can only be performed from the primary Admin Node.</p>
Access to Tenant Manager only	<ul style="list-style-type: none"> <li>• Primary or non-primary Admin Nodes</li> </ul>
S3 or Swift client access — Load Balancer service	<ul style="list-style-type: none"> <li>• Admin Nodes</li> <li>• Gateway Nodes</li> </ul>
S3 or Swift client access — CLB service	<ul style="list-style-type: none"> <li>• Gateway Nodes</li> </ul> <p><b>Note:</b> The CLB service is deprecated.</p>

### Limitations of using HA groups with Grid Manager or Tenant Manager

The failure of services for the Grid Manager or the Tenant Manager does not trigger failover within the HA group.

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 cannot be performed when the primary Admin Node is unavailable. During failover, you can use the Grid Manager to monitor your StorageGRID system.

### Limitations of using HA groups with the CLB service



The failure of the CLB service does not trigger failover within the HA group.



The CLB service is deprecated.

### Configuration options for HA groups

The following diagrams provide examples of different ways you can configure HA groups. Each option has advantages and disadvantages.



When creating multiple overlapping HA groups as shown in the Active-Active HA example, the total throughput scales with the number of nodes and HA groups. With three or more nodes and three or more HA groups, you also gain the ability to continue operations using any of the VIPs even during maintenance procedures that require you to take a node offline.

The table summarizes the benefits of each HA configuration shown in the diagram.

Configuration	Advantages	Disadvantages
Active-Backup HA	<ul style="list-style-type: none"> <li>• Managed by StorageGRID with no external dependencies.</li> <li>• Fast failover.</li> </ul>	<ul style="list-style-type: none"> <li>• Only one node in an HA group is active. At least one node per HA group will be idle.</li> </ul>
DNS Round Robin	<ul style="list-style-type: none"> <li>• Increased aggregate throughput.</li> <li>• No idle hosts.</li> </ul>	<ul style="list-style-type: none"> <li>• Slow failover, which could depend on client behavior.</li> <li>• Requires configuration of hardware outside of StorageGRID.</li> <li>• Needs a customer-implemented health check.</li> </ul>

Configuration	Advantages	Disadvantages
Active-Active	<ul style="list-style-type: none"> <li>Traffic is distributed across multiple HA groups.</li> <li>High aggregate throughput that scales with the number of HA groups.</li> <li>Fast failover.</li> </ul>	<ul style="list-style-type: none"> <li>More complex to configure.</li> <li>Requires configuration of hardware outside of StorageGRID.</li> <li>Needs a customer-implemented health check.</li> </ul>

### Creating a high availability group

You can create one or more high availability (HA) groups to provide highly available access to the services on Admin Nodes or Gateway Nodes.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

#### About this task

An interface must meet the following conditions to be included in an HA group:

- The interface must be for a Gateway Node or an Admin Node.
- The interface must belong to the Grid Network (eth0) or the Client Network (eth2).
- The interface must be configured with fixed or static IP addressing, not with DHCP.

#### Steps

- Select **Configuration > Network Settings > High Availability Groups**.

The High Availability Groups page appears.

#### High Availability Groups

High availability (HA) groups allow multiple nodes to participate in an active-backup group. HA groups maintain virtual IP addresses on the active node and switch to a backup node automatically if a node fails.



- Click **Create**.

The Create High Availability Group dialog box appears.

- Type a name and, if desired, a description for the HA group.
- Click **Select Interfaces**.

The Add Interfaces to High Availability Group dialog box appears. The table lists eligible nodes, interfaces, and IPv4 subnets.

## Add Interfaces to High Availability Group

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

Add to HA group	Node Name	Interface	IPv4 Subnet	Unavailable Reason
	g140-g1	eth0	172.16.0.0/21	This IP address is not in the same subnet as the selected interfaces
	g140-g1	eth2	47.47.0.0/21	This IP address is not in the same subnet as the selected interfaces
	g140-g2	eth0	172.16.0.0/21	This IP address is not in the same subnet as the selected interfaces
	g140-g2	eth2	47.47.0.0/21	This IP address is not in the same subnet as the selected interfaces
	g140-g3	eth0	172.16.0.0/21	This IP address is not in the same subnet as the selected interfaces
<input checked="" type="checkbox"/>	g140-g3	eth2	192.168.0.0/21	
	g140-g4	eth0	172.16.0.0/21	This IP address is not in the same subnet as the selected interfaces
<input checked="" type="checkbox"/>	g140-g4	eth2	192.168.0.0/21	

There are 2 interfaces selected.

Cancel Apply

An interface does not appear in the list if its IP address is assigned by DHCP.

5. In the **Add to HA group** column, select the check box for the interface you want to add to the HA group.

Note the following guidelines for selecting interfaces:

- You must select at least one interface.
- If you select more than one interface, all of the interfaces must be on either the Grid Network (eth0) or on the Client Network (eth2).
- All interfaces must be in the same subnet or in subnets with a common prefix.

IP addresses will be restricted to the smallest subnet (the one with the largest prefix).

- If you select interfaces on different types of nodes, and a failover occurs, only the services common to the selected nodes will be available on the virtual IPs.
  - Select two or more Admin Nodes for HA protection of the Grid Manager or the Tenant Manager.
  - Select two or more Admin Nodes, Gateway Nodes, or both for HA protection of the Load Balancer service.
  - Select two or more Gateway Nodes for HA protection of the CLB service.



The CLB service is deprecated.

## Add Interfaces to High Availability Group

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

Add to HA group	Node Name	Interface	IPv4 Subnet	Unavailable Reason
<input checked="" type="checkbox"/>	DC1-ADM1	eth0	10.96.100.0/23	
<input checked="" type="checkbox"/>	DC1-G1	eth0	10.96.100.0/23	
<input checked="" type="checkbox"/>	DC2-ADM1	eth0	10.96.100.0/23	

There are 3 interfaces selected.

**Attention:** You have selected nodes of different types that run different services. If a failover occurs, only the services common to all node types will be available on the virtual IPs.

Cancel

Apply

### 6. Click **Apply**.

The interfaces you selected are listed in the Interfaces section of the Create High Availability Group page. By default, the first interface in the list is selected as the Preferred Master.

## Create High Availability Group

### High Availability Group

Name

Description

### Interfaces

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

Select Interfaces

Node Name	Interface	IPv4 Subnet	Preferred Master
g140-g1	eth2	47.47.0.0/21	<input checked="" type="radio"/>
g140-g2	eth2	47.47.0.0/21	<input type="radio"/>

Displaying 2 interfaces.

### Virtual IP Addresses

Virtual IP Subnet: 47.47.0.0/21. All virtual IP addresses must be within this subnet. There must be at least 1 and no more than 10 virtual IP addresses.

Virtual IP Address 1



Cancel

Save

- If you want a different interface to be the preferred Master, select that interface in the **Preferred Master** column.

The preferred Master is the active interface unless a failure occurs that causes the VIP addresses to be reassigned to a Backup interface.



If the HA group provides access to the Grid Manager, you must select an interface on the primary Admin Node to be the preferred Master. Some maintenance procedures can only be performed from the primary Admin Node.

- In the Virtual IP Addresses section of the page, enter one to 10 virtual IP addresses for the HA group. Click the plus sign (+) to add multiple IP addresses.

You must provide at least one IPv4 address. Optionally, you can specify additional IPv4 and IPv6 addresses.

IPv4 addresses must be within the IPv4 subnet shared by all of the member interfaces.

9. Click **Save**.

The HA Group is created, and you can now use the configured virtual IP addresses.

### Related information

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

[Install VMware](#)

[Install Ubuntu or Debian](#)

[Managing load balancing](#)

### Editing a high availability group

You can edit a high availability (HA) group to change its name and description, add or remove interfaces, or add or update a virtual IP address.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

### About this task

Some of the reasons for editing an HA group include the following:

- Adding an interface to an existing group. The interface IP address must be within the same subnet as other interfaces already assigned to the group.
- Removing an interface from an HA group. For example, you cannot start a site or node decommission procedure if a node's interface for the Grid Network or the Client Network is used in an HA group.

### Steps

1. Select **Configuration > Network Settings > High Availability Groups**.

The High Availability Groups page appears.

#### High Availability Groups

High availability (HA) groups allow multiple nodes to participate in an active-backup group. HA groups maintain virtual IP addresses on the active node and switch to a backup node automatically if a node fails.

	Name	Description	Virtual IP Addresses	Interfaces
<input type="radio"/>	HA Group 1		47.47.4.219	g140-adm1:eth2 (preferred Master) g140-g1:eth2
<input type="radio"/>	HA Group 2		47.47.4.218 47.47.4.217	g140-g1:eth2 (preferred Master) g140-g2:eth2

Displaying 2 HA groups.

2. Select the HA group you want to edit, and click **Edit**.

The Edit High Availability Group dialog box appears.

3. Optionally, update the group's name or description.
4. Optionally, click **Select Interfaces** to change the interfaces for the HA Group.

The Add Interfaces to High Availability Group dialog box appears.

### Add Interfaces to High Availability Group

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

Add to HA group	Node Name	Interface	IPv4 Subnet	Unavailable Reason
	g140-g1	eth0	172.16.0.0/21	This IP address is not in the same subnet as the selected interfaces
	g140-g1	eth2	47.47.0.0/21	This IP address is not in the same subnet as the selected interfaces
	g140-g2	eth0	172.16.0.0/21	This IP address is not in the same subnet as the selected interfaces
	g140-g2	eth2	47.47.0.0/21	This IP address is not in the same subnet as the selected interfaces
	g140-g3	eth0	172.16.0.0/21	This IP address is not in the same subnet as the selected interfaces
<input checked="" type="checkbox"/>	g140-g3	eth2	192.168.0.0/21	
	g140-g4	eth0	172.16.0.0/21	This IP address is not in the same subnet as the selected interfaces
<input checked="" type="checkbox"/>	g140-g4	eth2	192.168.0.0/21	

There are 2 interfaces selected.

An interface does not appear in the list if its IP address is assigned by DHCP.

5. Select or unselect the check boxes to add or remove interfaces.

Note the following guidelines for selecting interfaces:

- You must select at least one interface.
- If you select more than one interface, all of the interfaces must be on either the Grid Network (eth0) or on the Client Network (eth2).
- All interfaces must be in the same subnet or in subnets with a common prefix.

IP addresses will be restricted to the smallest subnet (the one with the largest prefix).

- If you select interfaces on different types of nodes, and a failover occurs, only the services common to the selected nodes will be available on the virtual IPs.
  - Select two or more Admin Nodes for HA protection of the Grid Manager or the Tenant Manager.
  - Select two or more Admin Nodes, Gateway Nodes, or both for HA protection of the Load Balancer service.
  - Select two or more Gateway Nodes for HA protection of the CLB service.



The CLB service is deprecated.

6. Click **Apply**.

The interfaces you selected are listed in the Interfaces section of the page. By default, the first interface in the list is selected as the Preferred Master.

## Edit High Availability Group 'HA Group - Admin Nodes'

### High Availability Group

Name

Description

### Interfaces

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

Select Interfaces

Node Name	Interface	IPv4 Subnet	Preferred Master
DC1-ADM1	eth0	10.96.100.0/23	<input checked="" type="radio"/>
DC2-ADM1	eth0	10.96.100.0/23	<input type="radio"/>

Displaying 2 interfaces.

### Virtual IP Addresses

Virtual IP Subnet: 10.96.100.0/23. All virtual IP addresses must be within this subnet. There must be at least 1 and no more than 10 virtual IP addresses.

Virtual IP Address 1



Cancel

Save

7. If you want a different interface to be the preferred Master, select that interface in the **Preferred Master** column.

The preferred Master is the active interface unless a failure occurs that causes the VIP addresses to be reassigned to a Backup interface.



If the HA group provides access to the Grid Manager, you must select an interface on the primary Admin Node to be the preferred Master. Some maintenance procedures can only be performed from the primary Admin Node.

8. Optionally, update the virtual IP addresses for the HA group.

You must provide at least one IPv4 address. Optionally, you can specify additional IPv4 and IPv6 addresses.

IPv4 addresses must be within the IPv4 subnet shared by all of the member interfaces.

9. Click **Save**.

The HA Group is updated.

### Removing a high availability group

You can remove a high availability (HA) group that you are no longer using.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

#### About this task

If you remove an HA group, any S3 or Swift clients that are configured to use one of the group's virtual IP addresses will no longer be able to connect to StorageGRID. To prevent client disruptions, you should update all affected S3 or Swift 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 or using DHCP.

#### Steps

1. Select **Configuration > Network Settings > High Availability Groups**.

The High Availability Groups page appears.

#### High Availability Groups

High availability (HA) groups allow multiple nodes to participate in an active-backup group. HA groups maintain virtual IP addresses on the active node and switch to a backup node automatically if a node fails.

<input type="button" value="+ Create"/> <input type="button" value="Edit"/> <input type="button" value="Remove"/>				
	Name	Description	Virtual IP Addresses	Interfaces
<input type="radio"/>	HA Group 1		47.47.4.219	g140-adm1:eth2 (preferred Master) g140-g1:eth2
<input type="radio"/>	HA Group 2		47.47.4.218 47.47.4.217	g140-g1:eth2 (preferred Master) g140-g2:eth2

Displaying 2 HA groups.

2. Select the HA group you want to remove, and click **Remove**.

The Delete High Availability Group warning appears.

## Warning

### Delete High Availability Group

Are you sure you want to delete High Availability Group 'HA group 1'?

Cancel

OK

3. Click **OK**.

The HA group is removed.

## Configuring S3 API endpoint domain names

To support S3 virtual hosted-style requests, you must use the Grid Manager to configure the list of endpoint domain names that S3 clients connect to.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have confirmed that a grid upgrade is not in progress.



Do not 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 tasks:

- 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 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.

The certificate a client uses for HTTPS connections depends on how the client connects to the grid:

- If a client connects using the Load Balancer service, it uses the certificate for a specific load balancer endpoint.



Each load balancer endpoint has its own certificate, and each endpoint can be configured to recognize different endpoint domain names.

- If the client connects to a Storage Node or to the CLB service on a Gateway Node, the client uses a grid custom server certificate that has been updated to include all required endpoint domain names.



The CLB service is deprecated.

## Steps

1. Select **Configuration > Network Settings > Domain Names**.

The Endpoint Domain Names page appears.

Endpoint Domain Names

### Virtual Hosted-Style Requests

Enable support of S3 virtual hosted-style requests by specifying API endpoint domain names. Support is disabled if this list is empty. Examples: s3.example.com, s3.example.co.uk, s3-east.example.com

Endpoint 1	<input type="text" value="s3.example.com"/>	✕
Endpoint 2	<input type="text"/>	+ ✕

[Save](#)

2. Using the (+) icon to add additional fields, enter the list of S3 API endpoint domain names in the **Endpoint** fields.

If this list is empty, support for S3 virtual hosted-style requests is disabled.

3. Click **Save**.
4. Ensure that the server certificates that clients use match the required endpoint domain names.
  - For clients that use the Load Balancer service, update the certificate associated with the load balancer endpoint that the client connects to.
  - For clients that connect directly to Storage Nodes or that use the CLB service on Gateway Nodes, update the custom server certificate for the grid.
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.

## Related information

[Use S3](#)

[Viewing IP addresses](#)

[Creating a high availability group](#)

[Configuring a custom server certificate for connections to the Storage Node or the CLB service](#)

[Configuring load balancer endpoints](#)

## Enabling HTTP for client communications

By default, client applications use the HTTPS network protocol for all connections to Storage Nodes or to the deprecated CLB service on Gateway Nodes. You can optionally enable HTTP for these connections, for example, when testing a non-production grid.

### What you'll need

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

### About this task

Complete this task only if S3 and Swift clients need to make HTTP connections directly to Storage Nodes or to the deprecated CLB service on Gateway Nodes.

You do not need to complete this task 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 the information on configuring load balancer endpoints for more information.

See [Summary: IP addresses and ports for client connections](#) to learn which ports S3 and Swift clients use when connecting to Storage Nodes or to the deprecated CLB service using HTTP or HTTPS



Be careful when enabling HTTP for a production grid because requests will be sent unencrypted.

### Steps

1. Select **Configuration > System Settings > Grid Options**.
2. In the Network Options section, select the **Enable HTTP Connection** check box.

#### Network Options



3. Click **Save**.

### Related information

[Configuring load balancer endpoints](#)

[Use S3](#)

## Controlling which client operations are permitted

You can select the Prevent Client Modification grid option to deny specific HTTP client operations.

### What you'll need

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

### About this task

Prevent Client Modification is a system wide setting. When the Prevent Client Modification option is selected, the following requests are denied:

#### • S3 REST API

- Delete Bucket requests
- Any requests to modify an existing object's data, user-defined metadata, or S3 object tagging



This setting does not apply to buckets with versioning enabled. Versioning already prevents modifications to object data, user-defined metadata, and object tagging.


#### • Swift REST API

- Delete Container requests
- Requests to modify any existing object. For example, the following operations are denied: Put Overwrite, Delete, Metadata Update, and so on.

### Steps

1. Select **Configuration > System Settings > Grid Options**.
2. In the Network Options section, select the **Prevent Client Modification** check box.

#### Network Options

A screenshot of the 'Network Options' configuration page. The 'Prevent Client Modification' checkbox is checked and highlighted with a yellow rounded rectangle. Below it, the 'Enable HTTP Connection' checkbox is unchecked. At the bottom, the 'Network Transfer Encryption' section shows two radio button options: 'AES128-SHA' (unselected) and 'AES256-SHA' (selected). Each option has a small blue question mark icon to its right.

Prevent Client Modification

Enable HTTP Connection

Network Transfer Encryption  AES128-SHA  AES256-SHA

3. Click **Save**.

## Managing StorageGRID networks and connections

You can use the Grid Manager to configure and manage StorageGRID networks and connections.

See [Configuring S3 and Swift client connections](#) to learn how to connect S3 or Swift clients.

- [Guidelines for StorageGRID networks](#)
- [Viewing IP addresses](#)
- [Supported ciphers for outgoing TLS connections](#)
- [Changing network transfer encryption](#)
- [Configuring server certificates](#)
- [Configuring Storage proxy settings](#)
- [Configuring Admin proxy settings](#)
- [Managing traffic classification policies](#)
- [What link costs are](#)

### Guidelines for StorageGRID networks

StorageGRID 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.



To modify or add a network for a grid node, see the recovery and maintenance instructions. For more information about network topology, see the networking instructions.

#### 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 and Swift 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 grid node requires a dedicated network interface, IP address, subnet mask, and gateway for each network it is assigned to.
- A grid node cannot 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.

### Related information

[Maintain & recover](#)

[Network guidelines](#)

### Viewing 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 into the grid node at the command line and perform various maintenance procedures.

### What you'll need

You must be signed in to the Grid Manager using a supported browser.

### About this task

For information on changing IP addresses, see the recovery and maintenance instructions.

### Steps

1. Select **Nodes** > *grid node* > **Overview**.
2. Click **Show more** to the right of the IP Addresses title.

The IP addresses for that grid node are listed in a table.

**Node Information** 

**Name** SGA-lab11  
**Type** Storage Node  
**ID** 0b583829-6659-4c6e-b2d0-31461d22ba67

**Connection State**  Connected  
**Software Version** 11.4.0 (build 20200527.0043.61839a2)  
**IP Addresses** 192.168.4.138, 10.224.4.138, 169.254.0.1 [Show less](#) 

Interface	IP Address
eth0	192.168.4.138
eth0	fd20:331:331:0:2a0:98ff:fea1:831d
eth0	fe80::2a0:98ff:fea1:831d
eth1	10.224.4.138
eth1	fd20:327:327:0:280:e5ff:fe43:a99c
eth1	fd20:8b1e:b255:8154:280:e5ff:fe43:a99c
eth1	fe80::280:e5ff:fe43:a99c
hic2	192.168.4.138
hic4	192.168.4.138
mtc1	10.224.4.138
mtc2	169.254.0.1

## Related information

[Maintain & recover](#)

## 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 or Swift client applications.



TLS configuration options such as protocol versions, ciphers, key exchange algorithms, and MAC algorithms are not configurable in StorageGRID. Contact your NetApp account representative if you have specific requests about these settings.

### Supported TLS 1.2 cipher suites

The following TLS 1.2 cipher suites are supported:

- TLS\_ECDHE\_RSA\_WITH\_AES\_128\_GCM\_SHA256

- TLS\_ECDHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384
- TLS\_ECDHE\_ECDSA\_WITH\_AES\_128\_GCM\_SHA256
- TLS\_ECDHE\_ECDSA\_WITH\_AES\_256\_GCM\_SHA384
- TLS\_ECDHE\_RSA\_WITH\_CHACHA20\_POLY1305
- TLS\_ECDHE\_ECDSA\_WITH\_CHACHA20\_POLY1305
- TLS\_RSA\_WITH\_AES\_128\_GCM\_SHA256
- TLS\_RSA\_WITH\_AES\_256\_GCM\_SHA384

### Supported TLS 1.3 cipher suites

The following TLS 1.3 cipher suites are supported:

- TLS\_AES\_256\_GCM\_SHA384
- TLS\_CHACHA20\_POLY1305\_SHA256
- TLS\_AES\_128\_GCM\_SHA256

### Changing network transfer encryption

The StorageGRID system uses Transport Layer Security (TLS) to protect internal control traffic between grid nodes. The Network Transfer Encryption option sets the algorithm used by TLS to encrypt control traffic between grid nodes. This setting does not affect data encryption.

#### What you'll need

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


#### About this task

By default, network transfer encryption uses the AES256-SHA algorithm. Control traffic can also be encrypted using the AES128-SHA algorithm.

#### Steps

1. Select **Configuration > System Settings > Grid Options**.
2. In the Network Options section, change Network Transfer Encryption to **AES128-SHA** or **AES256-SHA** (default).

#### Network Options



3. Click **Save**.

## Configuring server certificates

You can customize the server certificates used by the StorageGRID system.

The StorageGRID system uses security certificates for multiple distinct purposes:

- **Management Interface Server Certificates:** Used to secure access to the Grid Manager, the Tenant Manager, the Grid Management API, and the Tenant Management API.
- **Storage API Server Certificates:** Used to secure access to the Storage Nodes and Gateway Nodes, which API client applications use to upload and download object data.

You can use the default certificates created during installation, or you can replace either, or both, of these default types of certificates with your own custom certificates.

### Supported types of custom server certificate

The StorageGRID system supports custom server certificates encrypted with RSA or ECDSA (Elliptic Curve Digital Signature Algorithm).

For more information on how StorageGRID secures client connections for the REST API, see the S3 or Swift implementation guides.

### Certificates for load balancer endpoints

StorageGRID manages the certificates used for load balancer endpoints separately. To configure load balancer certificates, see the instructions for configuring load balancer endpoints.

### Related information

[Use S3](#)

[Use Swift](#)

[Configuring load balancer endpoints](#)

### Configuring a custom server certificate for the Grid Manager and the Tenant Manager

You can replace the default StorageGRID server certificate with a single custom server certificate that allows users to access the Grid Manager and the Tenant Manager without encountering security warnings.

#### 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 server certificate and corresponding private key.

Because a single custom server 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 root CA certificate in the web browser they will use to access the Grid Manager and the Tenant Manager.



To ensure that operations are not disrupted by a failed server certificate, the **Expiration of server certificate for Management Interface** alert and the legacy Management Interface Certificate Expiry (MCEP) alarm are both triggered when this server certificate is about to expire. As required, you can view the number of days until the current service certificate expires by selecting **Support > Tools > Grid Topology**. Then, select **primary Admin Node > CMN > Resources**.



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

### Steps

1. Select **Configuration > Network Settings > Server Certificates**.
2. In the Management Interface Server Certificate section, click **Install Custom Certificate**.
3. Upload the required server certificate files:
  - **Server Certificate**: The custom server certificate file (.crt).
  - **Server 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 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.
4. Click **Save**.

The custom server certificates are used for all subsequent new client connections.

Select a tab to display detailed information about the default StorageGRID server certificate or a CA signed certificate that was uploaded.



After uploading a new certificate, allow up to one day for any related certificate expiration alerts (or legacy alarms) to clear.

5. Refresh the page to ensure the web browser is updated.

### Restoring the default server certificates for the Grid Manager and the Tenant Manager

You can revert to using the default server certificates for the Grid Manager and the Tenant Manager.

### Steps

1. Select **Configuration > Network Settings > Server Certificates**.
2. In the Manage Interface Server Certificate section, click **Use Default Certificates**.

3. Click **OK** in the confirmation dialog box.

When you restore the default server certificates, the custom server certificate files you configured are deleted and cannot be recovered from the system. The default server certificates are used for all subsequent new client connections.

4. Refresh the page to ensure the web browser is updated.

### Configuring a custom server certificate for connections to the Storage Node or the CLB service

You can replace the server certificate that is used for S3 or Swift client connections to the Storage Node or to the CLB service (deprecated) on Gateway Node. The replacement custom server certificate is specific to your organization.

#### 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, users might also need to install the root CA certificate in the S3 or Swift API client they will use to access the system, depending on the root Certificate Authority (CA) you are using.



To ensure that operations are not disrupted by a failed server certificate, the **Expiration of server certificate for Storage API Endpoints** alert and the legacy Storage API Service Endpoints Certificate Expiry (SCEP) alarm are both triggered when the root server certificate is about to expire. As required, you can view the number of days until the current service certificate expires by selecting **Support > Tools > Grid Topology**. Then, select **primary Admin Node > CMN > Resources**.

The custom certificates are only used if clients connect to StorageGRID using the deprecated CLB service on Gateway Nodes, or if they connect directly to Storage Nodes. S3 or Swift clients that connect to StorageGRID using the Load Balancer service on Admin Nodes or Gateway Nodes use the certificate configured for the load balancer endpoint.



The **Expiration of load balancer endpoint certificate** alert is triggered for load balancer endpoints that will expire soon.

#### Steps

1. Select **Configuration > Network Settings > Server Certificates**.
2. In the Object Storage API Service Endpoints Server Certificate section, click **Install Custom Certificate**.
3. Upload the required server certificate files:
  - **Server Certificate**: The custom server certificate file (.crt).
  - **Server 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 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.

4. Click **Save**.

The custom server certificate is used for all subsequent new API client connections.

Select a tab to display detailed information about the default StorageGRID server certificate or a CA signed certificate that was uploaded.



After uploading a new certificate, allow up to one day for any related certificate expiration alerts (or legacy alarms) to clear.

5. Refresh the page to ensure the web browser is updated.

### Related information

[Use S3](#)

[Use Swift](#)

[Configuring S3 API endpoint domain names](#)

### Restoring the default server certificates for the S3 and Swift REST API endpoints

You can revert to using the default server certificates for the S3 and Swift REST API endpoints.

### Steps

1. Select **Configuration > Network Settings > Server Certificates**.
2. In the Object Storage API Service Endpoints Server Certificate section, click **Use Default Certificates**.
3. Click **OK** in the confirmation dialog box.

When you restore the default server certificates for the object storage API endpoints, the custom server certificate files you configured are deleted and cannot be recovered from the system. The default server certificates are used for all subsequent new API client connections.

4. Refresh the page to ensure the web browser is updated.

### Copying the StorageGRID system's CA certificate

StorageGRID uses an internal Certificate Authority (CA) to secure internal traffic. This certificate does not change if you upload your own certificates.

### What you'll need

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

### About this task

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 > Network Settings > Server Certificates**.
2. In the **Internal CA Certificate** section, select all of the certificate text.

You must include -----BEGIN CERTIFICATE----- and -----END CERTIFICATE----- in your selection.

### Internal CA Certificate

StorageGRID uses an internal Certificate Authority (CA) to secure internal traffic. This certificate does not change if you upload your own certificates.

To export the internal CA certificate, copy all of the certificate text (starting with -----BEGIN CERTIFICATE----- and ending with -----END CERTIFICATE-----), and save it as a .pem file.



3. Right-click the selected text, and select **Copy**.
4. Paste the copied certificate into a text editor.
5. Save the file with the extension `.pem`.

For example: `storagegrid_certificate.pem`

## Configuring StorageGRID certificates for FabricPool

For S3 clients that perform strict hostname validation and do not 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.

### What you'll need

- You must have specific access permissions.
- You must be signed in to the Grid Manager using a supported 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 the instructions for configuring StorageGRID for FabricPool.



The separate Connection Load Balancer (CLB) service on Gateway Nodes is deprecated and no longer recommended for use with 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 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.

### Related information

[Configure StorageGRID for FabricPool](#)

### Generating a self-signed server certificate for the management interface

You can use a script to generate a self-signed server certificate for management API clients that require strict hostname validation.

### What you'll need

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

### About this task

In production environments, you should use a certificate that is signed by a known Certificate Authority (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.

### 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 certificate 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 API 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 365
```

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 > Server Certificates > Management Interface Server Certificate**.

### 7. Configure your management API client to use the public certificate you copied. Include the BEGIN and END tags.

## Configuring Storage proxy settings

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.

### What you'll need

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


### About this task

You can configure the settings for a single Storage proxy.

### Steps

1. Select **Configuration > Network Settings > Proxy Settings**.

The Storage Proxy Settings page appears. By default, **Storage** is selected in the sidebar menu.



2. Select the **Enable Storage Proxy** check box.

The fields for configuring a Storage proxy appear.

#### Storage Proxy Settings

If you are using platform services or Cloud Storage Pools, you can configure a non-transparent proxy server between Storage Nodes and the external S3 endpoints.

Enable Storage Proxy

Protocol  HTTP  SOCKS5

Hostname

Port (optional)

3. Select the protocol for the non-transparent 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.

You can leave this field blank if you use the default port for the protocol: 80 for HTTP or 1080 for SOCKS5.

6. Click **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.

#### After you finish

If you need to disable a Storage proxy, deselect the **Enable Storage Proxy** check box, and click **Save**.

#### Related information

[Networking and ports for platform services](#)

[Manage objects with ILM](#)

## Configuring Admin proxy settings

If you send AutoSupport messages using HTTP or HTTPS, you can configure a non-transparent proxy server between Admin Nodes and technical support (AutoSupport).

### What you'll need

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

### About this task


You can configure the settings for a single Admin proxy.

### Steps

1. Select **Configuration > Network Settings > Proxy Settings**.

The Admin Proxy Settings page appears. By default, **Storage** is selected in the sidebar menu.

2. From the sidebar menu, select **Admin**.



3. Select the **Enable Admin Proxy** check box.

#### Admin Proxy Settings

If you send AutoSupport messages using HTTPS or HTTP, you can configure a non-transparent proxy server between Admin Nodes and technical support.

Enable Admin Proxy	<input checked="" type="checkbox"/>
Hostname	<input type="text" value="myproxy.example.com"/>
Port	<input type="text" value="8080"/>
Username (optional)	<input type="text" value="root"/>
Password (optional)	<input type="password" value="••••••••"/>

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 the proxy username.

Leave this field blank if your proxy server does not require a username.

7. Optionally, enter the proxy password.

Leave this field blank if your proxy server does not require a password.

8. Click **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 proxy, deselect the **Enable Admin Proxy** check box, and click **Save**.

## Related information

[Specifying the protocol for AutoSupport messages](#)

## Managing traffic classification policies

To enhance your quality-of-service (QoS) offerings, you can create traffic classification policies to identify and monitor different types of network traffic. These policies can assist with traffic limiting and monitoring.

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 and optional limits

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
- Tenants
- Subnets (IPv4 subnets containing the client)
- Endpoints (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.

Optionally, you can set limits for a policy based on the following parameters:

- Aggregate Bandwidth In
- Aggregate Bandwidth Out
- Concurrent Read Requests
- Concurrent Write Requests
- Per-Request Bandwidth In
- Per-Request Bandwidth Out
- Read Request Rate
- Write Requests Rate



You can create policies to limit aggregate bandwidth or to limit per-request bandwidth. However, StorageGRID cannot limit both types of bandwidth at the same time. Aggregate bandwidth limits might impose an additional minor performance impact on non-limited traffic.

### Traffic limiting

When you have created traffic classification policies, traffic is limited according to the type of rules and limits you set. 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. 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.

### Using 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.

Traffic classification limits are implemented per load balancer. If traffic is distributed simultaneously across multiple load balancers, the total maximum rates are a multiple of the rate limits you specify.

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	Performance	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

### Creating traffic classification policies

You create traffic classification policies if you want to monitor, and optionally limit, network traffic by bucket, tenant, IP subnet, or load balancer endpoint. Optionally, you can set limits for a policy based on bandwidth, the number of concurrent requests, or the request rate.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.
- You must have created any load balancer endpoints you want to match.
- You must have created any tenants you want to match.

## Steps

1. Select **Configuration > Network Settings > Traffic Classification**.

The Traffic Classification Policies page appears.

### Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

<span>+ Create</span> <span>Edit</span> <span>Remove</span> <span>Metrics</span>			
	Name	Description	ID
<i>No policies found.</i>			

2. Click **Create**.

The Create Traffic Classification Policy dialog box appears.

## Create Traffic Classification Policy

### Policy

Name 

Description

### Matching Rules

Traffic that matches any rule is included in the policy.

Type	Inverse Match	Match Value
------	---------------	-------------

*No matching rules found.*

### Limits (Optional)

Type	Value	Units
------	-------	-------

*No limits found.*

Cancel

Save

3. In the **Name** field, enter a name for the policy.

Enter a descriptive name so you can recognize the policy.

4. Optionally, add a description for the policy in the **Description** field.

For example, describe what this traffic classification policy applies to and what it will limit.

5. Create one or more matching rules for the policy.

Matching rules control which entities will be affected by this traffic classification policy. For example, select Tenant if you want this policy to apply to the network traffic for a specific tenant. Or select Endpoint if you want this policy to apply to the network traffic on a specific load balancer endpoint.

- a. Click **Create** in the **Matching Rules** section.

The Create Matching Rule dialog box appears.

## Create Matching Rule

### Matching Rules

Type  -- Choose One -- 

Match Value  Choose type before providing match value

Inverse Match 

b. From the **Type** drop-down, select the type of entity to be included in the matching rule.

c. In the **Match Value** field, enter a match value based on the type of entity you chose.

- **Bucket:** Enter a bucket name.
- **Bucket Regex:** Enter a regular expression that will be 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.

- **CIDR:** Enter an IPv4 subnet, in CIDR notation, that matches the desired subnet.
- **Endpoint:** Select an endpoint from the list of existing endpoints. These are the load balancer endpoints you defined on the Load Balancer Endpoints page.
- **Tenant:** Select a tenant from the list of existing tenants. Tenant matching is based on the ownership of the bucket being accessed. Anonymous access to a bucket matches the tenant that owns the bucket.

d. If you want to match all network traffic *except* traffic consistent with the Type and Match Value just defined, select the **Inverse** check box. Otherwise, leave the check box unselected.

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**.



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.

e. Click **Apply**.

The rule is created and is listed in the Matching Rules table.

<span>+ Create</span> <span>Edit</span> <span>Remove</span>		
Type	Inverse Match	Match Value
Bucket Regex	<input checked="" type="checkbox"/>	control-ld+

Displaying 1 matching rule.

#### Limits (Optional)

<span>+ Create</span> <span>Edit</span> <span>Remove</span>			
Type	Value	Type	Units
No limits found.			

Cancel Save

f. Repeat these steps for each rule you want to create for the policy.

 Traffic that matches any rule is handled by the policy.

6. Optionally, create limits for the policy.

 Even if you do not create limits, StorageGRID collects metrics so that you can monitor network traffic that matches the policy.

a. Click **Create** in the **Limits** section.

The Create Limit dialog box appears.

### Create Limit

#### Limits (Optional)

Type 

Aggregate rate limits in use. Per-request rate limits are not available. 

Value 

Cancel Apply

b. From the **Type** drop-down, select the type of limit you want to apply to the policy.

In the following list, **In** refers to traffic from S3 or Swift clients to the StorageGRID load balancer, and **Out** refers to traffic from the load balancer to S3 or Swift clients.

- Aggregate Bandwidth In
- Aggregate Bandwidth Out
- Concurrent Read Requests
- Concurrent Write Requests
- Per-Request Bandwidth In
- Per-Request Bandwidth Out
- Read Request Rate
- Write Requests Rate



You can create policies to limit aggregate bandwidth or to limit per-request bandwidth. However, StorageGRID cannot limit both types of bandwidth at the same time. 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

c. In the **Value** field, enter a numerical value for the type of limit you chose.

The expected units are shown when you select a limit.

d. Click **Apply**.

The limit is created and is listed in the Limits table.

Type	Inverse Match	Match Value
<input checked="" type="radio"/> Bucket Regex	<input checked="" type="checkbox"/>	control-ld+

Displaying 1 matching rule.

#### Limits (Optional)

Type	Value	Units
<input checked="" type="radio"/> Aggregate Bandwidth Out	10000000000	Bytes/Second

Displaying 1 limit.

Cancel Save

e. Repeat these steps for each limit you want to add to the policy.

For example, if you want to create a 40 Gbps bandwidth limit for an SLA tier, create an Aggregate Bandwidth In limit and an Aggregate Bandwidth Out limit and set each one to 40 Gbps.



To convert megabytes per second to gigabits per second, multiply by eight. For example, 125 MB/s is equivalent to 1,000 Mbps or 1 Gbps.

7. When you are finished creating rules and limits, click **Save**.

The policy is saved and is listed in the Traffic Classification Policies table.

#### Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

Name	Description	ID
<input type="radio"/> ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/> Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b

Displaying 2 traffic classification policies.

S3 and Swift client traffic is now handled according to the traffic classification policies. You can view traffic charts and verify that the polices are enforcing the traffic limits you expect.

#### Related information

[Managing load balancing](#)

[Viewing network traffic metrics](#)

## Editing a 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

### Steps

1. Select **Configuration > Network Settings > 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.

<a href="#">+ Create</a>	<a href="#">✎ Edit</a>	<a href="#">✕ Remove</a>	<a href="#">📊 Metrics</a>
	Name	Description	ID
<input type="radio"/>	ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/>	Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bdc894b

Displaying 2 traffic classification policies.

2. Select the radio button to the left of the policy you want to edit.
3. Click **Edit**.

The Edit Traffic Classification Policy dialog box appears.

## Edit Traffic Classification Policy "Fabric Pools"

### Policy

Name 

Fabric Pools

Description (optional)

Monitor Fabric Pools

### Matching Rules

Traffic that matches any rule is included in the policy.

Type	Inverse Match	Match Value
<input checked="" type="radio"/> CIDR		10.10.152.0/24

Displaying 1 matching rule.

### Limits (Optional)

Type	Value	Units
------	-------	-------

No limits found.

Cancel

Save

4. Create, edit, or remove matching rules and limits as needed.
  - a. To create a matching rule or limit, click **Create**, and follow the instructions for creating a rule or creating a limit.
  - b. To edit a matching rule or limit, select the radio button for the rule or limit, click **Edit** in the **Matching Rules** section or the **Limits** section, and follow the instructions for creating a rule or creating a limit.
  - c. To remove a matching rule or limit, select the radio button for the rule or limit, and click **Remove**. Then, click **OK** to confirm that you want to remove the rule or limit.
5. When you are finished creating or editing a rule or a limit, click **Apply**.
6. When you are finished editing the policy, click **Save**.

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.

### Deleting a traffic classification policy

If you no longer need a traffic classification policy, you can delete it.

## What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

## Steps

1. Select **Configuration > Network Settings > 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.

	Name	Description	ID
<input type="radio"/>	ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/>	Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bdbc894b

Displaying 2 traffic classification policies.

2. Select the radio button to the left of the policy you want to delete.
3. Click **Remove**.

A Warning dialog box appears.



4. Click **OK** to confirm that you want to delete the policy.

The policy is deleted.

## Viewing network traffic metrics

You can monitor network traffic by viewing the graphs that are available from the Traffic Classification Policies page.

## What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access 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 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 Settings > 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	
<input type="radio"/> ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574	
<input checked="" type="radio"/> Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bdc894b	

Displaying 2 traffic classification policies.

2. Select the radio button to the left of the policy you want to view metrics for.
3. Click **Metrics**.

A new browser window opens, and the Traffic Classification Policy graphs appear. The graphs display metrics only for the traffic that matches the selected policy.

You can select other policies to view by using the **policy** pull-down.



The following graphs are included on the web page.


- Load Balancer 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.


- Load Balancer Request Completion Rate: This graph provides a 3-minute moving average of the number of completed 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.
- 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.
- 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.
- Write Request Rate by Object Size: This heatmap provides a 3-minute moving average of the rate at which write requests are completed based on object size. In this context, write requests refer only to PUT requests.
- Read Request Rate by Object Size: This heatmap provides a 3-minute moving average of the rate at which read requests are completed based on object size. In this context, read requests refer only to GET requests.

The colors in the heatmap indicate the relative frequency of an object size within an individual graph. The cooler colors (for example, purple and blue) indicate lower relative rates, and the warmer colors (for example, orange and red) indicate higher relative rates.

4. Hover the cursor over a line graph to see a pop-up of values on a specific part of the graph.



5. Hover the cursor over a heatmap 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.



6. Use the **Policy** pull-down in the upper left to select a different policy.

The graphs for the selected policy appear.

7. Alternatively, access the graphs from the **Support** menu.
  - a. Select **Support > Tools > Metrics**.
  - b. In the **Grafana** section of the page, select **Traffic Classification Policy**.
  - c. Select the policy from the pull-down on the upper left of the page.

Traffic classification policies are identified by their ID. Policy IDs are listed on the Traffic Classification Policies page.

8. Analyze the graphs to determine how often the policy is limiting traffic and whether you need to adjust the policy.

### Related information

[Monitor & troubleshoot](#)

### What link costs are


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.

- 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 CLB service on Gateway Nodes to direct client connections.



The CLB service is deprecated.

The diagram shows a three site grid that has link costs configured between sites:



- The CLB service on Gateway Nodes equally distribute 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 D2 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.

### Related information

[How load balancing works - CLB service](#)

### Updating link costs

You can update the link costs between data center sites to reflect latency between sites.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Grid Topology Page Configuration permission.

### Steps

1. Select **Configuration > Network Settings > Link Cost**.



## Link Cost

Updated: 2021-03-29 12:28:41 EDT

Site Names (1 - 2 of 2)

Site ID	Site Name	Actions
10	Data Center 1	
20	Data Center 2	

Show 50 Records Per Page  Previous « 1 » Next

### Link Costs

Link Source	Link Destination	Actions
10	20	

2. Select a site under **Link Source** and enter a cost value between 0 and 100 under **Link Destination**.

You cannot change the link cost if the source is the same as the destination.

To cancel changes, click **Revert**.

3. Click **Apply Changes**.

## Configuring AutoSupport

The AutoSupport feature enables your StorageGRID system to send health and status messages to technical support. Using AutoSupport can significantly speed 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 messages to be sent to one additional destination.

### Information included in AutoSupport messages

AutoSupport messages include information such as the following:

- StorageGRID software version
- Operating system version
- System-level and location-level attribute information
- Recent alerts and alarms (legacy system)
- Current status of all grid tasks, including historical data
- Events information as listed on the **Nodes > Grid Node > Events** page
- Admin Node database usage
- Number of lost or missing objects
- Grid configuration settings

- NMS entities
- Active ILM policy
- Provisioned grid specification file
- Diagnostic metrics

You can enable the AutoSupport feature and the individual AutoSupport options when you first install StorageGRID, or you can enable them later. If AutoSupport is not enabled, a message appears on the Grid ManagerDashboard. The message includes a link to the AutoSupport configuration page.

The AutoSupport feature is disabled. You should enable AutoSupport to allow StorageGRID to send health and status messages to technical support for proactive monitoring and troubleshooting.



You can select the “x” symbol  to close the message. The message will not appear again until your browser cache is cleared, even if AutoSupport remains disabled.

### Using 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.

You must enable AutoSupport if you want to use the Active IQ dashboards and functionality on the NetApp Support site.

[Active IQ Digital Advisor Documentation](#)

### Accessing AutoSupport settings

You configure AutoSupport using the Grid Manager (**Support > Tools > AutoSupport**). The **AutoSupport** page has two tabs: **Settings** and **Results**.

## AutoSupport

The AutoSupport feature enables your StorageGRID system to send periodic and event-driven health and status messages to technical support to allow proactive monitoring and troubleshooting. StorageGRID AutoSupport also enables the use of Active IQ for predictive recommendations.

Settings Results

---

### Protocol Details

Protocol ?  HTTPS  HTTP  SMTP

NetApp Support Certificate Validation ?

---

### AutoSupport Details

Enable Weekly AutoSupport ?

Enable Event-Triggered AutoSupport ?

Enable AutoSupport on Demand ?

---

### Additional AutoSupport Destination

Enable Additional AutoSupport Destination ?

## Protocols for sending AutoSupport messages

You can choose one of three protocols for sending AutoSupport messages:

- HTTPS
- HTTP
- SMTP

If you send AutoSupport messages using HTTPS or HTTP, you can configure a non-transparent proxy server between Admin Nodes and technical support.

If you use SMTP as the protocol for AutoSupport messages, you must configure an SMTP mail server.

## AutoSupport options

You can use any combination of the following options to send AutoSupport messages to technical support:

- **Weekly:** Automatically send AutoSupport messages once per week. Default setting: Enabled.
- **Event-triggered:** Automatically send AutoSupport messages every hour or when significant system events occur. Default setting: Enabled.
- **On Demand:** Allow technical support to request that your StorageGRID system send AutoSupport messages automatically, which is useful when they are actively working an issue (requires HTTPS AutoSupport transmission protocol). Default setting: Disabled.
- **User-triggered:** Manually send AutoSupport messages at any time.

## Related information

## Specifying the protocol for AutoSupport messages

You can use one of three protocols for sending AutoSupport messages.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access or Other Grid Configuration permission.
- If you will use the HTTPS or HTTP protocol for sending AutoSupport messages, you must have provided outbound internet access to the primary Admin Node, either directly or using a proxy server (inbound connections not required).
- If you will use the HTTPS or HTTP protocol and you want to use a proxy server, you must have configured an Admin proxy server.
- If you will use SMTP as the protocol for AutoSupport messages, you must have configured an SMTP mail server. The same mail server configuration is used for alarm email notifications (legacy system).

### About this task

AutoSupport messages can be sent using any of the following protocols:

- **HTTPS:** This is the default and recommended setting for new installations. The HTTPS protocol uses port 443. If you want to enable the AutoSupport on Demand feature, you must use the HTTPS protocol.
- **HTTP:** This protocol is not secure, unless it is used in a trusted environment where the proxy server converts to HTTPS when sending data over the internet. The HTTP protocol uses port 80.
- **SMTP:** Use this option if you want AutoSupport messages to be emailed. If you use SMTP as the protocol for AutoSupport messages, you must configure an SMTP mail server on the Legacy Email Setup page (**Support > Alarms (legacy) > Legacy Email Setup**).



SMTP was the only protocol available for AutoSupport messages before the StorageGRID 11.2 release. If you installed an earlier version of StorageGRID initially, SMTP might be the selected protocol.


The protocol you set is used for sending all types of AutoSupport messages.

### Steps

1. Select **Support > Tools > AutoSupport**.

The AutoSupport page appears, and the **Settings** tab is selected.

2. Select the protocol you want to use to send AutoSupport messages.



### 3. Select your choice for **Netapp Support Certificate Validation**.

- Use NetApp support certificate (default): Certificate validation ensures that the transmission of AutoSupport messages is secure. The NetApp support certificate is already installed with the StorageGRID software.
- Do not verify certificate: 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. Select **Save**.

All weekly, user-triggered, and event-triggered messages are sent using the selected protocol.

## Related information

[Configuring Admin proxy settings](#)

## Enabling AutoSupport on Demand

AutoSupport on Demand can assist in solving issues that technical support is actively working on. When you enable AutoSupport on Demand, technical support can request that AutoSupport messages be sent without the need for your intervention.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access or Other Grid Configuration permission.
- You must have enabled weekly AutoSupport messages.
- You must have set the transport protocol to HTTPS.

### About this task

When you enable this feature, technical support can request that your StorageGRID system send AutoSupport messages automatically. Technical support can also set the polling time interval for AutoSupport on Demand

queries.

Technical support cannot enable or disable AutoSupport on Demand.

## Steps

1. Select **Support > Tools > AutoSupport**.

The AutoSupport page appears with the **Settings** tab selected.

2. Select the HTTPS radio button in the **Protocol Details** section of the page.

The screenshot shows the AutoSupport configuration interface. At the top, there are two tabs: 'Settings' (selected) and 'Results'. Below the tabs is the 'Protocol Details' section, which includes a 'Protocol' dropdown menu with three options: 'HTTPS' (selected and highlighted with a yellow box), 'HTTP', and 'SMTP'. Below this is a 'NetApp Support Certificate Validation' dropdown menu with the option 'Use NetApp support certificate'. The 'AutoSupport Details' section contains three checkboxes: 'Enable Weekly AutoSupport' (checked and highlighted with a yellow box), 'Enable Event-Triggered AutoSupport' (unchecked), and 'Enable AutoSupport on Demand' (checked and highlighted with a yellow box). Below these is the 'Additional AutoSupport Destination' section, which has an unchecked checkbox for 'Enable Additional AutoSupport Destination'. At the bottom of the form are two buttons: 'Save' (highlighted in blue) and 'Send User-Triggered AutoSupport'.

3. Select the **Enable Weekly AutoSupport** check box.
4. Select the **Enable AutoSupport on Demand** check box.
5. Select **Save**.

AutoSupport on Demand is enabled, and technical support can send AutoSupport on Demand requests to StorageGRID.

## Disabling weekly AutoSupport messages

By default, the StorageGRID system is configured to send an AutoSupport message to NetApp Support once a week.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access or Other Grid Configuration permission.

### About this task

To determine when the weekly AutoSupport message is sent, see the **Next Scheduled Time** under **Weekly AutoSupport** on the **AutoSupport > Results** page.

Settings

Results

### Weekly AutoSupport

Next Scheduled Time  2021-02-12 00:20:00 EST

Most Recent Result  Idle (NetApp Support)

Last Successful Time  N/A (NetApp Support)

You can disable the automatic sending of an AutoSupport message at any time.

### Steps

1. Select **Support > Tools > AutoSupport**.

The AutoSupport page appears with the **Settings** tab selected.

2. Clear the **Enable Weekly AutoSupport** check box.

Settings

Results

### Protocol Details

Protocol   HTTPS  HTTP  SMTP

NetApp Support Certificate Validation  Use NetApp support certificate

### AutoSupport Details

Enable Weekly AutoSupport 

Enable Event-Triggered AutoSupport 

AutoSupport On Demand can only be enabled when the protocol is HTTPS and Weekly AutoSupport is enabled. When you enable AutoSupport on Demand, technical support can request that your StorageGRID system send AutoSupport messages automatically.

### Additional AutoSupport Destination

Enable Additional AutoSupport Destination 

Save

Send User-Triggered AutoSupport

3. Select **Save**.

### Disabling event-triggered AutoSupport messages

By default, the StorageGRID system is configured to send an AutoSupport message to NetApp Support when an important alert or other significant system event occurs.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

- You must have the Root Access or Other Grid Configuration permission.

### About this task

You can disable event-triggered AutoSupport messages at any time.



Event-triggered AutoSupport messages are also suppressed when you suppress email notifications system wide. (Select **Configuration > System Settings > Display Options**. Then, select **Notification Suppress All**.)

### Steps

1. Select **Support > Tools > AutoSupport**.

The AutoSupport page appears with the **Settings** tab selected.

2. Clear the **Enable Event-Triggered AutoSupport** check box.

The screenshot shows the 'AutoSupport Settings' page. At the top, there are two tabs: 'Settings' (selected) and 'Results'. Below the tabs is the 'Protocol Details' section, which includes a 'Protocol' dropdown menu with radio buttons for 'HTTPS' (selected), 'HTTP', and 'SMTP'. Below this is a 'NetApp Support Certificate Validation' dropdown menu with the option 'Use NetApp support certificate'. The 'AutoSupport Details' section contains two checkboxes: 'Enable Weekly AutoSupport' (unchecked) and 'Enable Event-Triggered AutoSupport' (checked, highlighted with a yellow box). Below these checkboxes is a light blue informational box that reads: 'AutoSupport On Demand can only be enabled when the protocol is HTTPS and Weekly AutoSupport is enabled. When you enable AutoSupport on Demand, technical support can request that your StorageGRID system send AutoSupport messages automatically.' At the bottom of the page is the 'Additional AutoSupport Destination' section, which includes an unchecked checkbox for 'Enable Additional AutoSupport Destination'. At the very bottom, there are two buttons: 'Save' (highlighted in blue) and 'Send User-Triggered AutoSupport'.

3. Select **Save**.

### Manually triggering an AutoSupport message

To assist technical support in troubleshooting issues with your StorageGRID system, you can manually trigger an AutoSupport message to be sent.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access or Other Grid Configuration permission.

### Steps

1. Select **Support > Tools > AutoSupport**.

The AutoSupport page appears with the **Settings** tab selected.

## 2. Select **Send User-Triggered AutoSupport**.

StorageGRID attempts to send an AutoSupport message to technical support. 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 message again.



After sending an User-triggered AutoSupport message, refresh the AutoSupport page in your browser after 1 minute to access the most recent results.

## Adding an additional AutoSupport destination

When you enable AutoSupport, health and status messages are sent to NetApp support. You can specify one additional destinations for all AutoSupport messages.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access or Other Grid Configuration permission.

### About this task

To verify or change the protocol used to send AutoSupport messages, see the instructions for specifying an AutoSupport protocol.



You cannot use the SMTP protocol to send AutoSupport messages to an additional destination.

## [Specifying the protocol for AutoSupport messages](#)

### Steps

#### 1. Select **Support > Tools > AutoSupport**.

The AutoSupport page appears with the **Settings** tab selected.

#### 2. Select **Enable additional AutoSupport destination**.

The Additional AutoSupport Destination fields appear.

### Additional AutoSupport Destination

Enable Additional AutoSupport Destination

Hostname

Port

Certificate Validation

You are not using a TLS certificate to secure the connection to the additional AutoSupport destination.

Save

Send User-Triggered AutoSupport

3. Enter the server hostname or IP address of an additional AutoSupport destination server.



You can enter only one additional destination.

4. Enter the port used to connect to an additional AutoSupport destination server (default is port 80 for HTTP or port 443 for HTTPS).
5. To send your AutoSupport messages with certificate validation, select **Use custom CA bundle** in the **Certificate Validation** drop-down. Then, do one of the following:
  - Use an editing tool to copy and paste all the contents of each of the PEM-encoded CA certificate files into the **CA bundle** field, concatenated in certificate chain order. You must include `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` in your selection.

### Additional AutoSupport Destination

Enable Additional AutoSupport Destination

Hostname

Port

Certificate Validation

CA Bundle

Browse

- Select **Browse**, navigate to the file containing the certificates, and then select **Open** to upload the file. Certificate validation ensures that the transmission of AutoSupport messages is secure.
6. To send your AutoSupport messages without certificate validation, select **Do not verify certificate** in the

**Certificate Validation** drop-down.

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.

A caution message appears: "You are not using a TLS certificate to secure the connection to the additional AutoSupport destination."

7. Select **Save**.

All future weekly, event-triggered, and user-triggered AutoSupport messages will be sent to the additional destination.

## Sending E-Series AutoSupport messages through StorageGRID

You can send E-Series SANtricity System Manager AutoSupport messages to technical support through a StorageGRID Admin Node rather than the storage appliance's management port.

### What you'll need

- You are signed into the Grid Manager using a supported web browser.
- You have the Storage Appliance Administrator permission or Root Access permission.



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager.

### About this task

E-Series AutoSupport messages contain details of the storage hardware and are more specific than other AutoSupport messages sent by the StorageGRID system.

Configure a special proxy server address in SANtricity System Manager to cause the AutoSupport messages to be transmitted through a StorageGRID Admin Node without the use of the appliance's management port. AutoSupport messages transmitted in this way respect the Preferred Sender and Admin proxy settings which may have been configured in the Grid Manager.

If you want to configure the Admin proxy server in Grid Manager, see the instructions for configuring Admin proxy settings.

### Configuring Admin proxy settings




This procedure is only for configuring a StorageGRID proxy server for E-Series AutoSupport messages. For additional details on E-Series AutoSupport configuration information, see the E-Series documentation center.

[NetApp E-Series Systems Documentation Center](#)

### 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.

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.

Configure AutoSupport Delivery Method

Select AutoSupport dispatch delivery method...

HTTPS

HTTP

Email

**HTTPS delivery settings** Show destination address

Connect to support team...

Directly ?

via Proxy server ?

Host address ?

tunnel-host

Port number ?

10225

My proxy server requires authentication

via Proxy auto-configuration script (PAC) ?

Save Test Configuration Cancel

6. Select **HTTPS** for the delivery method.



The certificate that enables the HTTPS protocol 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 messages.

9. Enter `10225` for the **Port number**.

`10225` is the port number on the StorageGRID proxy server that receives AutoSupport messages 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.”

### Troubleshooting AutoSupport messages

If an attempt to send an AutoSupport message fails, the StorageGRID system takes different actions depending on the type of AutoSupport message. You can check the status of AutoSupport messages by selecting **Support > Tools > AutoSupport > Results**.



Event-triggered AutoSupport messages are suppressed when you suppress email notifications system wide. (Select **Configuration > System Settings > Display Options**. Then, select **Notification Suppress All**.)

When the AutoSupport message fails to send, “Failed” appears on the **Results** tab of the **AutoSupport** page.

## AutoSupport

The AutoSupport feature enables your StorageGRID system to send periodic and event-driven health and status messages to technical support to allow proactive monitoring and troubleshooting. StorageGRID AutoSupport also enables the use of Active IQ for predictive recommendations.

Settings

Results

### Weekly AutoSupport

Next Scheduled Time  2020-12-11 23:30:00 EST

Most Recent Result  Idle (NetApp Support)

Last Successful Time  N/A (NetApp Support)

### Event-Triggered AutoSupport

Most Recent Result  N/A (NetApp Support)

Last Successful Time  N/A (NetApp Support)

### User-Triggered AutoSupport

Most Recent Result  Failed (NetApp Support)

Last Successful Time  N/A (NetApp Support)

### AutoSupport On Demand

AutoSupport On Demand messages are only sent to NetApp Support.

Most Recent Result  N/A (NetApp Support)

Last Successful Time  N/A (NetApp Support)

### Weekly AutoSupport message failure

If a weekly AutoSupport message 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 message 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 message again at the next scheduled time.
5. Maintains the regular AutoSupport schedule if the message fails because the NMS service is unavailable, and if a message is sent before seven days pass.
6. When the NMS service is available again, sends an AutoSupport message immediately if a message has not been sent for seven days or more.

## User-triggered or event-triggered AutoSupport message failure

If a user-triggered or an event-triggered AutoSupport message 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 messages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.`
2. Does not attempt to send the message 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 messages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.`

Learn how to configure email server settings in the [monitor & troubleshoot instructions](#).

## Correcting an AutoSupport message 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 messages cannot be sent using SMTP protocol due to incorrect settings on the E-mail Server page.`

## Related information

[Monitor & troubleshoot](#)

## Managing Storage Nodes

Storage Nodes provide disk storage capacity and services. Managing Storage Nodes entails monitoring the amount of usable space on each node, using watermark settings, and applying Storage Node configuration settings.

- [What a Storage Node is](#)
- [Managing Storage Options](#)
- [Managing object metadata storage](#)
- [Configuring global settings for stored objects](#)
- [Storage Node configuration settings](#)
- [Managing full Storage Nodes](#)

## What a Storage Node is

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

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

#### **What the ADC service is**

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

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

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

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

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


#### **What the DDS service is**

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

#### **Object counts**

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

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



## Queries

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

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

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

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

## Consistency guarantees and controls

StorageGRID guarantees read-after-write consistency for newly created objects. Any GET operation following a successfully completed PUT operation will be able to read the newly written data. Overwrites of existing

objects, metadata updates, and deletes remain eventually consistent.

### **What the LDR service is**

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

The LDR service handles the following tasks:

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

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

### **Queries**

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

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

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

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

### **ILM activity**

Information lifecycle management (ILM) metrics allow you to monitor the rate at which objects are evaluated for ILM implementation. You can view these metrics on the Dashboard or on the Nodes page > ILM tab for each Storage Node.

### **Object stores**

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

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

Object Stores							
ID	Size	Available	Replicated Data	EC Data	Object Data (%)	Health	
0000	4.40 TB	1.35 TB	43.99 GB	0 bytes	1.00%	No Errors	
0001	1.97 TB	1.57 TB	44.76 GB	351.14 GB	20.09%	No Errors	
0002	1.97 TB	1.46 TB	43.29 GB	465.20 GB	25.81%	No Errors	
0003	1.97 TB	1.70 TB	43.51 GB	223.98 GB	13.58%	No Errors	
0004	1.97 TB	1.92 TB	44.03 GB	0 bytes	2.23%	No Errors	
0005	1.97 TB	1.46 TB	43.67 GB	463.36 GB	25.73%	No Errors	
0006	1.97 TB	1.92 TB	43.10 GB	1.61 GB	2.27%	No Errors	
0007	1.97 TB	1.35 TB	46.05 GB	575.24 GB	31.53%	No Errors	
0008	1.97 TB	1.81 TB	46.00 GB	112.84 GB	8.06%	No Errors	
0009	1.97 TB	1.57 TB	43.91 GB	352.72 GB	20.13%	No Errors	
000A	1.97 TB	1.70 TB	44.31 GB	226.81 GB	13.76%	No Errors	
000B	1.97 TB	1.92 TB	43.17 GB	780.07 MB	2.23%	No Errors	
000C	1.97 TB	1.58 TB	44.32 GB	339.56 GB	19.48%	No Errors	
000D	1.97 TB	1.82 TB	44.47 GB	107.34 GB	7.70%	No Errors	
000E	1.97 TB	1.68 TB	43.07 GB	241.70 GB	14.45%	No Errors	
000F	2.03 TB	1.50 TB	44.57 GB	475.47 GB	25.67%	No Errors	

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

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

### Metadata protection

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

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

### [Managing object metadata storage](#)

### Managing Storage Options

You can view and configure Storage Options using the Configuration menu in the Grid Manager. Storage Options include the object segmentation settings and the current values for storage watermarks. You can also view the S3 and Swift ports used by the deprecated CLB service on Gateway Nodes and by the LDR service on Storage Nodes.

For information on port assignments, see [Summary: IP addresses and ports for client connections](#).



**Object Segmentation**

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

**Storage Watermarks**

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


**Ports**

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

**What object segmentation is**

Object segmentation is the process of splitting up an object into a collection of smaller fixed-size objects in order 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.



If your StorageGRID system includes an Archive Node whose Target Type is Cloud Tiering — Simple Storage Service and the targeted archival storage system is Amazon Web Services (AWS), the Maximum Segment Size must be less than or equal to 4.5 GiB (4,831,838,208 bytes). This upper limit ensures that the AWS PUT limitation of five GBs is not exceeded. Requests to AWS that exceed this value fail.

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 are not necessarily stored on the same Storage Node. Container and segments can be stored on any Storage Node.

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

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

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

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

#### What Storage Volume watermarks are

StorageGRID uses Storage Volume watermarks to allow you to monitor the amount of usable space available on Storage Nodes. If the amount of space available on a node is less than a configured watermark setting, the Storage Status (SSTS) alarm is triggered so that you can determine if you need to add Storage Nodes.

To view the current settings for the Storage Volume watermarks, select **Configuration > Storage Options > Overview**.



### Storage Options Overview

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

#### Object Segmentation


Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

#### Storage Watermarks

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

The following figure represents a Storage Node that has three volumes and shows the relative position of the three Storage Volume watermarks. Within each Storage Node, StorageGRID reserves space on volume 0 for object metadata; any remaining space on that volume is used for object data. All other volumes are used

exclusively for object data, which includes replicated copies and erasure-coded fragments.



The Storage Volume watermarks are system-wide defaults that indicate the minimum amount of free space required on each volume in the Storage Node to prevent StorageGRID from changing the node's read-write behavior or triggering an alarm. Note that all volumes must reach the watermark before StorageGRID takes action. If some volumes have more than the minimum required amount of free space, the alarm is not triggered and the node's read-write behavior does not change.

### Storage Volume Soft Read-Only Watermark (VHWM)

The Storage Volume Soft Read-Only Watermark is the first watermark to indicate that a node's usable space for object data is becoming full. This watermark represents how much free space must exist on every volume in a Storage Node to prevent the node from going into "soft read-only mode." Soft 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.

If the amount of free space on each volume is less than the setting of this watermark, the Storage Status (SSTS) alarm is triggered at the Notice level, and the Storage Node transitions to soft read-only mode.

For example, suppose the Storage Volume Soft Read-Only Watermark is set to 10 GB, which is its default value. If less than 10 GB of free space remains on each volume in the Storage Node, the SSTS alarm is triggered at the Notice level, and the Storage Node transitions to soft read-only mode.

## Storage Volume Hard Read-Only Watermark (VROM)

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. This watermark represents how much free space must exist on every volume in a Storage Node to prevent the node from going in to "hard read-only mode." Hard read-only mode means that the Storage Node is read-only and no longer accepts write requests.

If the amount of free space on every volume in a Storage Node is less than the setting of this watermark, the Storage Status (SSTS) alarm is triggered at the Major level, and the Storage Node transitions to hard read-only mode.

For example, suppose the Storage Volume Hard Read-Only Watermark is set to 5 GB, which is its default value. If less than 5 GB of free space remains on each storage volume in the Storage Node, the SSTS alarm is triggered at the Major level, and the Storage Node transitions to hard read-only mode.

The value of the Storage Volume Hard Read-Only Watermark must be less than the value of the Storage Volume Soft Read-Only Watermark.

## Storage Volume Read-Write Watermark (VLWM)

The Storage Volume Read-Write Watermark only applies to Storage Nodes that have transitioned to read-only mode. This watermark determines when the Storage Node is allowed to become read-write again.

For example, suppose a Storage Node has transitioned to hard read-only mode. If the Storage Volume Read-Write Watermark is set to 30 GB (default), the free space on every storage volume in the Storage Node must increase from 5 GB to 30 GB before the node can become read-write again.

The value of the Storage Volume Read-Write Watermark must be greater than the value of the Storage Volume Soft Read-Only Watermark.

### Related information

[Managing full Storage Nodes](#)

## Managing object metadata storage

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

### What is object metadata?

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

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


- System metadata, including a unique ID for each object (UUID), the object name, the name of the S3 bucket or Swift container, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
- Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object.

- For replicated object copies, the current storage location of each copy.
- For erasure-coded object copies, the current storage location of each fragment.
- For object copies in a Cloud Storage Pool, the location of the object, including the name of the external bucket and the object's unique identifier.
- For segmented objects and multipart objects, segment identifiers and data sizes.

**How is object metadata stored?**


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

This figure represents the Storage Nodes at two sites. Each site has the same amount of object metadata, which is equally distributed across the 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 a number of 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 for StorageGRID 11.5 is based the following:

- 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 for StorageGRID 11.5
11.5	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)

Version used for initial StorageGRID installation	Amount of RAM on Storage Nodes	Default Metadata Reserved Space setting for StorageGRID 11.5
	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)

To view the Metadata Reserved Space setting for your StorageGRID system:

1. Select **Configuration > System Settings > Storage Options**.
2. In the Storage Watermarks table, locate **Metadata Reserved Space**.



## Storage Options Overview

Updated: 2021-02-23 11:58:33 MST

### Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

### Storage Watermarks

Description	Settings
Storage Volume Read-Write Watermark	30 GB
Storage Volume Soft Read-Only Watermark	10 GB
Storage Volume Hard Read-Only Watermark	5 GB
Metadata Reserved Space	8,000 GB

In the screenshot, the **Metadata Reserved Space** value is 8,000 GB (8 TB). This is the default setting for a new StorageGRID 11.5 installation in which each Storage Node has 128 GB or more of RAM.

#### Actual reserved space for metadata


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

Size of volume 0 for the node	Actual reserved space for metadata
Less than 500 GB (non production use)	10% of volume 0

Size of volume 0 for the node	Actual reserved space for metadata
500 GB or more	The smaller of these values: <ul style="list-style-type: none"> <li>• Volume 0</li> <li>• Metadata Reserved Space setting</li> </ul>

To view the actual reserved space for metadata on a particular Storage Node:

1. From the Grid Manager, select **Nodes > Storage Node**.
2. Select the **Storage** tab.
3. Hover 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 11.5 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.

The **Actual reserved** value corresponds to this Prometheus metric:

```
storagegrid_storage_utilization_metadata_reserved_bytes
```


#### Example for actual reserved metadata space

Suppose you install a new StorageGRID system using version 11.5. 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 11.5 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 summarizes how StorageGRID determines the allowed metadata space value for a Storage Node.


Actual reserved space for metadata	Allowed metadata space
4 TB or less	60% of actual reserved space for metadata, up to a maximum of 1.98 TB
More than 4 TB	$(\text{Actual reserved space for metadata} - 1 \text{ TB}) \times 60\%$ , up to a maximum of 2.64 TB



If your StorageGRID system stores (or is expected to store) more than 2.64 TB of metadata on any Storage Node, the allowed metadata space can be increased in some cases. If your Storage Nodes each have more than 128 GB of RAM and available free space on storage volume 0, contact your NetApp account representative. NetApp will review your requirements and increase the allowed metadata space for each Storage Node, if possible.

To view the allowed metadata space for a Storage Node:

1. From the Grid Manager, select **Node > Storage Node**.
2. Select the **Storage** tab.
3. Hover your cursor over the Storage Used — Object Metadata chart and locate the **Allowed** value.



In the screenshot, the **Allowed** value is 2.64 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 11.5. 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 11.5 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 2.64 TB. (This is the maximum value for actual reserved space.)

### 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 cannot 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:

[Monitor & troubleshoot](#)

- To increase the object metadata capacity for your system, you must add new Storage Nodes:

[Expand your grid](#)

#### Configuring global settings for stored objects

You can use Grid Options to configure the settings for all of the objects stored in your StorageGRID system, including stored object compression, stored object encryption, and stored object hashing.

- [Configuring stored object compression](#)
- [Configuring stored object encryption](#)
- [Configuring stored object hashing](#)

### Configuring stored object compression

You can use the Compress Stored Objects grid option to reduce the size of objects stored in StorageGRID, so that objects consume less storage.

#### What you'll need

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

#### About this task

The Compress Stored Objects grid option is disabled by default. If you enable this option, 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 this option, be aware of the following:

- You should not enable compression 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, enabling Compress Stored Objects will not further reduce an object's size.
- Do not enable compression if you are using NetApp FabricPool with StorageGRID.
- If the Compress Stored Objects grid option is enabled, S3 and Swift client applications should avoid performing GET Object operations that specify a range of bytes to be returned. These "range read" operations are inefficient because StorageGRID must effectively uncompress the objects to access the requested bytes. GET Object operations that request a small range of bytes from a very large object are especially inefficient; for example, it is inefficient to read a 10 MB range from a 50 GB compressed object.

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



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

#### Steps

1. Select **Configuration > System Settings > Grid Options**.
2. In the Stored Object Options section, select the **Compress Stored Objects** check box.

## Stored Object Options

Compress Stored Objects  

Stored Object Encryption   None  AES-128  AES-256

Stored Object Hashing   SHA-1  SHA-256

3. Click **Save**.

### Configuring stored object encryption

You can encrypt stored objects if you want to ensure that data cannot be retrieved in a readable form if an object store is compromised. By default, objects are not encrypted.

#### What you'll need

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

#### About this task

Stored object encryption enables the encryption of all object data as it is ingested through S3 or Swift. 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 are not encrypted.



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.

Stored objects can be encrypted using the AES-128 or AES-256 encryption algorithm.

The Stored Object Encryption setting applies only to S3 objects that have not been encrypted by bucket-level or object-level encryption.

#### Steps

1. Select **Configuration > System Settings > Grid Options**.
2. In the Stored Object Options section, change Stored Object Encryption to **None** (default), **AES-128**, or **AES-256**.

## Stored Object Options

Compress Stored Objects  

Stored Object Encryption   None  AES-128  AES-256

Stored Object Hashing   SHA-1  SHA-256

3. Click **Save**.

### Configuring stored object hashing

The Stored Object Hashing option specifies the hashing algorithm used to verify object integrity.

#### What you'll need

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

#### About this task

By default, object data is hashed using the SHA-1 algorithm. The SHA-256 algorithm requires additional CPU resources and is generally not recommended for integrity verification.



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.

#### Steps

1. Select **Configuration > System Settings > Grid Options**.
2. In the Stored Object Options section, change Stored Object Hashing to **SHA-1** (default) or **SHA-256**.

#### Stored Object Options



3. Click **Save**.

### Storage Node configuration settings

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



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

To access a Storage Node's configuration settings and counters:

1. Select **Support > Tools > Grid Topology**.
2. Select **site > Storage Node**.
3. Expand the Storage Node and select the service or component.

4. Select the **Configuration** tab.

The following tables summarize Storage Node configuration settings.

#### LDR

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

#### LDR > Data Store

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

#### LDR > Storage

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

#### LDR > Verification

Attribute Name	Code	Description
Reset Missing Objects Count	VCMI	Resets the count of Missing Objects Detected (OMIS). Use only after foreground verification completes. Missing replicated object data is restored automatically by the StorageGRID system.
Verify	FVOV	Select object stores on which to perform foreground verification.
Verification Rate	VPRI	Set the rate at which background verification takes place. See information on configuring the background verification rate.

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

#### LDR > Erasure Coding

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

#### LDR > Replication

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

### Related information

[Monitor & troubleshoot](#)

### Managing 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.

#### Adding 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.

#### Adding storage expansion shelves

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

## Adding 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.

### Related information

[Expand your grid](#)

## Managing Admin Nodes

Each site in a StorageGRID deployment can have one or more Admin Nodes.

- [What an Admin Node is](#)
- [Using multiple Admin Nodes](#)
- [Identifying the primary Admin Node](#)
- [Selecting a preferred sender](#)
- [Viewing notification status and queues](#)
- [How Admin Nodes show acknowledged alarms \(legacy system\)](#)
- [Configuring audit client access](#)

### What an Admin Node is

Admin Nodes provide management services such as system configuration, monitoring, and logging. Each grid must have one primary Admin Node and might have any number of non-primary Admin Nodes for redundancy.

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, maintenance procedures must be performed using the primary Admin Node.

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

Admin Nodes host the following services:

- AMS service
- CMN service
- NMS service
- Prometheus service
- Load Balancer and High Availability services (to support S3 and Swift client traffic)

Admin Nodes also support the Management Application Program Interface (mgmt-api) to process requests from the Grid Management API and the Tenant Management API.

### What the AMS service is

The Audit Management System (AMS) service tracks system activity and events.

**What the CMN service is**

The Configuration Management Node (CMN) service manages system-wide configurations of connectivity and protocol features needed by all services. In addition, the CMN service is used to run and monitor grid tasks. There is only one CMN service per StorageGRID deployment. The Admin Node that hosts the CMN service is known as the primary Admin Node.

**What the NMS service is**

The Network Management System (NMS) service powers the monitoring, reporting, and configuration options that are displayed through the Grid Manager, the StorageGRID system's browser-based interface.

**What the Prometheus service is**

The Prometheus service collects time series metrics from the services on all nodes.

**Related information**

[Using the Grid Management API](#)

[Use a tenant account](#)


[Managing load balancing](#)

[Managing high availability groups](#)

**Using 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 and alarms (legacy system) are still triggered, and email notifications and AutoSupport messages are still sent. However, having multiple Admin Nodes does not provide failover protection except for notifications and AutoSupport messages. In particular, alarm acknowledgments made from one Admin Node are not copied to other Admin Nodes.



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.



When using an HA group, access is interrupted if the Master 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.

#### Related information

[Managing high availability groups](#)

#### Identifying the primary Admin Node

The primary Admin Node hosts the CMN service. Some maintenance procedures can only be performed using the primary Admin Node.

## What you'll need

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

## Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **site > Admin Node**, and then click **+** to expand the topology tree and show the services hosted on this Admin Node.

The primary Admin Node hosts the CMN service.

3. If this Admin Node does not host the CMN service, check the other Admin Nodes.

## Selecting a preferred sender

If your StorageGRID deployment includes multiple Admin Nodes, you can select which Admin Node should be the preferred sender of notifications. By default, the primary Admin Node is selected, but any Admin Node can be the preferred sender.

## What you'll need

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

## About this task

The **Configuration > System Settings > Display Options** page shows which Admin Node is currently selected to be the preferred sender. The primary Admin Node is selected by default.

Under normal system operations, only the preferred sender sends the following notifications:

- AutoSupport messages
- SNMP notifications
- Alert emails
- Alarm emails (legacy system)

However, all other Admin Nodes (standby senders) monitor the preferred sender. If a problem is detected, a standby sender can also send these notifications.

Both the preferred sender and a standby sender might send notifications 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.
- After a 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 alarm notifications and AutoSupport messages, all Admin Nodes send the test email. When you test alert notifications, you must sign in to every Admin Node to verify connectivity.

## Steps

1. Select **Configuration > System Settings > Display Options**.
2. From the Display Options menu, select **Options**.
3. Select the Admin Node you want to set as the preferred sender from the drop-down list.



### Display Options

Updated: 2017-08-30 16:31:10 MDT

Current Sender	ADMIN-DC1-ADM1
Preferred Sender	ADMIN-DC1-ADM1
GUI Inactivity Timeout	900
Notification Suppress All	<input type="checkbox"/>

Apply Changes

4. Click **Apply Changes**.

The Admin Node is set as the preferred sender of notifications.

## Viewing notification status and queues

The 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**. Finally, select **site > Admin Node > NMS > Interface Engine**.

The screenshot shows the 'Overview: NMS (170-176) - Interface Engine' page. It has tabs for Overview, Alarms, Reports, and Configuration. The Overview tab is active. The page displays the following information:

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. Because notifications can be dropped from the notifications queue without being sent, it is possible that an alarm can be triggered without a notification being sent. In the event that a notification is dropped from the queue without being sent, the MINS (E-mail Notification Status) Minor alarm is triggered.

### How Admin Nodes show acknowledged alarms (legacy system)

When you acknowledge an alarm on one Admin Node, the acknowledged alarm is not copied to any other Admin Node. Because acknowledgments are not copied to other Admin Nodes, the Grid Topology tree might not look the same for each Admin Node.

This difference can be useful when connecting web clients. Web clients can have different views of the StorageGRID system based on the administrator needs.



Note that notifications are sent from the Admin Node where the acknowledgment occurs.

### Configuring audit client access

The Admin Node, through the Audit Management System (AMS) service, logs all audited system events to a log file available through the audit share, which is added to each Admin Node at installation. For easy access to audit logs, you can configure client access to audit shares for both CIFS and NFS.

The StorageGRID system uses positive acknowledgment to prevent loss of audit messages before they are written to the log file. A message remains queued at a service until the AMS service or an intermediate audit relay service has acknowledged control of it.

For more information, see the instructions for understanding audit messages.



If you have the option to use CIFS or NFS, choose NFS.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

### Related information

[What an Admin Node is](#)

[Review audit logs](#)

[Upgrade software](#)

### Configuring audit clients for CIFS

The procedure used to configure an audit client depends on the authentication method: Windows Workgroup or Windows Active Directory (AD). When added, the audit share is automatically enabled as a read-only share.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

### Related information

[Upgrade software](#)

### Configuring audit clients for Workgroup

Perform this procedure for each Admin Node in a StorageGRID deployment from which you want to retrieve audit messages.

### What you'll need

- You must have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You must have the `Configuration.txt` file (available in the SAID package).

### About this task

Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

### 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. Confirm that all services have a state of Running or Verified: `storagegrid-status`

If all services are not Running or Verified, resolve issues before continuing.

3. Return to the command line, press **Ctrl+C**.
4. Start the CIFS configuration utility: `config_cifs.rb`

```

-----
| Shares                | Authentication          | Config                  |
-----
| add-audit-share       | set-authentication      | validate-config        |
| enable-disable-share  | set-netbios-name        | help                   |
| add-user-to-share     | join-domain             | exit                   |
| remove-user-from-share | add-password-server     |                         |
| modify-group          | remove-password-server  |                         |
|                       | add-wins-server         |                         |
|                       | remove-wins-server      |                         |
-----

```

5. Set the authentication for the Windows Workgroup:

If authentication has already been set, an advisory message appears. If authentication has already been set, go to the next step.

- a. Enter: `set-authentication`
- b. When prompted for Windows Workgroup or Active Directory installation, enter: `workgroup`
- c. When prompted, enter a name of the Workgroup: `workgroup_name`
- d. When prompted, create a meaningful NetBIOS name: `netbios_name`

or

Press **Enter** to use the Admin Node's hostname as the NetBIOS name.

The script restarts the Samba server and changes are applied. This should take less than one minute. After setting authentication, add an audit client.

- e. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

6. Add an audit client:

- a. Enter: `add-audit-share`



The share is automatically added as read-only.

- b. When prompted, add a user or group: `user`
- c. When prompted, enter the audit user name: `audit_user_name`
- d. When prompted, enter a password for the audit user: `password`
- e. When prompted, re-enter the same password to confirm it: `password`

f. When prompted, press **Enter**.

The CIFS configuration utility is displayed.



There is no need to enter a directory. The audit directory name is predefined.

7. If more than one user or group is permitted to access the audit share, add the additional users:

a. Enter: `add-user-to-share`

A numbered list of enabled shares is displayed.

b. When prompted, enter the number of the audit-export share: `share_number`

c. When prompted, add a user or group: `user`

or `group`

d. When prompted, enter the name of the audit user or group: `audit_user` or `audit_group`

e. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

f. Repeat these substeps for each additional user or group that has access to the audit share.

8. Optionally, verify your configuration: `validate-config`

The services are checked and displayed. You can safely ignore the following messages:

```
Can't find include file /etc/samba/includes/cifs-interfaces.inc
Can't find include file /etc/samba/includes/cifs-filesystem.inc
Can't find include file /etc/samba/includes/cifs-custom-config.inc
Can't find include file /etc/samba/includes/cifs-shares.inc
rlimit_max: increasing rlimit_max (1024) to minimum Windows limit
(16384)
```

a. When prompted, press **Enter**.

The audit client configuration is displayed.

b. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

9. Close the CIFS configuration utility: `exit`

10. Start the Samba service: `service smb start`

11. If the StorageGRID deployment is a single site, go to the next step.

or

Optionally, if the StorageGRID deployment includes Admin Nodes at other sites, enable these audit share

as required:

- a. Remotely log in to a site's 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. Repeat the steps to configure the audit share for each additional Admin Node.
- c. Close the remote secure shell login to the remote Admin Node: `exit`

12. Log out of the command shell: `exit`

## Related information

[Upgrade software](#)

## Configuring audit clients for Active Directory

Perform this procedure for each Admin Node in a StorageGRID deployment from which you want to retrieve audit messages.

### What you'll need

- You must have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You must have the CIFS Active Directory username and password.
- You must have the `Configuration.txt` file (available in the SAID package).



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

### 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. Confirm that all services have a state of Running or Verified: `storagegrid-status`

If all services are not Running or Verified, resolve issues before continuing.

3. Return to the command line, press **Ctrl+C**.
4. Start the CIFS configuration utility: `config_cifs.rb`

Shares	Authentication	Config
<code>add-audit-share</code>	<code>set-authentication</code>	<code>validate-config</code>
<code>enable-disable-share</code>	<code>set-netbios-name</code>	<code>help</code>
<code>add-user-to-share</code>	<code>join-domain</code>	<code>exit</code>
<code>remove-user-from-share</code>	<code>add-password-server</code>	
<code>modify-group</code>	<code>remove-password-server</code>	
	<code>add-wins-server</code>	
	<code>remove-wins-server</code>	

5. Set the authentication for Active Directory: `set-authentication`

In most deployments, you must set the authentication before adding the audit client. If authentication has already been set, an advisory message appears. If authentication has already been set, go to the next step.

- When prompted for Workgroup or Active Directory installation: `ad`
- When prompted, enter the name of the AD domain (short domain name).
- When prompted, enter the domain controller's IP address or DNS hostname.
- When prompted, enter the full domain realm name.

Use uppercase letters.

- When prompted to enable winbind support, type `y`.

Winbind is used to resolve user and group information from AD servers.

- When prompted, enter the NetBIOS name.
- When prompted, press **Enter**.

The CIFS configuration utility is displayed.

6. Join the domain:

- If not already started, start the CIFS configuration utility: `config_cifs.rb`
- Join the domain: `join-domain`
- You are prompted to test if the Admin Node is currently a valid member of the domain. If this Admin Node has not previously joined the domain, enter: `no`
- When prompted, provide the Administrator's username: `administrator_username`

where `administrator_username` is the CIFS Active Directory username, not the StorageGRID username.

- When prompted, provide the Administrator's password: `administrator_password`

where `administrator_password` is the CIFS Active Directory password, not the StorageGRID password.

password.

- f. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

7. Verify that you have correctly joined the domain:

- a. Join the domain: `join-domain`

- b. When prompted to test if the server is currently a valid member of the domain, enter: `y`

If you receive the message “Join is OK,” you have successfully joined the domain. If you do not get this response, try setting authentication and joining the domain again.

- c. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

8. Add an audit client: `add-audit-share`

- a. When prompted to add a user or group, enter: `user`

- b. When prompted to enter the audit user name, enter the audit user name.

- c. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

9. If more than one user or group is permitted to access the audit share, add additional users: `add-user-to-share`

A numbered list of enabled shares is displayed.

- a. Enter the number of the audit-export share.

- b. When prompted to add a user or group, enter: `group`

You are prompted for the audit group name.

- c. When prompted for the audit group name, enter the name of the audit user group.

- d. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

- e. Repeat this step for each additional user or group that has access to the audit share.

10. Optionally, verify your configuration: `validate-config`

The services are checked and displayed. You can safely ignore the following messages:

- Can't find include file `/etc/samba/includes/cifs-interfaces.inc`
- Can't find include file `/etc/samba/includes/cifs-filesystem.inc`
- Can't find include file `/etc/samba/includes/cifs-interfaces.inc`
- Can't find include file `/etc/samba/includes/cifs-custom-config.inc`

- Can't find include file `/etc/samba/includes/cifs-shares.inc`
- `rlimit_max`: increasing `rlimit_max` (1024) to minimum Windows limit (16384)



Do not combine the setting 'security=ads' with the 'password server' parameter. (by default Samba will discover the correct DC to contact automatically).

- a. When prompted, press **Enter** to display the audit client configuration.
- b. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

11. Close the CIFS configuration utility: `exit`
12. If the StorageGRID deployment is a single site, go to the next step.

or

Optionally, if the StorageGRID deployment includes Admin Nodes at other sites, enable these audit shares as required:

- a. Remotely log in to a site's 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. Repeat these steps to configure the audit shares for each Admin Node.
- c. Close the remote secure shell login to the Admin Node: `exit`

13. Log out of the command shell: `exit`

## Related information

[Upgrade software](#)

## Adding a user or group to a CIFS audit share

You can add a user or group to a CIFS audit share that is integrated with AD authentication.

### What you'll need

- You must have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You must have the `Configuration.txt` file (available in the SAID package).

### About this task

The following procedure is for an audit share integrated with AD authentication.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

## 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. Confirm that all services have a state of Running or Verified. Enter: `storagegrid-status`

If all services are not Running or Verified, resolve issues before continuing.

3. Return to the command line, press **Ctrl+C**.
4. Start the CIFS configuration utility: `config_cifs.rb`

```
-----  
| Shares                | Authentication          | Config                  |  
-----  
| add-audit-share       | set-authentication      | validate-config        |  
| enable-disable-share  | set-netbios-name       | help                   |  
| add-user-to-share     | join-domain            | exit                   |  
| remove-user-from-share| add-password-server    |                         |  
| modify-group          | remove-password-server |                         |  
|                       | add-wins-server        |                         |  
|                       | remove-wins-server     |                         |  
-----
```

5. Start adding a user or group: `add-user-to-share`  
A numbered list of audit shares that have been configured is displayed.
6. When prompted, enter the number for the audit share (audit-export): `audit_share_number`  
You are asked if you would like to give a user or a group access to this audit share.
7. When prompted, add a user or group: `user` or `group`
8. When prompted for the user or group name for this AD audit share, enter the name.

The user or group is added as read-only for the audit share both in the server's operating system and in the CIFS service. The Samba configuration is reloaded to enable the user or group to access the audit client share.

9. When prompted, press **Enter**.

The CIFS configuration utility is displayed.

10. Repeat these steps for each user or group that has access to the audit share.

11. Optionally, verify your configuration: `validate-config`

The services are checked and displayed. You can safely ignore the following messages:

- Can't find include file `/etc/samba/includes/cifs-interfaces.inc`
- Can't find include file `/etc/samba/includes/cifs-filesystem.inc`
- Can't find include file `/etc/samba/includes/cifs-custom-config.inc`
- Can't find include file `/etc/samba/includes/cifs-shares.inc`
  - a. When prompted, press **Enter** to display the audit client configuration.
  - b. When prompted, press **Enter**.

12. Close the CIFS configuration utility: `exit`

13. Determine if you need to enable additional audit shares, as follows:

- If the StorageGRID deployment is a single site, go to the next step.
- If the StorageGRID deployment includes Admin Nodes at other sites, enable these audit shares as required:
  - a. Remotely log in to a site's 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. Repeat these steps to configure the audit shares for each Admin Node.
  - c. Close the remote secure shell login to the remote Admin Node: `exit`

14. Log out of the command shell: `exit`

## Removing a user or group from a CIFS audit share

You cannot remove the last user or group permitted to access the audit share.

### What you'll need

- You must have the `Passwords.txt` file with the root account passwords (available in the SAID package).
- You must have the `Configuration.txt` file (available in the SAID package).

### About this task

Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

### 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 CIFS configuration utility: `config_cifs.rb`

```
-----  
| Shares                | Authentication          | Config                  |  
-----  
| add-audit-share       | set-authentication      | validate-config        |  
| enable-disable-share  | set-netbios-name       | help                   |  
| add-user-to-share     | join-domain            | exit                   |  
| remove-user-from-share| add-password-server    |                        |  
| modify-group          | remove-password-server |                        |  
|                       | add-wins-server        |                        |  
|                       | remove-wins-server     |                        |  
-----
```

3. Start removing a user or group: `remove-user-from-share`

A numbered list of available audit shares for the Admin Node is displayed. The audit share is labeled `audit-export`.

4. Enter the number of the audit share: `audit_share_number`

5. When prompted to remove a user or a group: `user` or `group`

A numbered list of users or groups for the audit share is displayed.

6. Enter the number corresponding to the user or group you want to remove: `number`

The audit share is updated, and the user or group is no longer permitted access to the audit share. For example:

```
Enabled shares  
 1. audit-export  
Select the share to change: 1  
Remove user or group? [User/group]: User  
Valid users for this share  
 1. audituser  
 2. newaudituser  
Select the user to remove: 1  
  
Removed user "audituser" from share "audit-export".  
  
Press return to continue.
```

7. Close the CIFS configuration utility: `exit`
8. If the StorageGRID deployment includes Admin Nodes at other sites, disable the audit share at each site as required.
9. Log out of each command shell when configuration is complete: `exit`

## Related information

[Upgrade software](#)

## Changing a CIFS audit share user or group name

You can change the name of a user or a group for a CIFS audit share by adding a new user or group and then deleting the old one.

### About this task

Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

### Steps

1. Add a new user or group with the updated name to the audit share.
2. Delete the old user or group name.

## Related information

[Upgrade software](#)

[Adding a user or group to a CIFS audit share](#)

[Removing a user or group from a CIFS audit share](#)

## Verifying CIFS audit integration

The audit share is read-only. Log files are intended to be read by computer applications and verification does not include opening a file. It is considered sufficient verification that the audit log files appear in a Windows Explorer window. Following connection verification, close all windows.

### Configuring the audit client for NFS

The audit share is automatically enabled as a read-only share.

### What you'll need

- You must have the `Passwords.txt` file with the root/admin password (available in the SAID package).
- You must have the `Configuration.txt` file (available in the SAID package).
- The audit client must be using NFS Version 3 (NFSv3).

### About this task

Perform this procedure for each Admin Node in a StorageGRID deployment from which you want to retrieve audit messages.

### 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. Confirm that all services have a state of Running or Verified. Enter: `storagegrid-status`

If any services are not listed as Running or Verified, resolve issues before continuing.

3. Return to the command line. Press **Ctrl+C**.
4. Start the NFS configuration utility. Enter: `config_nfs.rb`

```

-----
| Shares                | Clients                | Config                |
-----
| add-audit-share      | add-ip-to-share       | validate-config      |
| enable-disable-share | remove-ip-from-share  | refresh-config       |
|                      |                       | help                 |
|                      |                       | exit                 |
-----

```

5. Add the audit client: `add-audit-share`
  - a. When prompted, enter the audit client's IP address or IP address range for the audit share: `client_IP_address`
  - b. When prompted, press **Enter**.
6. If more than one audit client is permitted to access the audit share, add the IP address of the additional user: `add-ip-to-share`
  - a. Enter the number of the audit share: `audit_share_number`
  - b. When prompted, enter the audit client's IP address or IP address range for the audit share: `client_IP_address`
  - c. When prompted, press **Enter**.
 

The NFS configuration utility is displayed.
  - d. Repeat these substeps for each additional audit client that has access to the audit share.
7. Optionally, verify your configuration.
  - a. Enter the following: `validate-config`

The services are checked and displayed.
  - b. When prompted, press **Enter**.

The NFS configuration utility is displayed.

c. Close the NFS configuration utility: `exit`

8. Determine if you must enable audit shares at other sites.

- If the StorageGRID deployment is a single site, go to the next step.
- If the StorageGRID deployment includes Admin Nodes at other sites, enable these audit shares as required:
  - a. Remotely log in to the site's 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. Repeat these steps to configure the audit shares for each additional Admin Node.
  - c. Close the remote secure shell login to the remote Admin Node. Enter: `exit`

9. Log out of the command shell: `exit`

NFS audit clients are granted access to an audit share based on their IP address. Grant access to the audit share to a new NFS audit client by adding its IP address to the share, or remove an existing audit client by removing its IP address.

### Adding an NFS audit client to an audit share

NFS audit clients are granted access to an audit share based on their IP address. Grant access to the audit share to a new NFS audit client by adding its IP address to the audit share.

#### What you'll need

- You must have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You must have the `Configuration.txt` file (available in the SAID package).
- The audit client must be using NFS Version 3 (NFSv3).

#### 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 NFS configuration utility: `config_nfs.rb`

Shares	Clients	Config
add-audit-share	add-ip-to-share	validate-config
enable-disable-share	remove-ip-from-share	refresh-config
		help
		exit

3. Enter: `add-ip-to-share`

A list of NFS audit shares enabled on the Admin Node is displayed. The audit share is listed as:  
`/var/local/audit/export`

4. Enter the number of the audit share: `audit_share_number`

5. When prompted, enter the audit client's IP address or IP address range for the audit share:  
`client_IP_address`

The audit client is added to the audit share.

6. When prompted, press **Enter**.

The NFS configuration utility is displayed.

7. Repeat the steps for each audit client that should be added to the audit share.

8. Optionally, verify your configuration: `validate-config`

The services are checked and displayed.

a. When prompted, press **Enter**.

The NFS configuration utility is displayed.

9. Close the NFS configuration utility: `exit`

10. If the StorageGRID deployment is a single site, go to the next step.

Otherwise, if the StorageGRID deployment includes Admin Nodes at other sites, optionally enable these audit shares as required:

a. Remotely log in to a site's 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. Repeat these steps to configure the audit shares for each Admin Node.

c. Close the remote secure shell login to the remote Admin Node: `exit`

11. Log out of the command shell: `exit`

## Verifying NFS audit integration

After you configure an audit share and add an NFS audit client, you can mount the audit client share and verify that the files are available from the audit share.

### Steps

1. Verify connectivity (or variant for the client system) using the client-side IP address of the Admin Node hosting the AMS service. Enter: `ping IP_address`

Verify that the server responds, indicating connectivity.

2. Mount the audit read-only share using a command appropriate to the client operating system. A sample Linux command is (enter on one line):

```
mount -t nfs -o hard,intr Admin_Node_IP_address:/var/local/audit/export  
myAudit
```

Use the IP address of the Admin Node hosting the AMS service and the predefined share name for the audit system. The mount point can be any name selected by the client (for example, `myAudit` in the previous command).

3. Verify that the files are available from the audit share. Enter: `ls myAudit /*`

where `myAudit` is the mount point of the audit share. There should be at least one log file listed.

## Removing an NFS audit client from the audit share

NFS audit clients are granted access to an audit share based on their IP address. You can remove an existing audit client by removing its IP address.

### What you'll need

- You must have the `Passwords.txt` file with the root/admin account password (available in the SAID package).
- You must have the `Configuration.txt` file (available in the SAID package).

### About this task

You cannot remove the last IP address permitted to access the audit share.

### 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 NFS configuration utility: `config_nfs.rb`

```
-----  
| Shares                | Clients                | Config                |  
-----  
| add-audit-share      | add-ip-to-share       | validate-config      |  
| enable-disable-share | remove-ip-from-share  | refresh-config       |  
|                       |                       | help                 |  
|                       |                       | exit                 |  
-----
```

3. Remove the IP address from the audit share: `remove-ip-from-share`

A numbered list of audit shares configured on the server is displayed. The audit share is listed as:  
`/var/local/audit/export`

4. Enter the number corresponding to the audit share: `audit_share_number`

A numbered list of IP addresses permitted to access the audit share is displayed.

5. Enter the number corresponding to the IP address you want to remove.

The audit share is updated, and access is no longer permitted from any audit client with this IP address.

6. When prompted, press **Enter**.

The NFS configuration utility is displayed.

7. Close the NFS configuration utility: `exit`

8. If your StorageGRID deployment is a multiple data center site deployment with additional Admin Nodes at the other sites, disable these audit shares as required:

a. Remotely log in to each site's 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. Repeat these steps to configure the audit shares for each additional Admin Node.

c. Close the remote secure shell login to the remote Admin Node: `exit`

9. Log out of the command shell: `exit`

## Changing the IP address of an NFS audit client

1. Add a new IP address to an existing NFS audit share.

2. Remove the original IP address.

## Related information

[Adding an NFS audit client to an audit share](#)

[Removing an NFS audit client from the audit share](#)

## Managing Archive Nodes

Optionally, each of your StorageGRID system's data center sites can be deployed with an Archive Node, which allows you to connect to a targeted external archival storage system, such as Tivoli Storage Manager (TSM).

After configuring connections to the external target, you can configure the Archive Node to optimize TSM performance, take an Archive Node offline when a TSM server is nearing capacity or unavailable, and configure replication and retrieve settings. You can also set Custom alarms for the Archive Node.

- [What an Archive Node is](#)
- [Configuring Archive Node connections to archival storage](#)
- [Setting Custom alarms for the Archive Node](#)
- [Integrating Tivoli Storage Manager](#)

### What an Archive Node is

The Archive Node provides an interface through which you can target an external archival storage system for the long term storage of object data. The Archive Node also monitors this connection and the transfer of object data between the StorageGRID system and the targeted external archival storage system.

The screenshot displays the StorageGRID management console interface. On the left, a tree view shows the 'Grid Topology' with 'Data Center 1' expanded to show nodes DC1-ADM1-98-160 through DC1-S3-98-164, with DC1-ARC1-98-165 selected. The main panel shows the 'Overview' tab for 'ARC (DC1-ARC1-98-165) - ARC', updated on 2015-09-30 10:29:18 PDT. Below this, a status table shows various components are 'Online' with 'No Errors'. A 'Node Information' section at the bottom lists device type, version, build, node ID, and site ID.

Component	Status	Errors
ARC State	Online	No Errors
ARC Status	Online	No Errors
Tivoli Storage Manager State	Online	No Errors
Tivoli Storage Manager Status	Online	No Errors
Store State	Online	No Errors
Store Status	Online	No Errors
Retrieve State	Online	No Errors
Retrieve Status	Online	No Errors
Inbound Replication Status	Online	No Errors
Outbound Replication Status	Online	No Errors

**Node Information**

Device Type:	Archive Node
Version:	10.2.0
Build:	20150928.2133.a27b3ab
Node ID:	19002524
Site ID:	10

Object data that cannot be deleted, but is not regularly accessed, can at any time be moved off of a Storage Node's spinning disks and onto external archival storage such as the cloud or tape. This archiving of object data is accomplished through the configuration of a data center site's Archive Node and then the configuration of ILM rules where this Archive Node is selected as the "target" for content placement instructions. The Archive Node does not manage archived object data itself; this is achieved by the external archive device.



Object metadata is not archived, but remains on Storage Nodes.

### What the ARC service is

The Archive Node's Archive (ARC) service provides the management interface you can use to configure connections to external archival storage, such as tape through TSM middleware.

It is the ARC service that interacts with an external archival storage system, sending object data for near-line storage and performing retrievals when a client application requests an archived object. When a client application requests an archived object, a Storage Node requests the object data from the ARC service. The ARC service makes a request to the external archival storage system, which retrieves the requested object data and sends it to the ARC service. The ARC service verifies the object data and forwards it to the Storage Node, which in turn returns the object to the requesting client application.

Requests for object data archived to tape through TSM middleware are managed for efficiency of retrievals. Requests can be ordered so that objects stored in sequential order on tape are requested in that same sequential order. Requests are then queued for submission to the storage device. Depending upon the archival device, multiple requests for objects on different volumes can be processed simultaneously.

### Configuring Archive Node connections to archival storage

When you configure an Archive Node to connect with an external archive, you must select the target type.

The StorageGRID system supports the archiving of object data to the cloud through an S3 interface or to tape through Tivoli Storage Manager (TSM) middleware.



Once the type of archival target is configured for an Archive Node, the target type cannot be changed.

- [Archiving to the cloud through the S3 API](#)
- [Archiving to tape through TSM middleware](#)
- [Configuring Archive Node retrieve settings](#)
- [Configuring Archive Node replication](#)

### Archiving to the cloud through the S3 API

You can configure an Archive Node to connect directly to Amazon Web Services (AWS) or to any other system that can interface to the StorageGRID system through the S3 API.



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. The **Cloud Tiering - Simple Storage Service (S3)** option is still supported, but you might prefer to implement Cloud Storage Pools instead.

If you are currently using an Archive Node with the **Cloud Tiering - Simple Storage Service (S3)** option, consider migrating your objects to a Cloud Storage Pool. See the instructions for managing objects with information lifecycle management.

## Related information

[Manage objects with ILM](#)

### Configuring connection settings for the S3 API

If you are connecting to an Archive Node using the S3 interface, you must configure the connection settings for the S3 API. Until these settings are configured, the ARC service remains in a Major alarm state as it is unable to communicate with the external archival storage system.



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. The **Cloud Tiering - Simple Storage Service (S3)** option is still supported, but you might prefer to implement Cloud Storage Pools instead.

If you are currently using an Archive Node with the **Cloud Tiering - Simple Storage Service (S3)** option, consider migrating your objects to a Cloud Storage Pool. See the instructions for managing objects with information lifecycle management.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have created a bucket on the target archival storage system:
  - The bucket must be dedicated to a single Archive Node. It cannot be used by other Archive Nodes or other applications.
  - The bucket must have the appropriate region selected for your location.
  - The bucket should be configured with versioning suspended.
- Object Segmentation must be enabled and the Maximum Segment Size must be less than or equal to 4.5 GiB (4,831,838,208 bytes). S3 API requests that exceed this value will fail if S3 is used as the external archival storage system.

### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.

Target Type: Cloud Tiering - Simple Storage Service (S3)

### Cloud Tiering (S3) Account

Bucket Name:	<input type="text" value="name"/>
Region:	Virginia or Pacific Northwest (us-east-1)
Endpoint:	<input type="text" value="https://10.10.10.123:8082"/> <input type="checkbox"/> Use AWS
Endpoint Authentication:	<input type="checkbox"/>
Access Key:	<input type="text" value="ABCD123EFG45AB"/>
Secret Access Key:	<input type="password" value="....."/>
Storage Class:	Standard (Default)

Apply Changes 

- Select **Cloud Tiering - Simple Storage Service (S3)** from the Target Type drop-down list.



Configuration settings are unavailable until you select a Target Type.

- Configure the cloud tiering (S3) account through which the Archive Node will connect to the target external S3 capable archival storage system.

Most of the fields on this page are self-explanatory. The following describes fields for which you might need guidance.

- **Region:** Only available if **Use AWS** is selected. The region you select must match the bucket's region.
- **Endpoint** and **Use AWS:** For Amazon Web Services (AWS), select **Use AWS**. **Endpoint** is then automatically populated with an endpoint URL based on the Bucket Name and Region attributes. For example:

```
https://bucket.region.amazonaws.com
```

For a non-AWS target, enter the URL of the system hosting the bucket, including the port number. For example:

```
https://system.com:1080
```

- **End Point Authentication:** Enabled by default. If the network to the external archival storage system is trusted, you can unselect the check box to disable endpoint SSL certificate and hostname verification for the targeted external archival storage system. If another instance of a StorageGRID system is the target archival storage device and the system is configured with publicly signed certificates, you can keep the check box selected.

- **Storage Class:** Select **Standard (Default)** for regular storage. Select **Reduced Redundancy** only for objects that can be easily recreated. **Reduced Redundancy** provides lower cost storage with less reliability. If the targeted archival storage system is another instance of the StorageGRID system, **Storage Class** controls how many interim copies of the object are made at ingest on the target system, if dual commit is used when objects are ingested there.

6. Click **Apply Changes**.

The specified configuration settings are validated and applied to your StorageGRID system. Once configured, the target cannot be changed.

### Related information

[Manage objects with ILM](#)

### Modifying connection settings for S3 API

After the Archive Node is configured to connect to an external archival storage system through the S3 API, you can modify some settings should the connection change.

#### What you'll need

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

#### About this task

If you change the Cloud Tiering (S3) account, you must ensure that the user access credentials have read/write access to the bucket, including all objects that were previously ingested by the Archive Node to the bucket.

#### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.

Target Type: Cloud Tiering - Simple Storage Service (S3)

### Cloud Tiering (S3) Account

Bucket Name:	<input type="text" value="name"/>
Region:	<span>Virginia or Pacific Northwest (us-east-1)</span>
Endpoint:	<input type="text" value="https://10.10.10.123:8082"/> <input type="checkbox"/> Use AWS
Endpoint Authentication:	<input type="checkbox"/>
Access Key:	<input type="text" value="ABCD123EFG45AB"/>
Secret Access Key:	<input type="password" value="•••••"/>
Storage Class:	<span>Standard (Default)</span>

**Apply Changes** 

4. Modify account information, as necessary.

If you change the storage class, new object data is stored with the new storage class. Existing object continue to be stored under the storage class set when ingested.



Bucket Name, Region, and Endpoint, use AWS values and cannot be changed.

5. Click **Apply Changes**.

### Modifying the Cloud Tiering Service state

You can control the Archive Node's ability read and write to the targeted external archival storage system that connects through the S3 API by changing the state of the Cloud Tiering Service.

#### What you'll need


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

#### About this task

You can effectively take the Archive Node offline by changing the Cloud Tiering Service State to **Read-Write Disabled**.

#### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC**.
3. Select **Configuration > Main**.



4. Select a **Cloud Tiering Service State**.
5. Click **Apply Changes**.

### Resetting the Store Failure Count for S3 API connection


If your Archive Node connects to an archival storage system through the S3 API, you can reset the Store Failure Count, which can be used to clear the ARVF (Store Failures) alarm.

#### What you'll need

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

#### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Store**.
3. Select **Configuration > Main**.



4. Select **Reset Store Failure Count**.
5. Click **Apply Changes**.

The Store Failures attribute resets to zero.

## Migrating objects from Cloud Tiering - S3 to a Cloud Storage Pool

If you are currently using the **Cloud Tiering - Simple Storage Service (S3)** feature to tier object data to an S3 bucket, consider migrating your objects to a Cloud Storage Pool instead. Cloud Storage Pools provide a scalable approach that takes advantage of all of the Storage Nodes in your StorageGRID system.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You have already stored objects in the S3 bucket configured for Cloud Tiering.



Before migrating object data, contact your NetApp account representative to understand and manage any associated costs.

### About this task

From an ILM perspective, a Cloud Storage Pool is similar to a storage pool. However, while storage pools consist of Storage Nodes or Archive Nodes within the StorageGRID system, a Cloud Storage Pool consists of an external S3 bucket.

Before migrating objects from Cloud Tiering - S3 to a Cloud Storage Pool, you must first create an S3 bucket and then create the Cloud Storage Pool in StorageGRID. Then, you can create a new ILM policy and replace the ILM rule used to store objects in the Cloud Tiering bucket with a cloned ILM rule that stores the same objects in the Cloud Storage Pool.



When objects are stored in a Cloud Storage Pool, copies of those objects cannot also be stored within StorageGRID. If the ILM rule you are currently using for Cloud Tiering is configured to store objects in multiple locations at the same time, consider whether you still want to perform this optional migration because you will lose that functionality. If you continue with this migration, you must create new rules instead of cloning the existing ones.

### Steps

1. Create a Cloud Storage Pool.

Use a new S3 bucket for the Cloud Storage Pool to ensure it contains only the data managed by the Cloud Storage Pool.

2. Locate any ILM rules in the active ILM policy that cause objects to be stored in the Cloud Tiering bucket.
3. Clone each of these rules.
4. In the cloned rules, change the placement location to the new Cloud Storage Pool.
5. Save the cloned rules.
6. Create a new policy that uses the new rules.
7. Simulate and activate the new policy.

When the new policy is activated and ILM evaluation occurs, the objects are moved from the S3 bucket configured for Cloud Tiering to the S3 bucket configured for the Cloud Storage Pool. The usable space on

the grid is not affected. After the objects are moved to the Cloud Storage Pool, they are removed from the Cloud Tiering bucket.

## Related information

[Manage objects with ILM](#)

### Archiving to tape through TSM middleware

You can configure an Archive Node to target a Tivoli Storage Manager (TSM) server that provides a logical interface for storing and retrieving object data to random or sequential access storage devices, including tape libraries.

The Archive Node's ARC service acts as a client to the TSM server, using Tivoli Storage Manager as middleware for communicating with the archival storage system.

### TSM management classes

Management classes defined by the TSM middleware outline how the TSM's backup and archive operations function, and can be used to specify rules for content that are applied by the TSM server. Such rules operate independently of the StorageGRID system's ILM policy, and must be consistent with the StorageGRID system's requirement that objects are stored permanently and are always available for retrieval by the Archive Node. After object data is sent to a TSM server by the Archive Node, the TSM lifecycle and retention rules are applied while the object data is stored to tape managed by the TSM server.

The TSM management class is used by the TSM server to apply rules for data location or retention after objects are sent to the TSM server by the Archive Node. For example, objects identified as database backups (temporary content that can be overwritten with newer data) could be treated differently than application data (fixed content that must be retained indefinitely).

### Configuring connections to TSM middleware

Before the Archive Node can communicate with Tivoli Storage Manager (TSM) middleware, you must configure a number of settings.

#### What you'll need

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

#### About this task

Until these settings are configured, the ARC service remains in a Major alarm state as it is unable to communicate with the Tivoli Storage Manager.

#### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.



## Configuration: ARC (DC1-ARC1-98-165) - Target

Updated: 2015-09-28 09:56:36 PDT

Target Type:

Tivoli Storage Manager State:

### Target (TSM) Account

Server IP or Hostname:

Server Port:

Node Name:

User Name:

Password:

Management Class:

Number of Sessions:

Maximum Retrieve Sessions:

Maximum Store Sessions:

Apply Changes

4. From the **Target Type** drop-down list, select **Tivoli Storage Manager (TSM)**.
5. For the **Tivoli Storage Manager State**, select **Offline** to prevent retrievals from the TSM middleware server.

By default, the Tivoli Storage Manager State is set to Online, which means that the Archive Node is able to retrieve object data from the TSM middleware server.

6. Complete the following information:
  - **Server IP or Hostname:** Specify the IP address or fully qualified domain name of the TSM middleware server used by the ARC service. The default IP address is 127.0.0.1.
  - **Server Port:** Specify the port number on the TSM middleware server that the ARC service will connect to. The default is 1500.
  - **Node Name:** Specify the name of the Archive Node. You must enter the name (arc-user) that you registered on the TSM middleware server.
  - **User Name:** Specify the user name the ARC service uses to log in to the TSM server. Enter the default user name (arc-user) or the administrative user you specified for the Archive Node.
  - **Password:** Specify the password used by the ARC service to log in to the TSM server.
  - **Management Class:** Specify the default management class to use if a management class is not specified when the object is being saved to the StorageGRID system, or the specified management class is not defined on the TSM middleware server.
  - **Number of Sessions:** Specify the number of tape drives on the TSM middleware server that are dedicated to the Archive Node. The Archive Node concurrently creates a maximum of one session per mount point plus a small number of additional sessions (less than five).

You must change this value to be the same as the value set for MAXNUMMP (maximum number of mount points) when the Archive Node was registered or updated. (In the register command, the default value of MAXNUMMP used is 1, if no value is set.)

You must also change the value of MAXSESSIONS for the TSM server to a number that is at least as large as the Number of Sessions set for the ARC service. The default value of MAXSESSIONS on the TSM server is 25.

- **Maximum Retrieve Sessions:** Specify the maximum number of sessions that the ARC service can open to the TSM middleware server for retrieve operations. In most cases, the appropriate value is Number of Sessions minus Maximum Store Sessions. If you need to share one tape drive for storage and retrieval, specify a value equal to the Number of Sessions.
- **Maximum Store Sessions:** Specify the maximum number of concurrent sessions that the ARC service can open to the TSM middleware server for archive operations.

This value should be set to one except when the targeted archival storage system is full and only retrievals can be performed. Set this value to zero to use all sessions for retrievals.

7. Click **Apply Changes**.

### Optimizing an Archive Node for TSM middleware sessions

You can optimize the performance of an Archive Node that connects to Tivoli Server Manager (TSM) by configuring the Archive Node's sessions.

#### What you'll need

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

#### About this task

Typically, the number of concurrent sessions that the Archive Node has open to the TSM middleware server is set to the number of tape drives the TSM server has dedicated to the Archive Node. One tape drive is allocated for storage while the rest are allocated for retrieval. However, in situations where a Storage Node is being rebuilt from Archive Node copies or the Archive Node is operating in Read-only mode, you can optimize TSM server performance by setting the maximum number of retrieve sessions to be the same as number of concurrent sessions. The result is that all drives can be used concurrently for retrieval, and, at most, one of these drives can also be used for storage if applicable.

#### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.
4. Change **Maximum Retrieve Sessions** to be the same as **Number of Sessions**.

Overview	Alarms	Reports	Configuration
Main	Alarms		



## Configuration: ARC (DC1-ARC1-98-165) - Target

Updated: 2015-09-28 09:56:36 PDT

---

Target Type:

Tivoli Storage Manager State:

### Target (TSM) Account

---

Server IP or Hostname:	<input type="text" value="10.10.10.123"/>
Server Port:	<input type="text" value="1500"/>
Node Name:	<input type="text" value="ARC-USER"/>
User Name:	<input type="text" value="arc-user"/>
Password:	<input type="password" value="•••••"/>
Management Class:	<input type="text" value="sg-mgmtclass"/>
Number of Sessions:	<input type="text" value="2"/>
Maximum Retrieve Sessions:	<input type="text" value="2"/>
Maximum Store Sessions:	<input type="text" value="1"/>

[Apply Changes !\[\]\(8c09a179b36b7742df87bace4677eaa5\_img.jpg\)](#)

5. Click **Apply Changes**.

### Configuring the archive state and counters for TSM

If your Archive Node connects to a TSM middleware server, you can configure an Archive Node's archive store state to Online or Offline. You can also disable the archive store when the Archive Node first starts up, or reset the failure count being tracked for the associated alarm.

#### What you'll need

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

#### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Store**.
3. Select **Configuration > Main**.

Overview Alarms Reports **Configuration**

Main Alarms

 **Configuration: ARC (DC1-ARC1-98-165) - Store**  
Updated: 2015-09-29 17:10:12 PDT

---

Store State

Archive Store Disabled on Startup

Reset Store Failure Count

**Apply Changes** 

4. Modify the following settings, as necessary:

- Store State: Set the component state to either:
  - Online: The Archive Node is available to process object data for storage to the archival storage system.
  - Offline: The Archive Node is not available to process object data for storage to the archival storage system.
- Archive Store Disabled on Startup: When selected, the Archive Store component remains in the Read-only state when restarted. Used to persistently disable storage to the targeted the archival storage system. Useful when the targeted the archival storage system is unable to accept content.
- Reset Store Failure Count: Reset the counter for store failures. This can be used to clear the ARVF (Stores Failure) alarm.

5. Click **Apply Changes**.

**Related information**

[Managing an Archive Node when TSM server reaches capacity](#)

**Managing an Archive Node when TSM server reaches capacity**

The TSM server has no way to notify the Archive Node when either the TSM database or the archival media storage managed by the TSM server is nearing capacity. The Archive Node continues to accept object data for transfer to the TSM server after the TSM server stops accepting new content. This content cannot be written to media managed by the TSM server. An alarm is triggered if this happens. This situation can be avoided through proactive monitoring of the TSM server.

**What you'll need**


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

**About this task**

To prevent the ARC service from sending further content to the TSM server, you can take the Archive Node offline by taking its **ARC > Store** component offline. This procedure can also be useful in preventing alarms when the TSM server is unavailable for maintenance.

## Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Store**.
3. Select **Configuration > Main**.



4. Change **Store State** to *Offline*.
5. Select **Archive Store Disabled on Startup**.
6. Click **Apply Changes**.

## Setting Archive Node to read-only if TSM middleware reaches capacity

If the targeted TSM middleware server reaches capacity, the Archive Node can be optimized to only perform retrievals.

### What you'll need

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

## Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Select **Configuration > Main**.
4. Change Maximum Retrieve Sessions to be the same as the number of concurrent sessions listed in Number of Sessions.
5. Change Maximum Store Sessions to 0.



Changing Maximum Store Sessions to 0 is not necessary if the Archive Node is Read-only. Store sessions will not be created.

6. Click **Apply Changes**.

## Configuring Archive Node retrieve settings

You can configure the retrieve settings for an Archive Node to set the state to Online or

Offline, or reset the failure counts being tracked for the associated alarms.

### What you'll need

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

### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Retrieve**.
3. Select **Configuration > Main**.

Configuration: ARC (DC1-ARC1-98-165) - Retrieve  
Updated: 2015-05-07 12:24:45 PDT

Retrieve State	Online
Reset Request Failure Count	<input type="checkbox"/>
Reset Verification Failure Count	<input type="checkbox"/>

Apply Changes

4. Modify the following settings, as necessary:
  - **Retrieve State:** Set the component state to either:
    - Online: The grid node is available to retrieve object data from the archival media device.
    - Offline: The grid node is not available to retrieve object data.
  - Reset Request Failures Count: Select the check box to reset the counter for request failures. This can be used to clear the ARRF (Request Failures) alarm.
  - Reset Verification Failure Count: Select the check box to reset the counter for verification failures on retrieved object data. This can be used to clear the ARRV (Verification Failures) alarm.
5. Click **Apply Changes**.

### Configuring Archive Node replication

You can configure the replication settings for an Archive Node and disable inbound and outbound replication, or reset the failure counts being tracked for the associated alarms.

### What you'll need

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

### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Replication**.

3. Select **Configuration > Main**.

Configuration: ARC (DC1-ARC1-98-165) - Replication  
Updated: 2015-05-07 12:21:53 PDT

Reset Inbound Replication Failure Count

Reset Outbound Replication Failure Count

**Inbound Replication**

Disable Inbound Replication

**Outbound Replication**

Disable Outbound Replication

Apply Changes

4. Modify the following settings, as necessary:

- **Reset Inbound Replication Failure Count:** Select to reset the counter for inbound replication failures. This can be used to clear the RIRF (Inbound Replications — Failed) alarm.
- **Reset Outbound Replication Failure Count:** Select to reset the counter for outbound replication failures. This can be used to clear the RORF (Outbound Replications — Failed) alarm.
- **Disable Inbound Replication:** Select to disable inbound replication as part of a maintenance or testing procedure. Leave cleared during normal operation.

When inbound replication is disabled, object data can be retrieved from the ARC service for replication to other locations in the StorageGRID system, but objects cannot be replicated to this ARC service from other system locations. The ARC service is read-only.

- **Disable Outbound Replication:** Select the check box to disable outbound replication (including content requests for HTTP retrievals) as part of a maintenance or testing procedure. Leave unchecked during normal operation.

When outbound replication is disabled, object data can be copied to this ARC service to satisfy ILM rules, but object data cannot be retrieved from the ARC service to be copied to other locations in the StorageGRID system. The ARC service is write-only.

5. Click **Apply Changes**.

### Setting Custom alarms for the Archive Node

You should establish Custom alarms for the ARQL and ARRL attributes that are used to monitor the speed and efficiency of object data retrieval from the archival storage system by the Archive Node.

- **ARQL:** Average Queue Length. The average time, in microseconds, that object data is queued for retrieval from the archival storage system.
- **ARRL:** Average Request Latency. The average time, in microseconds, needed by the Archive Node to

retrieve object data from the archival storage system.

The acceptable values for these attributes depend on how the archival storage system is configured and used. (Go to **ARC > Retrieve > Overview > Main.**) The values set for request timeouts and the number of sessions made available for retrieve requests are particularly influential.

After integration is complete, monitor the Archive Node's object data retrievals to establish values for normal retrieval times and queue lengths. Then, create Custom alarms for ARQL and ARRL that will trigger if an abnormal operating condition arises.

#### **Related information**

[Monitor & troubleshoot](#)

#### **Integrating Tivoli Storage Manager**

This section includes best practices and set-up information for integrating an Archive Node with a Tivoli Storage Manager (TSM) server, including Archive Node operational details that impact the configuration of the TSM server.

- [Archive Node configuration and operation](#)
- [Configuration best practices](#)
- [Completing the Archive Node setup](#)

#### **Archive Node configuration and operation**

Your StorageGRID system manages the Archive Node as a location where objects are stored indefinitely and are always accessible.

When an object is ingested, copies are made to all required locations, including Archive Nodes, based on the Information Lifecycle Management (ILM) rules defined for your StorageGRID system. The Archive Node acts as a client to a TSM server, and the TSM client libraries are installed on the Archive Node by the StorageGRID software installation process. Object data directed to the Archive Node for storage is saved directly to the TSM server as it is received. The Archive Node does not stage object data before saving it to the TSM server, nor does it perform object aggregation. However, the Archive Node can submit multiple copies to the TSM server in a single transaction when data rates warrant.

After the Archive Node saves object data to the TSM server, the object data is managed by the TSM server using its lifecycle/retention policies. These retention policies must be defined to be compatible with the operation of the Archive Node. That is, object data saved by the Archive Node must be stored indefinitely and must always be accessible by the Archive Node, unless it is deleted by the Archive Node.

There is no connection between the StorageGRID system's ILM rules and the TSM server's lifecycle/retention policies. Each operates independently of the other; however, as each object is ingested into the StorageGRID system, you can assign it a TSM management class. This management class is passed to the TSM server along with object data. Assigning different management classes to different object types permits you to configure the TSM server to place object data in different storage pools, or to apply different migration or retention policies as required. For example, objects identified as database backups (temporary content that can be overwritten with newer data) might be treated differently than application data (fixed content that must be retained indefinitely).

The Archive Node can be integrated with a new or an existing TSM server; it does not require a dedicated TSM server. TSM servers can be shared with other clients, provided that the TSM server is sized appropriately for the maximum expected load. TSM must be installed on a server or virtual machine separate from the Archive

Node.

It is possible to configure more than one Archive Node to write to the same TSM server; however, this configuration is only recommended if the Archive Nodes write different sets of data to the TSM server. Configuring more than one Archive Node to write to the same TSM server is not recommended when each Archive Node writes copies of the same object data to the archive. In the latter scenario, both copies are subject to a single point of failure (the TSM server) for what are supposed to be independent, redundant copies of object data.

Archive Nodes do not make use of the Hierarchical Storage Management (HSM) component of TSM.

### **Configuration best practices**

When you are sizing and configuring your TSM server there are best practices you should apply to optimize it to work with the Archive Node.

When sizing and configuring the TSM server, you should consider the following factors:

- Because the Archive Node does not aggregate objects before saving them to the TSM server, the TSM database must be sized to hold references to all objects that will be written to the Archive Node.
- Archive Node software cannot tolerate the latency involved in writing objects directly to tape or other removable media. Therefore, the TSM server must be configured with a disk storage pool for the initial storage of data saved by the Archive Node whenever removable media are used.
- You must configure TSM retention policies to use event-based retention. The Archive Node does not support creation-based TSM retention policies. Use the following recommended settings of `retmin=0` and `retver=0` in the retention policy (which indicates that retention begins when the Archive Node triggers a retention event, and is retained for 0 days after that). However, these values for `retmin` and `retver` are optional.

The disk pool must be configured to migrate data to the tape pool (that is, the tape pool must be the `NXTSTGPOOL` of the disk pool). The tape pool must not be configured as a copy pool of the disk pool with simultaneous write to both pools (that is, the tape pool cannot be a `COPYSTGPOOL` for the disk pool). To create offline copies of the tapes containing Archive Node data, configure the TSM server with a second tape pool that is a copy pool of the tape pool used for Archive Node data.

### **Completing the Archive Node setup**

The Archive Node is not functional after you complete the installation process. Before the StorageGRID system can save objects to the TSM Archive Node, you must complete the installation and configuration of the TSM server and configure the Archive Node to communicate with the TSM server.

For more information about optimizing TSM retrieval and store sessions, see information about managing archival storage.

- [Managing Archive Nodes](#)

Refer to the following IBM documentation, as necessary, as you prepare your TSM server for integration with the Archive Node in a StorageGRID system:

- [IBM Tape Device Drivers Installation and User's Guide](#)
- [IBM Tape Device Drivers Programming Reference](#)

## Installing a new TSM server

You can integrate the Archive Node with either a new or an existing TSM server. If you are installing a new TSM server, follow the instructions in your TSM documentation to complete the installation.



An Archive Node cannot be co-hosted with a TSM server.

## Configuring the TSM server

This section includes sample instructions for preparing a TSM server following TSM best practices.

The following instructions guide you through the process of:

- Defining a disk storage pool, and a tape storage pool (if required) on the TSM server
- Defining a domain policy that uses the TSM management class for the data saved from the Archive Node, and registering a node to use this domain policy

These instructions are provided for your guidance only; they are not intended to replace TSM documentation, or to provide complete and comprehensive instructions suitable for all configurations. Deployment specific instructions should be provided by a TSM administrator who is familiar both with your detailed requirements, and with the complete set of TSM Server documentation.

### Defining TSM tape and disk storage pools

The Archive Node writes to a disk storage pool. To archive content to tape, you must configure the disk storage pool to move content to a tape storage pool.

#### About this task

For a TSM server, you must define a tape storage pool and a disk storage pool within Tivoli Storage Manager. After the disk pool is defined, create a disk volume and assign it to the disk pool. A tape pool is not required if your TSM server uses disk-only storage.

You must complete a number of steps on your TSM server before you can create a tape storage pool. (Create a tape library and at least one drive in the tape library. Define a path from the server to the library and from the server to the drives, and then define a device class for the drives.) The details of these steps can vary depending upon the hardware configuration and storage requirements of the site. For more information, see the TSM documentation.

The following set of instructions illustrates the process. You should be aware that the requirements for your site could be different depending on the requirements of your deployment. For configuration details and for instructions, see the TSM documentation.



You must log onto the server with administrative privileges and use the `dsmdmcc` tool to execute the following commands.

#### Steps

1. Create a tape library.

```
define library tapelibrary libtype=scsi
```

Where *tapelibrary* is an arbitrary name chosen for the tape library, and the value of *libtype* can vary depending upon the type of tape library.

2. Define a path from the server to the tape library.

```
define path servername tapelibrary srctype=server desttype=library device=lib-  
devicename
```

- *servername* is the name of the TSM server
- *tapelibrary* is the tape library name you defined
- *lib-devicename* is the device name for the tape library

3. Define a drive for the library.

```
define drive tapelibrary drivename
```

- *drivename* is the name you want to specify for the drive
- *tapelibrary* is the tape library name you defined

You might want to configure an additional drive or drives, depending upon your hardware configuration. (For example, if the TSM server is connected to a Fibre Channel switch that has two inputs from a tape library, you might want to define a drive for each input.)

4. Define a path from the server to the drive you defined.

```
define path servername drivename srctype=server desttype=drive  
library=tapelibrary device=drive-dname
```

- *drive-dname* is the device name for the drive
- *tapelibrary* is the tape library name you defined

Repeat for each drive that you have defined for the tape library, using a separate *drivename* and *drive-dname* for each drive.

5. Define a device class for the drives.

```
define devclass DeviceClassName devtype=lto library=tapelibrary  
format=tapetype
```

- *DeviceClassName* is the name of the device class
- *lto* is the type of drive connected to the server
- *tapelibrary* is the tape library name you defined
- *tapetype* is the tape type; for example, *ultrium3*

6. Add tape volumes to the inventory for the library.

```
checkin libvolume tapelibrary
```

*tapelibrary* is the tape library name you defined.

## 7. Create the primary tape storage pool.

```
define stgpool SGWSTapePool DeviceClassName description=description  
collocate=filespace maxxscratch=XX
```

- *SGWSTapePool* is the name of the Archive Node's tape storage pool. You can select any name for the tape storage pool (as long as the name uses the syntax conventions expected by the TSM server).
- *DeviceClassName* is the name of the device class name for the tape library.
- *description* is a description of the storage pool that can be displayed on the TSM server using the query stgpool command. For example: "Tape storage pool for the Archive Node."
- *collocate=filespace* specifies that the TSM server should write objects from the same file space into a single tape.
- *XX* is one of the following:
  - The number of empty tapes in the tape library (in the case that the Archive Node is the only application using the library).
  - The number of tapes allocated for use by the StorageGRID system (in instances where the tape library is shared).

## 8. On a TSM server, create a disk storage pool. At the TSM server's administrative console, enter

```
define stgpool SGWSDiskPool disk description=description  
maxsize=maximum_file_size nextstgpool=SGWSTapePool highmig=percent_high  
lowmig=percent_low
```

- *SGWSDiskPool* is the name of the Archive Node's disk pool. You can select any name for the disk storage pool (as long as the name uses the syntax conventions expected by the TSM).
- *description* is a description of the storage pool that can be displayed on the TSM server using the query stgpool command. For example, "Disk storage pool for the Archive Node."
- *maximum\_file\_size* forces objects larger than this size to be written directly to tape, rather than being cached in the disk pool. It is recommended to set *maximum\_file\_size* to 10 GB.
- *nextstgpool=SGWSTapePool* refers the disk storage pool to the tape storage pool defined for the Archive Node.
- *percent\_high* sets the value at which the disk pool begins to migrate its contents to the tape pool. It is recommended to set *percent\_high* to 0 so that data migration begins immediately
- *percent\_low* sets the value at which migration to the tape pool stops. It is recommended to set *percent\_low* to 0 to clear out the disk pool.

## 9. On a TSM server, create a disk volume (or volumes) and assign it to the disk pool.

```
define volume SGWSDiskPool volume_name formatsize=size
```

- *SGWSDiskPool* is the disk pool name.
- *volume\_name* is the full path to the location of the volume (for example, `/var/local/arc/stage6.dsm`) on the TSM server where it writes the contents of the disk pool in preparation for transfer to tape.
- *size* is the size, in MB, of the disk volume.

For example, to create a single disk volume such that the contents of a disk pool fill a single tape, set the value of size to 200000 when the tape volume has a capacity of 200 GB.

However, it might be desirable to create multiple disk volumes of a smaller size, as the TSM server can write to each volume in the disk pool. For example, if the tape size is 250 GB, create 25 disk volumes with a size of 10 GB (10000) each.

The TSM server preallocates space in the directory for the disk volume. This can take some time to complete (more than three hours for a 200 GB disk volume).

## Defining a domain policy and registering a node

You need to define a domain policy that uses the TSM management class for the data saved from the Archive Node, and then register a node to use this domain policy.



Archive Node processes can leak memory if the client password for the Archive Node in Tivoli Storage Manager (TSM) expires. Ensure that the TSM server is configured so the client username/password for the Archive Node never expires.

When registering a node on the TSM server for the use of the Archive Node (or updating an existing node), you must specify the number of mount points that the node can use for write operations by specifying the MAXNUMMP parameter to the REGISTER NODE command. The number of mount points is typically equivalent to the number of tape drive heads allocated to the Archive Node. The number specified for MAXNUMMP on the TSM server must be at least as large as the value set for the **ARC > Target > Configuration > Main > Maximum Store Sessions** for the Archive Node, which is set to a value of 0 or 1, as concurrent store sessions are not supported by the Archive Node.

The value of MAXSESSIONS set for the TSM server controls the maximum number of sessions that can be opened to the TSM server by all client applications. The value of MAXSESSIONS specified on the TSM must be at least as large as the value specified for **ARC > Target > Configuration > Main > Number of Sessions** in the Grid Manager for the Archive Node. The Archive Node concurrently creates at most one session per mount point plus a small number (< 5) of additional sessions.

The TSM node assigned to the Archive Node uses a custom domain policy `tsm-domain`. The `tsm-domain` domain policy is a modified version of the “standard” domain policy, configured to write to tape and with the archive destination set to be the StorageGRID system’s storage pool (`SGWSDiskPool`).



You must log in to the TSM server with administrative privileges and use the `dsmadm` tool to create and activate the domain policy.

## Creating and activating the domain policy

You must create a domain policy and then activate it to configure the TSM server to save data sent from the Archive Node.

### Steps

1. Create a domain policy.

```
copy domain standard tsm-domain
```

2. If you are not using an existing management class, enter one of the following:

```
define policyset tsm-domain standard
```

```
define mgmtclass tsm-domain standard default
```

*default* is the default management class for the deployment.

3. Create a copygroup to the appropriate storage pool. Enter (on one line):

```
define copygroup tsm-domain standard default type=archive  
destination=SGWSDiskPool retinit=event retmin=0 retver=0
```

*default* is the default Management Class for the Archive Node. The values of `retinit`, `retmin`, and `retver` have been chosen to reflect the retention behavior currently used by the Archive Node



Do not set `retinit` to `retinit=create`. Setting `retinit=create` blocks the Archive Node from deleting content since retention events are used to remove content from the TSM server.

4. Assign the management class to be the default.

```
assign defmgmtclass tsm-domain standard default
```

5. Set the new policy set as active.

```
activate policyset tsm-domain standard
```

Ignore the “no backup copy group” warning that appears when you enter the activate command.

6. Register a node to use the new policy set on the TSM server. On the TSM server, enter (on one line):

```
register node arc-user arc-password passexp=0 domain=tsm-domain  
MAXNUMMP=number-of-sessions
```

`arc-user` and `arc-password` are same client node name and password as you define on the Archive Node, and the value of `MAXNUMMP` is set to the number of tape drives reserved for Archive Node store sessions.



By default, registering a node creates an administrative user ID with client owner authority, with the password defined for the node.

## Migrating data into StorageGRID

You can migrate large amounts of data to the StorageGRID system while simultaneously using the StorageGRID system for day-to-day operations.

The following section is a guide to understanding and planning a migration of large amounts of data into the StorageGRID system. It is not a general guide to data migration, and it does not include detailed steps for performing a migration. Follow the guidelines and instructions in this section to ensure that data is migrated efficiently into the StorageGRID system without interfering with day-to-day operations, and that the migrated data is handled appropriately by the StorageGRID system.

- [Confirming capacity of the StorageGRID system](#)

- [Determining the ILM policy for migrated data](#)
- [Impact of migration on operations](#)
- [Scheduling data migration](#)
- [Monitoring data migration](#)
- [Creating custom notifications for migration alarms](#)

## Confirming capacity of the StorageGRID system

Before migrating large amounts of data into the StorageGRID system, confirm that the StorageGRID system has the disk capacity to handle the anticipated volume.

If the StorageGRID system includes an Archive Node and a copy of migrated objects has been saved to nearline storage (such as tape), ensure that the Archive Node's storage has sufficient capacity for the anticipated volume of migrated data.

As part of the capacity assessment, look at the data profile of the objects you plan to migrate and calculate the amount of disk capacity required. For details about monitoring the disk capacity of your StorageGRID system, see the instructions for monitoring and troubleshooting StorageGRID.

### Related information

[Monitor & troubleshoot](#)

[Managing Storage Nodes](#)

## Determining the ILM policy for migrated data

The StorageGRID system's ILM policy determines how many copies are made, the locations to which copies are stored, and for how long these copies are retained. An ILM policy consists of a set of ILM rules that describe how to filter objects and manage object data over time.

Depending on how migrated data is used and your requirements for migrated data, you might want to define unique ILM rules for migrated data that are different from the ILM rules used for day-to-day operations. For example, if there are different regulatory requirements for day-to-day data management than there are for the data that is included in the migration, you might want a different number of copies of the migrated data on a different grade of storage.

You can configure rules that apply exclusively to migrated data if it is possible to uniquely distinguish between migrated data and object data saved from day-to-day operations.

If you can reliably distinguish between the types of data using one of the metadata criteria, you can use this criteria to define an ILM rule that applies only to migrated data.

Before beginning data migration, ensure that you understand the StorageGRID system's ILM policy and how it will apply to migrated data, and that you have made and tested any changes to the ILM policy.



An ILM policy that has been incorrectly specified can cause unrecoverable data loss. Carefully review all changes you make to an ILM policy before activating it to make sure the policy will work as intended.

### Related information

### Impact of migration on operations

A StorageGRID system is designed to provide efficient operation for object storage and retrieval, and to provide excellent protection against data loss through the seamless creation of redundant copies of object data and metadata.

However, data migration must be carefully managed according to the instructions in this chapter to avoid having an impact on day-to-day system operations, or, in extreme cases, placing data at risk of loss in case of a failure in the StorageGRID system.

Migration of large quantities of data places additional load on the system. When the StorageGRID system is heavily loaded, it responds more slowly to requests to store and retrieve objects. This can interfere with store and retrieve requests which are integral to day-to-day operations. Migration can also cause other operational issues. For example, when a Storage Node is nearing capacity, the heavy intermittent load due to batch ingest can cause the Storage Node to cycle between read-only and read-write, generating notifications.

If the heavy loading persists, queues can develop for various operations that the StorageGRID system must perform to ensure full redundancy of object data and metadata.

Data migration must be carefully managed according to the guidelines in this document to ensure safe and efficient operation of the StorageGRID system during migration. When migrating data, ingest objects in batches or continuously throttle ingest. Then, continuously monitor the StorageGRID system to ensure that various attribute values are not exceeded.

### Scheduling data migration

Avoid migrating data during core operational hours. Limit data migration to evenings, weekends, and other times when system usage is low.

If possible, do not schedule data migration during periods of high activity. However, if it is not practical to completely avoid the high activity period, it is safe to proceed as long as you closely monitor the relevant attributes and take action if they exceed acceptable values.

### Related information

[Monitoring data migration](#)

### Monitoring data migration

Data migration must be monitored and adjusted as necessary to ensure data is placed according to the ILM policy within the required timeframe.

This table lists the attributes you must monitor during data migration, and the issues that they represent.

If you use traffic classification policies with rate limits to throttle ingest, you can monitor the observed rate in conjunction with the statistics described in the following table and reduce the limits if necessary.

Monitor	Description
Number of objects waiting for ILM evaluation	<ol style="list-style-type: none"> <li>1. Select <b>Support &gt; Tools &gt; Grid Topology</b>.</li> <li>2. Select <b>deployment &gt; Overview &gt; Main</b>.</li> <li>3. In the ILM Activity section, monitor the number of objects shown for the following attributes: <ul style="list-style-type: none"> <li>◦ <b>Awaiting - All (XQUZ)</b>: The total number of objects awaiting ILM evaluation.</li> <li>◦ <b>Awaiting - Client (XCQZ)</b>: The total number of objects awaiting ILM evaluation from client operations (for example, ingest).</li> </ul> </li> <li>4. If the number of objects shown for either of these attributes exceeds 100,000, throttle the ingest rate of objects to reduce the load on the StorageGRID system.</li> </ol>
Targeted archival system's storage capacity	If the ILM policy saves a copy of the migrated data to a targeted archival storage system (tape or the cloud), monitor the capacity of the targeted archival storage system to ensure that there is sufficient capacity for the migrated data.
<b>Archive Node &gt; ARC &gt; Store</b>	If an alarm for the <b>Store Failures (ARVF)</b> attribute is triggered, the targeted archival storage system might have reached capacity. Check the targeted archival storage system and resolve any issues that triggered an alarm.

### Creating custom notifications for migration alarms

You might want StorageGRID to send alert notifications or alarm (legacy system) notifications to the system administrator responsible for monitoring migration if certain values exceed recommended thresholds.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have configured email settings for alert (or alarm) notifications.

#### Steps

1. Create a custom alert rule or a Global Custom alarm for each Prometheus metric or StorageGRID attribute you want to monitor during data migration.

Alerts are triggered based on Prometheus metric values. Alarms are triggered based on attribute values. See the instructions for monitoring and troubleshooting StorageGRID for more information.

2. Disable the custom alert rule or the Global Custom alarm after data migration is complete.

Note that Global Custom alarms override Default alarms.

#### Related information

[Monitor & troubleshoot](#)

# Manage objects with ILM

Learn how to manage objects with information lifecycle rules and policies and how to use S3 Object Lock to comply with regulations for object retention.

- [Managing objects with information lifecycle management](#)
- [Managing objects with S3 Object Lock](#)
- [Example ILM rules and policies](#)

## Managing objects with information lifecycle management

You manage the objects in a StorageGRID system by configuring information lifecycle management (ILM) rules and policies. The ILM rules and policies instruct StorageGRID how to create and distribute copies of object data and how to manage those copies over time.

Designing and implementing ILM rules and the ILM policy 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.

- [How ILM operates throughout an object's life](#)
- [What an ILM policy is](#)
- [What an ILM rule is](#)
- [Creating storage grades, storage pools, EC profiles, and regions](#)
- [Creating an ILM rule](#)
- [Creating an ILM policy](#)
- [Working with ILM rules and ILM policies](#)

### 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 or Swift 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


the instructions for monitoring and troubleshooting StorageGRID.

- **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 cannot 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.



### Related information

[Monitor & troubleshoot](#)

### How objects are ingested

StorageGRID protects objects during ingest either by performing synchronous placement or by performing dual commit, as specified in the ILM rule that matches the objects.

When an S3 or Swift client stores an object to the grid, StorageGRID ingests the object using one of these two methods:

- **Synchronous placement:** StorageGRID immediately creates all object copies that are needed to meet ILM requirements. StorageGRID sends an “ingest successful” message to the client when all copies are created.

If StorageGRID cannot immediately create all object copies (for example, because a required location is temporarily unavailable), it either sends an “ingest failed” message to the client, or it falls back to creating interim object copies and evaluating ILM later, depending on the choice you made when you created the ILM rule.

- **Dual commit:** StorageGRID immediately creates two interim copies of the object, each on a different Storage Node, and sends an “ingest successful” message to the client. StorageGRID then queues the object for ILM evaluation.

When StorageGRID performs the ILM evaluation, it first checks to see if the interim copies satisfy the placement instructions in the ILM rule. For example, the two interim copies might satisfy the instructions in a two-copy ILM rule, but they would not satisfy the instructions in an erasure-coding rule. If the interim copies do not satisfy the ILM instructions, StorageGRID creates new object copies and deletes any interim copies that are not needed.

If StorageGRID cannot create two interim copies (for example, if a network issue prevents the second copy from being made), StorageGRID does not retry. Ingest fails.



S3 or Swift clients can specify that StorageGRID create a single interim copy at ingest by specifying `REDUCED_REDUNDANCY` for the storage class. See the instructions for implementing an S3 or Swift client for more information.

By default, StorageGRID uses synchronous placement to protect objects during ingest.

### Related information

[Data-protection options for ingest](#)

[Use S3](#)

[Use Swift](#)

### Data-protection options for ingest

When you create an ILM rule, you specify one of three options for protecting objects at ingest: Dual commit, Balanced, or Strict. 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.

#### 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.

## When to use the Dual commit option

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 do not 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.

## 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 cannot create all copies, for example, because a required storage location is temporarily unavailable. The client must retry the operation.

## When to use the Strict option

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 center.

[Example 5: ILM rules and policy for Strict ingest behavior](#)

## Balanced


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 cannot immediately make all copies, it uses Dual commit instead.

## When to use the Balanced option

Use the Balanced option to achieve the best combination of data protection, grid performance, and ingest success. Balanced is the default option in the ILM rule wizard.

## Flowchart of three ingest options

The flowchart shows what happens when objects are matched by an ILM rule that uses one of these ingest options.



## Related information

[How objects are ingested](#)

## Advantages, disadvantages, and limitations of the data-protection 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.

### 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 for 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 and 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 cannot 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 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 do not 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 are not made immediately. Placement changes are made when ILM is re-evaluated by normal background ILM processes. If required placement changes cannot 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 or Strict options

The Balanced or Strict options cannot be used for ILM rules that have any of these placement instructions:

- Placement in a Cloud Storage Pool at day 0.
- Placement in an Archive Node at day 0.
- Placements in a Cloud Storage Pool or an Archive Node when the rule has a User Defined Creation Time as its Reference Time.

These restrictions exist because StorageGRID cannot synchronously make copies to a Cloud Storage Pool or an Archive Node, and a User Defined Creation Time could resolve to the present.

## How ILM rules and consistency controls interact to affect data protection

Both your ILM rule and your choice of consistency control 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 control used when an object is stored affects the initial placement of object metadata. 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 level and ingest behavior can provide better initial data protection and more predictable system responses.

Here is a brief summary of the consistency controls 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.
- **available** (eventual consistency for HEAD operations): Behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations.



Before selecting a consistency level, read the full description of these settings in the instructions for creating an S3 or Swift client application. You should understand the benefits and limitations before changing the default value.

## Example of how the consistency control and ILM rule can interact

Suppose you have a two-site grid with the following ILM rule and the following consistency level setting:

- **ILM rule**: Create two object copies, one at the local site and one at a remote site. The Strict ingest behavior is selected.
- **Consistency level**: “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 level, 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 cannot be retrieved.

The inter-relationship between consistency levels and ILM rules can be complex. Contact NetApp if you require assistance.

### **Related information**

[What replication is](#)

[What erasure coding is](#)

[What erasure-coding schemes are](#)

[Example 5: ILM rules and policy for Strict ingest behavior](#)

[Use S3](#)

[Use Swift](#)

### **How objects are stored (replication or erasure coding)**

StorageGRID can protect objects against loss either by storing replicated copies or by storing erasure-coded copies. You specify the type of copies to create in the placement instructions of ILM rules.


- [What replication is](#)
- [Why you should not use single-copy replication](#)
- [What erasure coding is](#)
- [What erasure-coding schemes are](#)
- [Advantages, disadvantages, and requirements for erasure coding](#)

### **What replication is**

Replication is one of two methods used by StorageGRID to store object data. When objects match an ILM rule that uses replication, the system creates exact copies of object data and stores the copies on Storage Nodes or Archive 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 a storage pool is](#)

[Using multiple storage pools for cross-site replication](#)


#### 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.




Do not 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 erasure coding is

Erasur coding is the second method used by StorageGRID to store object data. 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.


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 requires a minimum of nine Storage Nodes, with three Storage Nodes at each of three different sites. 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 data center 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 a storage pool is](#)

[What erasure-coding schemes are](#)

[Configuring Erasure Coding profiles](#)

#### What erasure-coding schemes are

When you configure the Erasure Coding profile for an ILM rule, you select an available erasure-coding scheme based on how many Storage Nodes and sites make up the storage pool you plan to use. Erasure-coding schemes control how many data fragments and how many parity fragments are created for each object.

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. An object can sustain up to  $m$  lost or corrupt fragments.  $k$  fragments are needed to retrieve or repair an object.

When configuring an Erasure Coding profile, use the following guidelines for storage pools:

- The storage pool must include three or more sites, or exactly one site.



You cannot configure an Erasure Coding profile 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](#)
- Do not use the default storage pool, All Storage Nodes, or a storage pool that includes the default site, All Sites.

- The storage pool should include at least  $k+m + 1$  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.

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 overall more computationally efficient, as fewer fragments are created and distributed (or retrieved) per object, can show better performance due to the larger fragment size, and can require fewer nodes be added in an expansion when more storage is required. (See the instructions for expanding StorageGRID for information on planning a storage expansion.)

### 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%

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
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%

Erasure-coding scheme ( $k+m$ )	Minimum number of Storage Nodes	Recommended number of Storage Nodes	Storage overhead
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%

### Related information

[Expand your grid](#)

### 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 required. For example, if you use an erasure-coding scheme of 6+3, you must have at least three Storage Nodes at three different sites. In contrast, if you simply replicate object data, you require only one Storage Node for each copy.
- Increased cost and complexity of storage expansions. To expand a deployment that uses replication, you simply 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 the instructions for expanding StorageGRID.
- 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 larger than 1 MB in size.



Due to the overhead of managing the number of fragments associated with an erasure-coded copy, do not use erasure coding for objects 200 KB or smaller.

- 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.

## Related information

[Expand your grid](#)

## 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 have three primary ways 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 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 cannot be deleted by any method.
  - Before an object version's retain-until-date is reached, that version cannot 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.

#### Managing 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.
- An S3 or Swift 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 use 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, Cloud Storage Pools, or Archive Nodes).

### How S3 bucket lifecycle and ILM interact

The Expiration action in an S3 bucket lifecycle always overrides ILM settings. 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.

## Related information

[Managing objects with S3 Object Lock](#)

[Use S3](#)

[What ILM rule placement instructions are](#)

## 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	<p>When <b>any</b> of the following conditions are true:</p> <ul style="list-style-type: none"><li>• Automatic object deletion has been triggered by one of the following events:<ul style="list-style-type: none"><li>◦ The expiration date or number of days in the lifecycle configuration for an S3 bucket is reached.</li><li>◦ The last time period specified in an ILM rule elapses.</li></ul></li></ul> <p><b>Note:</b> Objects in a bucket that has S3 Object Lock enabled cannot be deleted if they are under a legal hold or if a retain-until-date has been specified but not yet met.</p> <ul style="list-style-type: none"><li>• An S3 or Swift client requests deletion and one or more of these conditions is true:<ul style="list-style-type: none"><li>◦ Copies cannot be deleted within 30 seconds because, for example, an object location is temporarily unavailable.</li><li>◦ Background deletion queues are idle.</li></ul></li></ul>

Method of performing deletion	When used
Objects are removed immediately (synchronous deletion)	<p>When an S3 or Swift client makes a delete request and <b>all</b> of the following conditions are met:</p> <ul style="list-style-type: none"> <li>• All copies can be removed within 30 seconds.</li> <li>• Background deletion queues contain objects to process.</li> </ul>

When S3 or Swift clients make delete requests, StorageGRID begins by adding a number of 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.

### Understanding the impact of how StorageGRID deletes 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 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.

You should be aware that the amount of time required to free space after an object is deleted depends on a number of 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 do not delete the current version of the object and do not free space. Instead, an object delete request simply creates a delete marker as the current version of the object, which makes the previous version of the object “noncurrent.”

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, you must do one of the following:

- **S3 client request:** Specify the object version number 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 cannot be recovered.
- **ILM:** Add two ILM rules to your ILM policy. Use **Noncurrent Time** as the Reference Time in the first rule to match the noncurrent versions of the object. Use **Ingest Time** in the second rule to match the current version. The **Noncurrent Time** rule must appear in the policy above the **Ingest Time** rule.

## Related information

[Use S3](#)

[Example 4: ILM rules and policy for S3 versioned objects](#)

## What an ILM policy is

An information lifecycle management (ILM) policy is an ordered set of ILM rules that determines how the StorageGRID system manages object data over time.

## How an ILM policy evaluates objects

The active ILM policy for your StorageGRID system controls the placement, duration, and data protection 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 do not 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 cannot use any filters.

## Example ILM policy

This example ILM policy uses three ILM rules.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### 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.

	Default	Rule Name	Tenant Account	Actions
		Rule 1: 3 replicated copies for Tenant A	Tenant A (58889986524346589742)	
		Rule 2: Erasure coding for objects greater than 1 MB	—	
	<input checked="" type="checkbox"/>	Rule 3: 2 copies 2 data centers (default)	—	

In this example, Rule 1 matches all objects belonging to Tenant A. These objects are stored as three replicated copies at three sites. Objects belonging to other tenants are not matched by Rule 1, so they are evaluated against Rule 2.

Rule 2 matches all objects from other tenants but only if they are larger than 1 MB. These larger objects are stored using 6+3 erasure coding at three sites. Rule 2 does not match objects 1 MB or smaller, so these objects are evaluated against Rule 3.

Rule 3 is the last and default rule in the policy, and it does not use filters. Rule 3 makes 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 proposed, active, and historical policies are

Every StorageGRID system must have one active ILM policy. A StorageGRID system might also have one proposed ILM policy and any number of historical policies.

When you first create an ILM policy, you create a proposed policy by selecting one or more ILM rules and arranging them in a specific order. After you have simulated the proposed policy to confirm its behavior, you activate it to create the active policy.

When you activate a new 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.

Activating the proposed policy causes the previously active policy to become a historical policy. Historical ILM policies cannot be deleted.



## Related information

[Creating an ILM policy](#)

## What an ILM rule is

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.

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, in Cloud Storage Pools, or on Archive Nodes.
- 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. The copies are evenly distributed across all Storage Nodes.

## 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 defines what happens when an S3 or Swift client saves an object to the grid. 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.

**Example ILM rule**

This example ILM rule applies to the objects belonging to Tenant A. It makes two replicated copies of those objects and stores each copy at a different site. The two copies are retained “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.

This rule uses 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.

**Two copies at two sites for Tenant A**

<b>Description:</b>	Applies only to Tenant A
<b>Ingest Behavior:</b>	Balanced
<b>Tenant Accounts:</b>	Tenant A (34176783492629515782)
<b>Reference Time:</b>	Ingest Time
<b>Filtering Criteria:</b>	Matches all objects.

**Retention Diagram:**

The diagram illustrates the retention policy for two sites. A vertical line marks 'Day 0' as the trigger point. Below this line, two horizontal bars represent the retention duration for Site 1 and Site 2. Site 1's bar is blue and Site 2's bar is orange. Both bars start at Day 0 and extend to the right, ending at a point labeled 'Forever'. Small cylinder icons are placed at the Day 0 trigger point for each site.

**Related information**

[Data-protection options for ingest](#)

[What a storage pool is](#)

[What a Cloud Storage Pool is](#)

[How objects are stored \(replication or erasure coding\)](#)

[What ILM rule filtering is](#)

## What ILM rule placement instructions are

### What ILM rule filtering is

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 do not match the filters in another rule.

Basic filters allow you to apply different rules to large, distinct groups of objects. The basic filters on the Define Basics page of the Create ILM Rule wizard allow you to apply a rule to specific tenant accounts, specific S3 buckets or Swift containers, or both.

Create ILM Rule Step 1 of 3: Define Basics

Name

Description

Tenant Accounts (optional)

Bucket Name

[Advanced filtering... \(0 defined\)](#)

These 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.

The **Advanced Filtering** page of the Create ILM Rule wizard gives 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)
- S3 bucket region (Location Constraint)
- Object size
- User metadata
- S3 object tags

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 also combine basic and advanced 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 advanced filters in each rule to identify the specific type of objects that the rule applies to.

### What ILM rule placement instructions are

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 a placement instruction, you specify when the placement applies (the time period), which type of copies to create (replicated or erasure coded), and where to store the copies (one or more storage locations). Within a single rule you can specify multiple placements for one time period, and placement instructions for more than one time period:

- To specify more than one object placement during a single time period, click the plus sign icon **+** to add more than one line for that time period.
- To specify object placements for more than one time period, click the **Add** button to add the next time period. Then, specify one or more lines within the time period.

The example shows the Define Placements page of the Create ILM Rule wizard.

Placements Sort by start day

From day 0 store for 365 days Add Remove

Type replicated Location DC1 x DC2 x Add Pool Copies 2 + x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

Type erasure coded Location All 3 sites (6 plus 3) Copies 1 1 + x

From day 365 store forever Add Remove

Type replicated Location Archive x Add Pool Copies 2 Temporary location -- Optional -- 2 + x

<b>1</b>	The first placement instruction has two lines for the first year: <ol style="list-style-type: none"><li>1. The first line creates two replicated object copies at two data center sites.</li><li>2. The second line creates a 6+3 erasure-coded copy using three data center sites.</li></ol>
<b>2</b>	The second placement instruction creates two archived 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.

## Creating storage grades, storage pools, EC profiles, and regions

Before you can create the ILM rules for your StorageGRID system, you must define object storage locations, determine the types of copies you want, and optionally configure S3 regions.

- [Creating and assigning storage grades](#)
- [Configuring storage pools](#)
- [Using Cloud Storage Pools](#)
- [Configuring Erasure Coding profiles](#)
- [Configuring regions \(optional and S3 only\)](#)

### Creating and assigning 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, instead of on all nodes at the site. For example, you might want certain objects to be stored on your fastest Storage Nodes, such as StorageGRID all-flash storage appliances.

#### What you'll need

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

#### About this task

If you use more than one type of storage, you can optionally create a storage grade to identify each type. Creating storage grades allows you to select a specific type of Storage Node when configuring storage pools.

If storage grade is not a concern (for example, all Storage Nodes are identical), you can skip this procedure and use the All Storage Nodes default storage grade when configuring storage pools.

When you add a new Storage Node in an expansion, that node is added to the All Storage Nodes default storage grade. As a result:

- If an ILM rule uses a storage pool with the All Storage Nodes grade, the new node can be used immediately after the expansion completes.
- If an ILM rule uses a storage pool with a custom storage grade, the new node will not be used until you manually assign the custom storage grade to the node, as described below.



When creating storage grades, do not create more storage grades than necessary. For example, do not create one 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. Create a storage grade:
  - a. For each storage grade you need to define, click **Insert**  to add a row and enter a label for the storage grade.

The Default storage grade cannot be modified. It is reserved for new Storage Nodes added during a StorageGRID system expansion.



### Storage Grade Definitions

Storage Grade	Label	Actions
0	Default	
1	<input type="text" value="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 

- b. To edit an existing storage grade, click **Edit**  and modify the label as required.



You cannot delete storage grades.

- c. Click **Apply Changes**.

These storage grades are now available for assignment to Storage Nodes.

3. Assign a storage grade to a Storage Node:
  - a. For each Storage Node's LDR service, click **Edit**  and select a storage grade from the list.

## Storage Grades



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. Do not change this assignment after the ILM policy is activated. If the assignment is changed, data is stored based on the new storage grade.

b. Click **Apply Changes**.

### Configuring storage pools

When defining an ILM rule, you use storage pools to specify where objects are stored . Before creating a storage pool, you must review the storage pool guidelines.

- [What a storage pool is](#)
- [Guidelines for creating storage pools](#)
- [Using multiple storage pools for cross-site replication](#)
- [Using a storage pool as a temporary location \(deprecated\)](#)
- [Creating a storage pool](#)
- [Viewing storage pool details](#)
- [Editing a storage pool](#)
- [Removing a storage pool](#)

### What a storage pool is

A storage pool is a logical grouping of Storage Nodes or Archive Nodes. You configure storage pools to determine where the StorageGRID system stores object data and the type of storage used.

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. When you configure ILM rules for replication, you select one or more storage pools that include either Storage Nodes or Archive Nodes. When you create Erasure Coding profiles, you select a storage pool that includes Storage Nodes.

## Guidelines for creating storage pools

When configuring and using storage pools, follow these guidelines.

### Guidelines for all storage pools

- StorageGRID includes a default storage pool, All Storage Nodes, that uses the default site, All Sites, and the default storage grade, All Storage Nodes. The All Storage Nodes storage pool is automatically updated whenever you add new data center sites.



Using the All Storage Nodes storage pool or the All Sites site is not recommended because these items are automatically updated to include any new sites you add in an expansion, which might not be the behavior you want. Before using the All Storage Nodes storage pool or the default site, carefully review the guidelines for replicated and erasure-coded copies.

- Keep storage pool configurations as simple as possible. Do not 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.

### Guidelines for storage pools used for replicated copies

- Create a different storage pool for each site. Then, specify one or more site-specific storage pools in the placement instructions for each rule. 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 for the new site. Then, update ILM rules to control which objects are stored on the new site.
- In general, do not use the default storage pool, All Storage Nodes, or any storage pool that includes the default site, All Sites.

### Guidelines for storage pools used for erasure-coded copies

- You cannot use Archive Nodes for erasure-coded data.
- The number of Storage Nodes and sites contained in the storage pool determine which erasure-coding schemes are available.
- If a storage pool includes only two sites, you cannot use that storage pool for erasure coding. No erasure-coding schemes are available for a storage pool that has two sites.
- In general, do not use the default storage pool, All Storage Nodes, or any storage pool that includes the default site, All Sites in any Erasure Coding profile.



If your grid includes only one site, you are prevented from using the All Storage Nodes storage pool or the All Sites default site in an Erasure Coding profile. This behavior prevents the Erasure Coding profile from becoming invalid if a second site is added.

- If you have high throughput requirements, creating 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 Strict or Balanced are selected as the Ingest Behavior) or could lead to ILM queue backlogs (when Dual Commit is selected as the Ingest Behavior).
- 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.

### **Guidelines for storage pools used for archived copies**

- You cannot create a storage pool that includes both Storage Nodes and Archive Nodes. Archived copies require a storage pool that only includes Archive Nodes.
- When using a storage pool that includes Archive Nodes, you should also maintain at least one replicated or erasure-coded copy on a storage pool that includes Storage Nodes.
- If the global S3 Object Lock setting is enabled and you are creating a compliant ILM rule, you cannot use a storage pool that includes Archive Nodes. See the instructions for managing objects with S3 Object Lock.
- If an Archive Node's Target Type is Cloud Tiering - Simple Storage Service (S3), the Archive Node must be in its own storage pool. See the instructions for administering StorageGRID.

### **Related information**

[What replication is](#)

[What erasure coding is](#)

[What erasure-coding schemes are](#)

[Using multiple storage pools for cross-site replication](#)

[Using a storage pool as a temporary location \(deprecated\)](#)

[Managing objects with S3 Object Lock](#)


[Administer StorageGRID](#)

### **Using multiple storage pools for cross-site replication**


If your StorageGRID deployment includes more than one site, you can enable site-loss protection by creating a storage pool for each site and specifying both storage pools in the rule's placement instructions. For example, if you configure an ILM rule to make two replicated copies and specify storage pools at two sites, one copy of each object will be placed at each site. If you configure a rule to make two copies and specify three storage

pools, the copies are distributed to balance disk usage among the storage pools, while ensuring that the two copies are stored at different sites.

The following example illustrates what can happen if an ILM rule places replicated object copies to a single storage pool containing Storage Nodes from two sites. Because the system uses any available nodes in the storage pool when it places the replicated copies, it might place all copies of some objects within only one of the sites. In this example, the system stored two copies of object AAA on Storage Nodes at Site 1, and two copies of object CCC on Storage Nodes at Site 2. Only object BBB is protected if one of the sites fails or becomes inaccessible.



In contrast, this example illustrates how objects are stored when you use multiple storage pools. In the example, the ILM rule specifies that two replicated copies of each object be created, and that the copies be distributed to two storage pools. Each storage pool contains all Storage Nodes at one site. Because a copy of each object is stored at each site, object data is protected from site failure or inaccessibility.



When using multiple storage pools, keep the following rules in mind:

- If you are creating n copies, you must add n or more storage pools. For example, if a rule is configured to make three copies, you must specify three or more 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, the system distributes the copies to keep disk usage among the pools balanced and to ensure that two or more copies are not stored in the same storage pool.
- 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 do not contain the same Storage Nodes.

### Using a storage pool as a temporary location (deprecated)

When you create an ILM rule with an object placement that includes a single storage pool, you are prompted to specify a second storage pool to use as a temporary location.

Temporary locations have been deprecated and will be removed in a future release. You should not select a storage pool as a temporary location for a new ILM rule.



If you select the Strict ingest behavior (Step 3 of the Create ILM Rule wizard), the temporary location is ignored.

### Related information

[Data-protection options for ingest](#)

### Creating 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have 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. Select the plus icon **+** to add an entry for each site.



Do not include the default 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, do not create a storage pool that includes different storage grades at a single site.

## Guidelines for creating storage pools

### Steps

1. Select **ILM > Storage Pools**.

The Storage Pools page appears and lists all defined storage pools.

Storage Pools

#### Storage Pools

A storage pool is a logical group of Storage Nodes or Archive Nodes and is used in ILM rules to determine where object data is stored.

Create	Edit	Remove	View Details				
Name	Used Space	Free Space	Total Capacity	ILM Usage			
All Storage Nodes	1.10 MB	102.90 TB	102.90 TB	Used in 1 ILM rule			

Displaying 1 storage pool.

#### Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

Create	Edit	Remove	Clear Error		
--------	------	--------	-------------	--	--

No Cloud Storage Pools found.

The list includes the system-default storage pool, All Storage Nodes, which uses the system-default site, All Sites, and the default storage grade, All Storage Nodes.



Because the All Storage Nodes storage pool is automatically updated whenever you add new data center sites, using this storage pool in ILM rules is not recommended.

2. To create a new storage pool, select **Create**.

The Create Storage Pool dialog box appears.

## Create Storage Pool

- For replication and single-site erasure coding, create a storage pool for each site.
- For erasure coding at three or more sites, click + to add each site to a single storage pool.
- Do not add more than one storage grade for a single site.

Name

Site  Storage Grade

Viewing Storage Pool -		
Site Name	Archive Nodes	Storage Nodes

Cancel

Save

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 and Archive Nodes in the table are automatically updated.

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 default All Storage Nodes storage grade includes all Storage Nodes at the selected site. The default Archive Nodes storage grade includes all Archive 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 **+** to add an entry for each site to the storage pool.

## Create Storage Pool

- For replication and single-site erasure coding, create a storage pool for each site.
- For erasure coding at three or more sites, select + to add each site to a single storage pool.
- Do not select more than one storage grade for a single site.

Name:

Site:  Storage Grade:

Site:  Storage Grade:

Site:  Storage Grade:

### Viewing Storage Pool - All 3 Sites for Erasure Coding

Site Name	Archive Nodes	Storage Nodes
Data Center 1	0	3
Data Center 2	0	3
Data Center 3	0	3

You are creating a multi-site storage pool, which should not be used for replication or single-site erasure coding.

Cancel

Save



You are prevented from creating duplicate entries or from creating a storage pool that includes both the **Archive Nodes** storage grade and any storage grade that contains Storage Nodes.

You are warned if you add more than one entry for a site but with different storage grades.

To remove an entry, select ✘.

7. When you are satisfied with your selections, select **Save**.

The new storage pool is added to the list.

### Related information

[Guidelines for creating storage pools](#)

### Viewing 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.

### What you'll need

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

## Steps

### 1. Select **ILM > Storage Pools**.

The Storage Pools page appears. This page lists all defined storage pools.

Storage Pools

#### Storage Pools

A storage pool is a logical group of Storage Nodes or Archive Nodes and is used in ILM rules to determine where object data is stored.

+ Create Edit Remove View Details					
Name	Used Space	Free Space	Total Capacity	ILM Usage	
<input checked="" type="radio"/> All Storage Nodes	1.88 MB	2.80 TB	2.80 TB	Used in 1 ILM rule	
<input type="radio"/> DC1	621.77 KB	932.42 GB	932.42 GB	Used in 2 ILM rules	
<input type="radio"/> DC2	675.82 KB	932.42 GB	932.42 GB	Used in 2 ILM rules	
<input type="radio"/> DC3	578.95 KB	932.42 GB	932.42 GB	Used in 1 ILM rule	
<input type="radio"/> All 3 Sites	1.88 MB	2.80 TB	2.80 TB	Used in 1 ILM rule and 1 EC profile	
<input type="radio"/> Archive	—	—	—	—	

Displaying 6 storage pools.

#### Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

+ Create Edit Remove Clear Error			
No Cloud Storage Pools found.			

The table includes the following information for each storage pool that includes Storage Nodes:

- **Name:** The unique display name of the storage pool.
- **Used Space:** The amount of space that is currently being used to store objects in the storage pool.
- **Free Space:** The amount of space that remains available to store objects in the storage pool.
- **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.



You cannot remove a storage pool if it is being used.

### 2. To view details about a specific storage pool, select its radio button and select **View Details**.

The Storage Pool Details modal appears.

### 3. View the **Nodes Included** tab to learn about the Storage Nodes or Archive Nodes included in the storage pool.

## Storage Pool Details - DC1

Nodes Included

ILM Usage

Number of Nodes: 3

Storage Grade: All Storage Nodes

Node Name	Site Name	Used (%)
DC1-S1	Data Center 1	0.000%
DC1-S2	Data Center 1	0.000%
DC1-S3	Data Center 1	0.000%

Close

The table includes the following information for each node:

- Node Name
- Site Name
- Used (%): For Storage Nodes, the percentage of the total usable space for object data that has been used. This value does not include object metadata.



The same Used (%) value is also shown in the Storage Used - Object Data chart for each Storage Node (select **Nodes** > **Storage Node** > **Storage**).

4. Select the **ILM Usage** tab to determine if the storage pool is currently being used in any ILM rules or Erasure Coding profiles.

In this example, the DC1 storage pool is used in three ILM rules: two rules that are in the active ILM policy and one rule that is not in the active policy.

## Storage Pool Details - DC1

Nodes Included

ILM Usage

### ILM Rules Using the Storage Pool

The following ILM rules in the active ILM policy (Example ILM policy) use this storage pool.

- 3 copies for Account01
- 2 copies for smaller objects

1 ILM rule that is not in the active ILM policy uses this storage pool.

If you want to remove this storage pool, you must delete or edit every rule where it is used. Go to the [ILM Rules page](#).

### EC Profiles Using the Storage Pool

No Erasure Coding profiles use this storage pool.

Close



You cannot remove a storage pool if it is used in an ILM rule.

In this example, the All 3 Sites storage pool is used in an Erasure Coding profile. In turn, that Erasure

Coding profile is used by one ILM rule in the active ILM policy.

Storage Pool Details - All 3 Sites

Nodes Included | ILM Usage

---

### ILM Rules Using the Storage Pool

The following ILM rules in the active ILM policy (Example ILM policy) use this storage pool.

- EC larger objects

If you want to remove this storage pool, you must delete or edit every rule where it is used. Go to the [ILM Rules page](#)

---

### EC Profiles Using the Storage Pool

The following Erasure Coding profiles use this storage pool.

Profile Name	Profile Status
6 plus 3	Used in 1 ILM Rule

[Close](#)



You cannot remove a storage pool if it is used in an Erasure Coding profile.

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.

6. When you are done viewing storage pool details, select **Close**.

## Related information

[Working with ILM rules and ILM policies](#)

## Editing a storage pool

You can edit a storage pool to change its name or to update sites and storage grades.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have 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 must have considered how your changes will affect object data placement.

### About this task

If you are adding a new storage grade to a storage pool that is used in the active ILM policy, be aware that the Storage Nodes in the new storage grade will not be used automatically. To force StorageGRID to use a new storage grade, you must activate a new ILM policy after saving the edited storage pool.

### Steps

1. Select **ILM > Storage Pools**.

The Storage Pools page appears.

2. Select the radio button for the storage pool you want to edit.

You cannot edit the All Storage Nodes storage pool.

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 since the change would make the erasure-coding scheme invalid.

6. Select **Save**.

### After you finish

If you added a new storage grade to a storage pool used in the active ILM policy, activate a new ILM policy to force StorageGRID to use the new storage grade. For example, clone your existing ILM policy and then activate the clone.

### Removing a storage pool

You can remove a storage pool that is not being used.

#### What you'll need

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

#### Steps

1. Select **ILM > Storage Pools**.

The Storage Pools page appears.

2. Look at the ILM Usage column in the table to determine whether you can remove the storage pool.

You cannot remove a storage pool if it is being used in an ILM rule or in an Erasure Coding profile. As required, select **View Details > ILM Usage** to determine where a storage pool is used.

3. If the storage pool you want to remove is not being used, select the radio button.
4. Select **Remove**.
5. Select **OK**.

### Using Cloud Storage Pools

You can use Cloud Storage Pools to move StorageGRID objects to an external storage location, such as S3 Glacier or Microsoft Azure Blob storage. Moving objects outside of the grid lets you take advantage of a low-cost storage tier for long-term archive.

- [What a Cloud Storage Pool is](#)
- [Lifecycle of a Cloud Storage Pool object](#)
- [When to use Cloud Storage Pools](#)
- [Considerations for Cloud Storage Pools](#)
- [Comparing Cloud Storage Pools and CloudMirror replication](#)
- [Creating a Cloud Storage Pool](#)
- [Editing a Cloud Storage Pool](#)
- [Removing a Cloud Storage Pool](#)
- [Troubleshooting Cloud Storage Pools](#)

## What a Cloud Storage Pool is

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, 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 or Archive Nodes within the StorageGRID system, a Cloud Storage Pool consists of an external bucket (S3) or container (Azure Blob storage).

The following table compares storage pools to Cloud Storage Pools and shows the high-level similarities and differences.

	<b>Storage pool</b>	<b>Cloud Storage Pool</b>
How is it created?	Using the <b>ILM &gt; Storage Pools</b> option in Grid Manager.  You must set up storage grades before you can create the storage pool.	Using the <b>ILM &gt; Storage Pools</b> 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 or Archive Nodes within StorageGRID.	<p>In an Amazon S3 bucket or Azure Blob storage container that is external to the StorageGRID system.</p> <p>If the Cloud Storage Pool is an Amazon S3 bucket:</p> <ul style="list-style-type: none"> <li>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 POST Object restore API.</li> <li>You can create Cloud Storage Pools for use with AWS Commercial Cloud Services (C2S), which supports the AWS Secret Region.</li> </ul> <p>If the Cloud Storage Pool is an Azure Blob storage container, StorageGRID transitions the object to the Archive tier.</p> <p><b>Note:</b> In general, do not configure Azure Blob Storage lifecycle management for the container used for a Cloud Storage Pool. POST Object restore operations on objects in the Cloud Storage Pool can be affected by the configured lifecycle.</p>
What controls object placement?	An ILM rule in the active ILM policy.	An ILM rule in the active ILM policy.
What data protection method is used?	Replication or erasure coding.	Replication.
How many copies of each object are allowed?	Multiple.	<p>One copy in the Cloud Storage Pool and, optionally, one or more copies in StorageGRID.</p> <p><b>Note:</b> You cannot store an object in more than one Cloud Storage Pool at any given time.</p>
What are the advantages?	Objects are quickly accessible at any time.	Low-cost storage.

## 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.

### Related information

[S3: Lifecycle of a Cloud Storage Pool object](#)


[Azure: Lifecycle of a Cloud Storage Pool object](#)

### S3: Lifecycle of a Cloud Storage Pool object

The figure shows the lifecycle stages of an object that is stored in an S3 Cloud Storage Pool.



In the figure and explanations, “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.



#### 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 GET Object 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 POST Object restore API.



Do not use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support POST Object restore requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to S3 Glacier storage. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).

- During the transition, the client application can use an S3 HEAD Object 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 POST Object restore 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 POST Object restore request. Instead, the local copy can be retrieved directly, using a GET Object request.

## 5. Object retrieved


Once an object has been restored, the client application can issue a GET Object request to retrieve the restored object.

### Related information

[Use S3](#)

### Azure: Lifecycle of a Cloud Storage Pool object

The figure shows 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



Do not use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support POST Object restore requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to the Azure Blob storage Archive tier. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).

### 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 POST Object

restore request to restore a retrievable copy to the Azure Cloud Storage Pool.

When StorageGRID receives the POST Object Restore, it temporarily transitions the object to the Azure Blob storage Cool tier. As soon as the expiration date in the POST Object restore 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 POST Object restore request. Instead, the local copy can be retrieved directly, using a GET Object request.

## 5. Object retrieved

Once an object has been restored to the Azure Cloud Storage Pool, the client application can issue a GET Object request to retrieve the restored object.

### When to use Cloud Storage Pools

Cloud Storage Pools can provide significant benefits in several use cases.

#### Backing up StorageGRID data in an 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 POST Object restore 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.

#### Tiering 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 Pools 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 tenant 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 Pools, 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.
- Do not use Cloud Storage Pools for objects that have been ingested by Swift clients. Swift does not support POST Object restore requests, so StorageGRID will not be able to retrieve any Swift objects that have been transitioned to S3 Glacier storage or the Azure Blob storage Archive tier. Issuing a Swift GET object request to retrieve these objects will fail (403 Forbidden).
- Using Cloud Storage Pools with FabricPool is not supported because of the added latency to retrieve an object from the Cloud Storage Pool target.

## Information required to create a Cloud Storage Pool

Before you can create a Cloud Storage Pool, you must create the external S3 bucket or the external Azure Blob storage container that you will use for the Cloud Storage Pool. Then, when you create the Cloud Storage Pool in StorageGRID, you must specify the following information:

- The provider type: Amazon S3 or Azure Blob storage.
- If you select Amazon S3, whether the Cloud Storage Pool is for use with the AWS Secret Region (**CAP (C2S Access Portal)**).
- The exact name of the bucket or container.
- The service endpoint needed to access the bucket or container.
- The authentication needed to access the bucket or container:
  - **S3**: Optionally, an access key ID and secret access key.

- **C2S:** The complete URL for obtaining temporary credentials from the CAP server; a server CA certificate, a client certificate, a private key for the client certificate, and, if the private key is encrypted, the passphrase for decrypting it.
- **Azure Blob storage:** An account name and account key. These credentials must have full permission for the container.
- Optionally, a custom CA certificate to verify TLS connections to the bucket or container.

### 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 simply 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 bucket policy 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 policy 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 POST Object restore 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 cannot 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

Do not 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

[Creating a Cloud Storage Pool](#)

[S3: Specifying authentication details for a Cloud Storage Pool](#)

[C2S S3: Specifying authentication details for a Cloud Storage Pool](#)

[Azure: Specifying authentication details for a Cloud Storage Pool](#)

[Administer StorageGRID](#)

### Comparing 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?	A Cloud Storage Pool 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 on-premise, you can keep only one copy within StorageGRID and send a copy to the Cloud Storage Pool.	The CloudMirror replication service 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.

	<b>Cloud Storage Pool</b>	<b>CloudMirror replication service</b>
How is it set up?	Cloud Storage Pools are defined in the same way as storage pools, using the Grid Manager or the Grid Management API. A Cloud Storage Pool 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?	<ul style="list-style-type: none"> <li>• Any compatible S3 infrastructure (including Amazon S3)</li> <li>• Azure Blob Archive tier</li> </ul>	<ul style="list-style-type: none"> <li>• Any compatible S3 infrastructure (including Amazon S3)</li> </ul>
What causes objects to be moved to the destination?	One or more ILM rules in the active ILM policy. 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 are not 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.
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 in 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.

	<b>Cloud Storage Pool</b>	<b>CloudMirror replication service</b>
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.

**Related information**

[Administer StorageGRID](#)

**Creating a Cloud Storage Pool**

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 or Azure Blob Storage), and the information StorageGRID needs to access the external bucket or container.

**What you'll need**

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have reviewed the guidelines for configuring Cloud Storage Pools.
- The external bucket or container referenced by the Cloud Storage Pool must exist.
- You must have all of the authentication information needed to access the bucket or container.

**About this task**

A Cloud Storage Pool specifies a single external S3 bucket or Azure Blob storage 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.

**Steps**

1. Select **ILM > Storage Pools**.

The Storage Pools page appears. This page includes two sections: Storage Pools and Cloud Storage Pools.

## Storage Pools

### Storage Pools

A storage pool is a logical group of Storage Nodes or Archive Nodes and is used in ILM rules to determine where object data is stored.

Name	Used Space	Free Space	Total Capacity	ILM Usage
All Storage Nodes	1.10 MB	102.90 TB	102.90 TB	Used in 1 ILM rule

Displaying 1 storage pool.

### Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

+ Create Edit Remove Clear Error

No Cloud Storage Pools found.

- In the Cloud Storage Pools section of the page, click **Create**.

The Create Cloud Storage Pool dialog box appears.

### Create Cloud Storage Pool

Display Name

Provider Type

Bucket or Container

Cancel Save

- Enter the following information:

Field	Description
Display 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: <ul style="list-style-type: none"><li>Amazon S3 (select this option for an S3 or C2S S3 Cloud Storage Pool)</li><li>Azure Blob Storage</li></ul> <p><b>Note:</b> When you select a Provider Type, the Service Endpoint, Authentication and Server Verification sections appear at the bottom on the page.</p>

Field	Description
Bucket or Container	The name of the external S3 bucket or Azure container that was created for the Cloud Storage Pool. The name you specify here must exactly match the bucket or container's name or Cloud Storage Pool creation will fail. You cannot change this value after the Cloud Storage Pool is saved.

4. Complete the Service Endpoint, Authentication and Server Verification sections of the page, based on the selected provider type.
  - [S3: Specifying authentication details for a Cloud Storage Pool](#)
  - [C2S S3: Specifying authentication details for a Cloud Storage Pool](#)
  - [Azure: Specifying authentication details for a Cloud Storage Pool](#)

### **S3: Specifying authentication details for a Cloud Storage Pool**

When you create a Cloud Storage Pool for S3, you must select the type of authentication that is required for the Cloud Storage Pool endpoint. You can specify Anonymous or enter an Access Key ID and Secret Access Key.

#### **What you'll need**

- You must have entered the basic information for the Cloud Storage Pool and specified **Amazon S3** as the provider type.

## Create Cloud Storage Pool

Display Name ⓘ S3 Cloud Storage Pool

Provider Type ⓘ Amazon S3 ▼

Bucket or Container ⓘ my-s3-bucket

### Service Endpoint

Protocol ⓘ  HTTP  HTTPS

Hostname ⓘ example.com or 0.0.0.0

Port (optional) ⓘ 443

### Authentication

Authentication Type ⓘ ▼

### Server Verification

Certificate Validation ⓘ Use operating system CA certificate ▼

Cancel

Save

- If you are using access key authentication, you must know the Access Key ID and Secret Access Key for the external S3 bucket.

### Steps

1. In the **Service Endpoint** section, provide the following information:
  - a. Select which protocol to use when connecting to the Cloud Storage Pool.

The default protocol is HTTPS.

- b. Enter the server hostname or IP address of the Cloud Storage Pool.

For example:

`s3-aws-region.amazonaws.com`



Do not include the bucket name in this field. You include the bucket name in the **Bucket or Container** field.

c. Optionally, specify the port that should be used when connecting to the Cloud Storage Pool.

Leave this field blank to use the default port: port 443 for HTTPS or port 80 for HTTP.

2. In the **Authentication** section, select the type of authentication that is required for the Cloud Storage Pool endpoint.

Option	Description
Access Key	An Access Key ID and Secret Access Key are required to access the Cloud Storage Pool bucket.
Anonymous	Everyone has access to the Cloud Storage Pool bucket. An Access Key ID and Secret Access Key are not required.
CAP (C2S Access Portal)	Used for C2S S3 only. Go to <a href="#">C2S S3: Specifying authentication details for a Cloud Storage Pool</a> .

3. If you selected Access Key, enter the following information:

Option	Description
Access Key ID	The Access Key ID for the account that owns the external bucket.
Secret Access Key	The associated Secret Access Key.

4. In the Server Verification section, select which method should be used to validate the certificate for TLS connections to the Cloud Storage Pool:

Option	Description
Use operating system CA certificate	Use the default CA certificates installed on the operating system to secure connections.
Use custom CA certificate	Use a custom CA certificate. Click <b>Select New</b> , and upload the PEM-encoded CA certificate.
Do not verify certificate	The certificate used for the TLS connection is not verified.

5. Click **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the bucket and the service endpoint exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the bucket to identify the bucket 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 you specified does not already exist.

## Error

### 422: Unprocessable Entity

Validation failed. Please check the values you entered for errors.

Cloud Pool test failed. Could not create or update Cloud Pool. Error from endpoint: NoSuchBucket: The specified bucket does not exist. status code: 404, request id: 4211567681, host id:

OK

See the instructions for troubleshooting Cloud Storage Pools, resolve the issue, and then try saving the Cloud Storage Pool again.

### Related information

[Troubleshooting Cloud Storage Pools](#)

### C2S S3: Specifying authentication details for a Cloud Storage Pool

To use the Commercial Cloud Services (C2S) S3 service as a Cloud Storage Pool, you must configure C2S Access Portal (CAP) as the authentication type, so that StorageGRID can request temporary credentials to access the S3 bucket in your C2S account.

### What you'll need

- You must have entered the basic information for an Amazon S3 Cloud Storage Pool, including the service endpoint.
- You must know 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.
- You must have a server CA certificate issued by an appropriate Government Certificate Authority (CA). StorageGRID uses this certificate to verify the identity of the CAP server. The server CA certificate must use PEM encoding.
- You must have a client certificate issued by an appropriate Government Certificate Authority (CA). StorageGRID uses this certificate to identify itself to the CAP server. The client certificate must use PEM encoding and must have been granted access to your C2S account.
- You must have a PEM-encoded private key for the client certificate.
- If the private key for the client certificate is encrypted, you must have the passphrase for decrypting it.

### Steps

1. In the **Authentication** section, select **CAP (C2S Access Portal)** from the **Authentication Type** drop-down.

The CAP C2S authentication fields appear.

## Create Cloud Storage Pool

Display Name ⓘ

Provider Type ⓘ

Bucket or Container ⓘ

### Service Endpoint

Protocol ⓘ  HTTP  HTTPS

Hostname ⓘ

Port (optional) ⓘ

### Authentication

Authentication Type ⓘ

Temporary Credentials URL ⓘ

Server CA Certificate ⓘ

Client Certificate ⓘ

Client Private Key ⓘ

Client Private Key Passphrase (optional) ⓘ

### Server Verification

Certificate Validation ⓘ

Cancel

Save

2. Provide the following information:

- a. For **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. For **Server CA Certificate**, click **Select New**, and upload the PEM-encoded CA certificate that StorageGRID will use to verify the CAP server.
- c. For **Client Certificate**, click **Select New**, and upload the PEM-encoded certificate that StorageGRID will use to identify itself to the CAP server.
- d. For **Client Private Key**, click **Select New**, and upload the PEM-encoded private key for the client certificate.

If the private key is encrypted, the traditional format must be used. (PKCS #8 encrypted format is not supported.)

- 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.

3. In the Server Verification section, provide the following information:

- a. For **Certificate Validation**, select **Use custom CA certificate**.
- b. Click **Select New**, and upload the PEM-encoded CA certificate.

4. Click **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the bucket and the service endpoint exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the bucket to identify the bucket 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 you specified does not already exist.

## Error

### 422: Unprocessable Entity

Validation failed. Please check the values you entered for errors.

Cloud Pool test failed. Could not create or update Cloud Pool. Error from endpoint: NoSuchBucket: The specified bucket does not exist. status code: 404, request id: 4211567681, host id:

OK

See the instructions for troubleshooting Cloud Storage Pools, resolve the issue, and then try saving the Cloud Storage Pool again.

### Related information

### Azure: Specifying authentication details for a Cloud Storage Pool

When you create a Cloud Storage Pool for Azure Blob storage, you must specify an account name and account key for the external container that StorageGRID will use to store objects.

#### What you'll need

- You must have entered the basic information for the Cloud Storage Pool and specified **Azure Blob Storage** as the provider type. **Shared Key** appears in the **Authentication Type** field.

### Create Cloud Storage Pool

Display Name ⓘ Azure Cloud Storage Pool

Provider Type ⓘ Azure Blob Storage ▼

Bucket or Container ⓘ my-azure-container

---

#### Service Endpoint

URI ⓘ https://myaccount.blob.core.windows.net

---

#### Authentication

Authentication Type ⓘ Shared Key

Account Name ⓘ

Account Key ⓘ

---

#### Server Verification

Certificate Validation ⓘ Use operating system CA certificate ▼

Cancel Save

- You must know the Uniform Resource Identifier (URI) used to access the Blob storage container used for the Cloud Storage Pool.

- You must know the name of the storage account and the secret key. You can use the Azure portal to find these values.

## Steps

1. In the **Service Endpoint** section, enter the Uniform Resource Identifier (URI) used to access the Blob storage container used for the Cloud Storage Pool.

Specify the URI in one of the following formats:

- `https://host:port`
- `http://host:port`

If you do not specify a port, by default port 443 is used for HTTPS URIs and port 80 is used for HTTP URIs.

### Example URI for Azure Blob storage container:

`https://myaccount.blob.core.windows.net`

2. In the **Authentication** section, provide the following information:
  - a. For **Account Name**, enter the name of the Blob storage account that owns the external service container.
  - b. For **Account Key**, enter the secret key for the Blob storage account.



For Azure endpoints, you must use Shared Key authentication.

3. In the **Server Verification** section, select which method should be used to validate the certificate for TLS connections to the Cloud Storage Pool:

Option	Description
Use operating system CA certificate	Use the default CA certificates installed on the operating system to secure connections.
Use custom CA certificate	Use a custom CA certificate. Click <b>Select New</b> , and upload the PEM-encoded certificate.
Do not verify certificate	The certificate used for the TLS connection is not verified.

4. Click **Save**.

When you save a Cloud Storage Pool, StorageGRID does the following:

- Validates that the container and the URI exist and that they can be reached using the credentials that you specified.
- Writes a marker file to the 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 container you specified does not already exist.

See the instructions for troubleshooting Cloud Storage Pools, resolve the issue, and then try saving the Cloud Storage Pool again.

## Related information

[Troubleshooting Cloud Storage Pools](#)

## Editing a Cloud Storage Pool

You can edit a Cloud Storage Pool to change its name, service endpoint, or other details; however, you cannot change the S3 bucket or Azure container for a Cloud Storage Pool.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have reviewed the guidelines for configuring Cloud Storage Pools.

### Steps

1. Select **ILM > Storage Pools**.

The Storage Pools page appears. The Cloud Storage Pools table lists the existing Cloud Storage Pools.

#### Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

	Pool Name	URI	Pool Type	Container	Used in ILM Rule	Last Error
<input checked="" type="radio"/>	azure-endpoint	https://storagegrid.blob.core.windows.net	azure	azure-3	✓	
<input type="radio"/>	s3-endpoint	https://s3.amazonaws.com	s3	s3-1	✓	

Displaying 2 pools.

2. Select the radio button for the Cloud Storage Pool you want to edit.
3. Click **Edit**.
4. As required, change the display name, service endpoint, authentication credentials, or certificate validation method.



You cannot 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 select **View Current** to review the certificate that is currently in use.

5. Click **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.

## Related information

[Considerations for Cloud Storage Pools](#)

[Troubleshooting Cloud Storage Pools](#)

## Removing a Cloud Storage Pool

You can remove a Cloud Storage Pool that is not used in an ILM rule and that does not contain object data.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- 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 “Troubleshooting 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. Do not remove this file, which is named `x-ntap-sgws-cloud-pool-uuid`.

- You have already removed any ILM rules that might have used the pool.

### Steps

1. Select **ILM > Storage Pools**.

The Storage Pools page appears.

2. Select the radio button for a Cloud Storage Pool that is not currently used in an ILM rule.

You cannot remove a Cloud Storage Pool if it is used in an ILM rule. The **Remove** button is disabled.

#### Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

Pool Name	URI	Pool Type	Container	Used in ILM Rule	Last Error
<input checked="" type="radio"/> azure-endpoint	https://storagegrid.blob.core.windows.net	azure	azure-3	✓	
<input type="radio"/> s3-endpoint	https://s3.amazonaws.com	s3	s3-1	✓	

Displaying 2 pools.

3. Click **Remove**.

A confirmation warning is displayed.

## Warning

### Remove Cloud Storage Pool

Are you sure you want to remove this Cloud Storage Pool: My Cloud Storage Pool?

Cancel

OK

4. Click **OK**.

The Cloud Storage Pool is removed.

## Related information

[Troubleshooting Cloud Storage Pools](#)

## Troubleshooting Cloud Storage Pools

If you encounter errors when creating, editing, or deleting a Cloud Storage Pool, use these troubleshooting steps to help resolve the issue.

### Determining if an error has occurred

StorageGRID performs a simple health check on every Cloud Storage Pool once a minute to ensure that the Cloud Storage Pool can be accessed and that it is functioning correctly. If the 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.

#### Cloud Storage Pools

You can add Cloud Storage Pools to ILM rules to store objects outside of the StorageGRID system. A Cloud Storage Pool defines how to access the external bucket or container where objects will be stored.

Pool Name	URI	Pool Type	Container	Used in ILM Rule	Last Error
<input checked="" type="radio"/> S3	10.96.106.142:18082	s3	s3	✓	Endpoint failure: DC2-S1-106-147: Could not create or update Cloud Storage Pool. Error from endpoint: RequestError: send request failed caused by: Get https://10.96.106.142:18082/s3-targetbucket/x-ntap-sgws-cloud-pool-uuid: net/http: request canceled while waiting for connection (Client.Timeout exceeded while awaiting headers) 8 minutes ago
<input type="radio"/> Azure	http://pboerkoe@10.96.100.254:10000/d-evstoreaccount1	azure	azure	✓	

Displaying 2 pools.

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 Pool page (select **ILM > Storage Pools**), review the error messages in the Last Error column, and refer to the troubleshooting guidelines below.

### Checking 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 radio button for the endpoint, and click **Clear Error**. A confirmation message

indicates that StorageGRID has cleared the error for the Cloud Storage Pool.

Error successfully cleared. This error might reappear if the underlying problem is not resolved.



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: 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 does not have 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 these steps to correct the issue:

- Check to make sure that no one in your organization is also using this Cloud Storage Pool.
- Delete the `x-ntap-sgws-cloud-pool-uuid` file and try configuring the Cloud Storage Pool again.

### **Error: Could not create or update Cloud Storage Pool. Error from endpoint**

You might encounter this error when you try to create or edit a Cloud Storage Pool. This error indicates that some kind of 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 the HTTP protocol 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.
- 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 do not 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 cannot 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 cannot 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.

## Related information

[Lifecycle of a Cloud Storage Pool object](#)

## Configuring Erasure Coding profiles

You configure Erasure Coding profiles by associating a storage pool with an erasure-coding scheme, such as 6+3. Then, when you configure the placement instructions for an ILM rule, you can select the Erasure Coding profile. If an object matches the rule, data and parity fragments are created and distributed to the storage locations in the storage pool according to the erasure-coding scheme.

- [Creating an Erasure Coding profile](#)
- [Renaming an Erasure Coding profile](#)
- [Deactivating an Erasure Coding profile](#)

## Creating an Erasure Coding profile

To create an Erasure Coding profile, you associate a storage pool containing Storage Nodes with an erasure-coding scheme. This association determines the number of data and parity fragments created and where the system distributes these fragments.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have created a storage pool that includes exactly one site or a storage pool that includes three or more sites. No erasure-coding schemes are available for a storage pool that has only two sites.

### About this task

The storage pools used in Erasure Coding profiles must include exactly one site or three or more sites. If you want to provide site redundancy, the storage pool must have at least three sites.



You must select a storage pool that contains Storage Nodes. You cannot use Archive Nodes for erasure-coded data.

### Steps

1. Select **ILM > Erasure Coding**.

The Erasure Coding Profiles page appears.

## Erasure Coding Profiles ?

An Erasure Coding profile determines how many data and parity fragments are created and where those fragments are stored.

To create an Erasure Coding profile, select a [storage pool](#) and an erasure coding scheme. The storage pool must include Storage Nodes from exactly one site or from three or more sites. If you want to provide site redundancy, the storage pool must include nodes from at least three sites.

To deactivate an Erasure Coding profile that you no longer plan to use, first remove it from all ILM rules. Then, if the profile is still associated with object data, wait for those objects to be moved to new locations based on the new rules in the active ILM policy. Depending on the number of objects and the size of your StorageGRID system, it might take weeks or even months for the objects to be moved.

See [Managing objects with information lifecycle management](#) for important details.

<a href="#">+ Create</a>	<a href="#">Rename</a>	<a href="#">Deactivate</a>	Profile	Status	Storage Pool	Storage Nodes	Sites	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
No Erasure Coding profiles found.											

### 2. Click **Create**.

The Create EC Profile dialog box appears.

### Create EC Profile

You cannot change the selected scheme and storage pool after saving the profile.

Profile Name ?

Storage Pool ?

[Cancel](#) [Save](#)

### 3. Enter a unique name for the Erasure Coding profile.

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.



The Erasure Coding profile name is appended to the storage pool name in the placement instruction for an ILM rule.

From day  store

Type  Location  Copies

**Erasure Coding profile name** [Add](#) [Remove](#)

**Storage pool name**

### 4. Select the storage pool you created for this Erasure Coding profile.



If your grid currently includes only one site, you are prevented from using the default storage pool, All Storage Nodes, or any storage pool that includes the default site, All Sites. This behavior prevents the Erasure Coding profile from becoming invalid if a second site is added.



If a storage pool includes exactly two sites, you cannot use that storage pool for erasure coding. No erasure-coding schemes are available for a storage pool that has two sites.

When you select a storage pool, the list of available erasure-coding schemes is shown, based on the number of Storage Nodes and sites in the pool.

Create EC Profile

You cannot change the selected scheme and storage pool after saving the profile.

Profile Name ?

Storage Pool ?

9 Storage Nodes across 3 site(s)

Scheme

	Erasure Code <span>?</span>	Storage Overhead (%) <span>?</span>	Storage Node Redundancy <span>?</span>	Site Redundancy <span>?</span>
<input checked="" type="radio"/>	6+3	50%	3	Yes
<input type="radio"/>	2+1	50%	1	Yes
<input type="radio"/>	4+2	50%	2	Yes

The following information is listed for each available erasure-coding scheme:

- **Erasure Code:** The name of the erasure-coding scheme in the following format: data fragments + parity fragments.
- **Storage Overhead (%):** The additional storage required for parity fragments relative to the object's data size. Storage Overhead = Total number of parity fragments / Total number of data fragments.
- **Storage Node Redundancy:** The number of Storage Nodes that can be lost while still maintaining the ability to retrieve object data.
- **Site Redundancy:** Whether the selected erasure code allows the object data to be retrieved if a site is lost.

To support site redundancy, the selected storage pool must include multiple sites, each with enough Storage Nodes to allow any site to be lost. For example, to support site redundancy using a 6+3 erasure-coding scheme, the selected storage pool must include at least three sites with at least three Storage Nodes at each site.

Messages are displayed in these cases:

- The storage pool you selected does not provide site redundancy. The following message is expected when the selected storage pool includes only one site. You can use this Erasure Coding profile in ILM rules to protect against node failures.

Scheme

	Erasure Code <span>?</span>	Storage Overhead (%) <span>?</span>	Storage Node Redundancy <span>?</span>	Site Redundancy <span>?</span>
<input checked="" type="radio"/>	2+1	50%	1	No

The selected storage pool and erasure coding scheme cannot protect object data from loss if a site is lost.  
To provide site redundancy, the storage pool must have at least three sites.

- The storage pool you selected does not satisfy the requirements for any erasure-coding scheme. For example, the following message is expected when the selected storage pool includes exactly two sites. If you want to use erasure coding to protect object data, you must select a storage pool with exactly one site or a storage pool with three or more sites.

Scheme

Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
No erasure coding schemes are supported for the selected storage pool because it contains two sites. You must select a storage pool that contains exactly one site or a storage pool that contains at least three sites.			

- Your grid includes only one site and you selected the default storage pool, All Storage Nodes, or any storage pool that includes the default site, All Sites.

Create EC Profile

You cannot change the selected scheme and storage pool after saving the profile.

Profile Name

Storage Pool

3 Storage Nodes across 1 site(s)

Scheme

Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
No erasure coding schemes are available for the selected storage pool. The storage pool includes the <b>All Sites</b> site, so it cannot be used in an Erasure Coding profile for a one-site grid.			

Cancel Save

- The erasure-coding scheme and storage pool you selected overlap with another Erasure Coding profile.

## Create EC Profile

You cannot change the selected scheme and storage pool after saving the profile.

Profile Name

Storage Pool

9 Storage Nodes across 3 site(s)

### Scheme

	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
<input type="radio"/>	6+3	50%	3	Yes
<input checked="" type="radio"/>	2+1	50%	1	Yes
<input type="radio"/>	4+2	50%	2	Yes

The selected storage pool and erasure coding scheme overlap an existing Erasure Coding profile. Use caution if you apply this new profile to objects already protected by the other profile. When a new profile is applied to existing erasure-coded objects, entirely new erasure-coded fragments are created, which might cause resource issues.

Cancel Save

In this example, a warning message appears because another Erasure Coding profile is using the 2+1 scheme and the storage pool for the other profile also uses one of the sites in the All 3 Sites storage pool.

While you are not prevented from creating this new profile, you must be very careful when you start using it in the ILM policy. If this new profile is applied to existing erasure-coded objects already protected by the other profile, StorageGRID will create an entirely new set of object fragments. It will not reuse the existing 2+1 fragments. Resource issues might occur when you migrate from one Erasure Coding profile to the other, even though the erasure-coding schemes are the same.

5. If more than one erasure-coding scheme is listed, select the one you want to use.

When deciding which erasure-coding scheme to use, you should balance fault tolerance (achieved by having more parity segments) against the network traffic requirements for repairs (more fragments equals more network traffic). For example, when deciding between a 4+2 scheme and 6+3 scheme, select the 6+3 scheme if additional parity and fault tolerance are required. Select the 4+2 scheme if network resources are constrained to reduce network usage during node repairs.

6. Click **Save**.

## Renaming an Erasure Coding profile

You might want to rename an Erasure Coding profile to make it more obvious what the profile does.


### What you'll need

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

### Steps

1. Select **ILM > Erasure Coding**.

The Erasure Coding Profiles page appears. The **Rename** and **Deactivate** buttons are both disabled.



	Profile	Status	Storage Pool	Storage Nodes	Sites	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
<input type="radio"/>	DC1 2-1		DC1	3	1	2+1	50	1	No
<input type="radio"/>	DC2 2-1		DC2	3	1	2+1	50	1	No
<input type="radio"/>	DC3 2-1		DC3	3	1	2+1	50	1	No
<input checked="" type="radio"/>	All sites 6-3	Deactivated	All 3 Sites	9	3	6+3	50	3	Yes

2. Select the profile you want to rename.

The **Rename** and **Deactivate** buttons become enabled.

3. Click **Rename**.

The Rename EC Profile dialog box appears.



**Rename EC Profile**

Profile Name

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.



From day  store

Type  Location  Copies

**Erasure Coding profile name** (points to '6 plus 3')

**Storage pool name** (points to 'All 3 sites')



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. Click **Save**.

### Deactivating 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.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

- You must have specific access permissions.
- You must have confirmed 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

When you deactivate an Erasure Coding profile, the profile still appears on the Erasure Coding Profiles page, but its status is **Deactivated**.

Profile	Status	Storage Pool	Storage Nodes	Sites	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
DC1 2-1		DC1	3	1	2+1	50	1	No
DC2 2-1		DC2	3	1	2+1	50	1	No
DC3 2-1		DC3	3	1	2+1	50	1	No
All sites 6-3	Deactivated	All 3 Sites	9	3	6+3	50	3	Yes

You can no longer use an Erasure Coding profile that has been deactivated. A deactivated profile is not shown when you create the placement instructions for an ILM rule. You cannot reactivate a deactivated profile.

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 **ILM > Erasure Coding**.

The Erasure Coding Profiles page appears. The **Rename** and **Deactivate** buttons are both disabled.

2. Review the **Status** column to confirm that the Erasure Coding profile you want to deactivate is not used in any ILM rules.

You cannot deactivate an Erasure Coding profile if it is used in any ILM rule. In the example, the **2\_1 EC Profile** is used in at least one ILM rule.

Profile	Status	Storage Pool	Storage Nodes	Sites	Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
2_1 EC Profile	Used In ILM Rule	DC1	3	1	2+1	50	1	No
Site 1 EC Profile	Deactivated	DC1	3	1	2+1	50	1	No

3. If the profile is used in an ILM rule, follow these steps:

- a. Select **ILM > Rules**.

- b. For each rule listed, select the radio button and review the retention diagram to determine if the rule uses the Erasure Coding profile you want to deactivate.

In the example, the **Three site EC for larger objects** rule uses a storage pool called **All 3 Sites** and the **All sites 6-3** Erasure Coding profile. Erasure Coding profiles are represented by this icon: 

## ILM Rules

Information lifecycle management (ILM) rules determine how and where object data is stored over time. Every object ingested into StorageGRID is evaluated against the ILM rules that make up the active ILM policy. Use this page to manage and view ILM rules. You cannot edit or remove an ILM rule that is used by an active or proposed ILM policy.

+ Create
Clone
Edit
Remove

Name	Used In Active Policy	Used In Proposed Policy
<input type="radio"/> 2 copy replication for smaller objects	✓	
<input checked="" type="radio"/> Three site EC for larger objects	✓	
<input type="radio"/> Make 2 Copies		

Three site EC for larger objects

**Description:** 6-3 erasure coding at 3 sites for objects larger than 200 KB

**Ingest Behavior:** Balanced

**Reference Time:** Ingest Time

**Filtering Criteria:**

Matches all of the following metadata:
 

System Metadata
Object Size (MB)
greater than
0.2

**Retention Diagram:**

The retention diagram shows a horizontal timeline starting at 'Day 0' with a 'Trigger' event. A yellow box labeled 'All 3 Sites (All sites 6-3)' is positioned at the start of the timeline. A blue arrow points to the right, labeled 'Forever' under the 'Duration' section.

- c. If the ILM rule uses the Erasure Coding profile you want to deactivate, determine if the rule is used in either the active ILM policy or a proposed policy.

In the example, the **Three site EC for larger objects** rule is used in the active 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	<ol style="list-style-type: none"> <li>1. Edit or delete all affected ILM rules. If you edit the rule, remove all placements that use the Erasure Coding profile.</li> <li>2. Continue with this procedure.</li> </ol>	<a href="#">Working with ILM rules and ILM policies</a>

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 the active ILM policy	<ol style="list-style-type: none"> <li>1. Clone the active policy.</li> <li>2. Remove the ILM rule that uses the Erasure Coding profile.</li> <li>3. Add one or more new ILM rules to ensure objects are protected.</li> <li>4. Save, simulate, and activate the new policy.</li> <li>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.</li> </ol> <p><b>Note:</b> 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.</p> <p>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.</p> <ol style="list-style-type: none"> <li>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.</li> <li>7. Continue with this procedure.</li> </ol>	<ul style="list-style-type: none"> <li>• <a href="#">Creating an ILM policy</a></li> <li>• <a href="#">Working with ILM rules and ILM policies</a></li> </ul>
In an ILM rule that is currently in a proposed ILM policy	<ol style="list-style-type: none"> <li>1. Edit the proposed policy.</li> <li>2. Remove the ILM rule that uses the Erasure Coding profile.</li> <li>3. Add one or more new ILM rules to ensure all objects are protected.</li> <li>4. Save the proposed policy.</li> <li>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.</li> <li>6. Continue with this procedure.</li> </ol>	<ul style="list-style-type: none"> <li>• <a href="#">Creating an ILM policy</a></li> <li>• <a href="#">Working with ILM rules and ILM policies</a></li> </ul>

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 in a historical ILM policy	<ol style="list-style-type: none"> <li>1. Edit or delete the rule. If you edit the rule, remove all placements that use the Erasure Coding profile. (The rule will now appear as a historical rule in the historical policy.)</li> <li>2. Continue with this procedure.</li> </ol>	<ul style="list-style-type: none"> <li>• <a href="#">Working with ILM rules and ILM policies</a></li> </ul>

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 EC Profile dialog box appears.



5. If you are sure you want to deactivate the profile, select **Deactivate**.

- 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.
- 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.

#### Configuring 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.

#### What you'll need

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

#### About this task

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 and must contain at least 2 and no more than 32 characters. 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 cannot delete or modify a region if it is currently used within the active ILM policy or the proposed ILM policy.
- If the region used as the advanced filter in an ILM rule is invalid, it is still possible to add that rule to the proposed policy. However, an error occurs if you attempt to save or activate the proposed 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 cannot be modified or removed.

### Regions (optional and S3 only)

Define any regions you want to use for the Location Constraint advanced filter in ILM rules. Then, use these exact names when creating S3 buckets. (Region names are case sensitive.)

---

Region 1	<input type="text" value="us-east-1 (required)"/>	
Region 2	<input type="text" value="us-west-1"/>	<b>+</b> <b>x</b>

### 2. To add a region:


- a. Click the insert icon **+** to the right of the last entry.

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, click the delete icon **x**.

An error message appears if you attempt to remove a region that is currently used in the active policy or the proposed policy.



4. When you are done making changes, click **Save**.

You can now select these regions from the **Location Constraint** list on the Advanced Filtering page of the Create ILM rule wizard.

## Related information

[Using advanced filters in ILM rules](#)

## Creating an ILM rule

ILM rules allow you to manage the placement of object data over time. To create an ILM rule, you use the Create ILM Rule wizard.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- If you want to specify which tenant accounts this rule applies to, you must have the Tenant Accounts permission or you must 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 bucket for S3 or by container for Swift.
- If you are creating replicated copies, you must have configured any storage pools or Cloud Storage Pools you plan to use.
- If you are creating erasure-coded copies, you must have configured an Erasure Coding profile.
- You must be familiar with the [data-protection options for ingest](#).
- If you need to create a compliant rule for use with S3 Object Lock, you must be familiar with the [requirements for S3 Object Lock](#).



To create the default ILM rule for a policy, use this procedure instead: [Creating a default ILM rule](#).

## 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 option to use for data protection option at ingest (Balanced, Strict, or Dual commit)

## Steps

1. Select **ILM > Rules**.

The ILM Rules page appears, with the stock rule, Make 2 Copies, selected.

### ILM Rules

Information lifecycle management (ILM) rules determine how and where object data is stored over time. Every object ingested into StorageGRID is evaluated against the ILM rules that make up the active ILM policy. Use this page to manage and view ILM rules. You cannot edit or remove an ILM rule that is used by an active or proposed ILM policy.

The screenshot shows the ILM Rules management interface. At the top, there are buttons for '+ Create', 'Clone', 'Edit', and 'Remove'. Below is a table with the following structure:

Name	Used In Active Policy	Used In Proposed Policy
Make 2 Copies	✓	

Below the table, the configuration for the selected rule 'Make 2 Copies' is shown:

- Ingest Behavior:** Dual commit
- Reference Time:** Ingest Time
- Filtering Criteria:** Matches all objects.
- Retention Diagram:** A diagram showing a trigger at Day 0 for 'All Storage Nodes' with a duration of 'Forever'.



The ILM Rules page looks slightly different if the global S3 Object Lock setting has been enabled for the StorageGRID system. The summary table includes a **Compliant** column, and the details for the selected rule include a **Compliant** field.

2. Select **Create**.

Step 1 (Define Basics) of the Create ILM Rule wizard appears. You use the Define basics page to define which objects the rule applies to.

## Related information

[Use S3](#)

[Use Swift](#)

[Configuring Erasure Coding profiles](#)

[Configuring storage pools](#)

[Using Cloud Storage Pools](#)

[Data-protection options for ingest](#)

[Managing objects with S3 Object Lock](#)

### Step 1 of 3: Define basics

Step 1 (Define Basics) of the Create ILM Rule wizard allows you to define the rule's basic and advanced filters.

#### 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.

Create ILM Rule Step 1 of 3: Define Basics

Name

Description

Tenant Accounts (optional)

Bucket Name  Value

[Advanced filtering... \(0 defined\)](#)

Cancel Next

#### Steps

1. Enter a unique name for the rule in the **Name** field.

You must enter between 1 and 64 characters.

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.

Name

Description

3. Optionally, select one or more S3 or Swift tenant accounts to which this rule applies. If this rule applies to all tenants, leave this field blank.

If you do not have either the Root Access permission or the Tenant Accounts permission, you cannot 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 or Swift containers to which this rule applies.

If **matches all** is selected (default), the rule applies to all S3 buckets or Swift containers.

5. Optionally, select **Advanced filtering** to specify additional filters.

If you do not configure advanced filtering, the rule applies to all objects that match the basic filters.



If this rule will create erasure-coded copies, select **Advanced filtering**. Then, add the **Object Size (MB)** advanced filter and set it to **greater than 0.2**. The size filter ensures that objects that are 2 MB or smaller will not be erasure coded.

6. Select **Next**.

Step 2 (Define Placements) appears.

### Related information

[What ILM rule filtering is](#)

[Using advanced filters in ILM rules](#)

[Step 2 of 3: Define placements](#)

### Using 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 (microseconds)	<ul style="list-style-type: none"><li>• equals</li><li>• does not equal</li><li>• less than</li><li>• less than or equals</li><li>• greater than</li><li>• greater than or equals</li></ul>	Time and date the object was ingested.  <b>Note:</b> 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 are not moved unnecessarily.

Metadata type	Supported operators	Metadata value
Key	<ul style="list-style-type: none"> <li>• equals</li> <li>• does not equal</li> <li>• contains</li> <li>• does not contain</li> <li>• starts with</li> <li>• does not start with</li> <li>• ends with</li> <li>• does not end with</li> </ul>	<p>All or part of a unique S3 or Swift object key.</p> <p>For example, you might want to match objects that end with <code>.txt</code> or start with <code>test-object/</code>.</p>
Last Access Time (microseconds)	<ul style="list-style-type: none"> <li>• equals</li> <li>• does not equal</li> <li>• less than</li> <li>• less than or equals</li> <li>• greater than</li> <li>• greater than or equals</li> <li>• exists</li> <li>• does not exist</li> </ul>	<p>Time and date the object was last retrieved (read or viewed).</p> <p><b>Note:</b> If you plan to use last access time as an advanced filter, Last Access Time updates must be enabled for the S3 bucket or Swift container.</p> <p><a href="#">Using Last Access Time in ILM rules</a></p>
Location Constraint (S3 only)	<ul style="list-style-type: none"> <li>• equals</li> <li>• does not equal</li> </ul>	<p>The region where an S3 bucket was created. Use <b>ILM &gt; Regions</b> to define the regions that are shown.</p> <p><b>Note:</b> 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.</p> <p><a href="#">Configuring regions (optional and S3 only)</a></p>
Object Size (MB)	<ul style="list-style-type: none"> <li>• equals</li> <li>• not equals</li> <li>• less than</li> <li>• less than or equals</li> <li>• greater than</li> <li>• greater than or equals</li> </ul>	<p>The object's size in MB.</p> <p>To filter on object sizes smaller than 1 MB, type in a decimal value. For example, set the <b>Object Size (MB)</b> advanced filter to <b>greater than 0.2</b> for any rule that makes erasure-coded copies. This setting ensures that erasure coding is not used for objects 200 KB or smaller.</p> <p><b>Note:</b> Your browser type and locale settings control whether you need to use a period or a comma as the decimal separator.</p>

Metadata type	Supported operators	Metadata value
User Metadata	<ul style="list-style-type: none"> <li>contains</li> <li>ends with</li> <li>equals</li> <li>exists</li> <li>does not contain</li> <li>does not end with</li> <li>does not equal</li> <li>does not exist</li> <li>does not start with</li> <li>starts with</li> </ul>	<p>Key-value pair, where <b>User Metadata Name</b> is the key and <b>User Metadata Value</b> is the value.</p> <p>For example, to filter on objects that have user metadata of <code>color=blue</code>, specify <code>color</code> for <b>User Metadata Name</b>, <code>equals</code> for the operator, and <code>blue</code> for <b>User Metadata Value</b>.</p> <p><b>Note:</b> User-metadata names are not case sensitive; user-metadata values are case sensitive.</p>
Object Tag (S3 only)	<ul style="list-style-type: none"> <li>contains</li> <li>ends with</li> <li>equals</li> <li>exists</li> <li>does not contain</li> <li>does not end with</li> <li>does not equal</li> <li>does not exist</li> <li>does not start with</li> <li>starts with</li> </ul>	<p>Key-value pair, where <b>Object Tag Name</b> is the key and <b>Object Tag Value</b> is the value.</p> <p>For example, to filter on objects that have an object tag of <code>Image=True</code>, specify <code>Image</code> for <b>Object Tag Name</b>, <code>equals</code> for the operator, and <code>True</code> for <b>Object Tag Value</b>.</p> <p><b>Note:</b> Object tag names and object tag values are case sensitive. You must enter these items exactly as they were defined for the object.</p>

### Specifying 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.

## Advanced Filtering

Use advanced filtering if you want a rule to apply only to specific objects. You can filter objects based on their system metadata, user metadata, or object tags (S3 only). When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the advanced filter.

### Objects between 10 and 100 MB

Matches all of the following metadata:

Object Size (MB)	greater than or equals	10	+	x	
Object Size (MB)	less than or equals	100	+	x	
+					x

Cancel

Remove Filters

Save

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 a 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.

## Advanced Filtering

Use advanced filtering if you want a rule to apply only to specific objects. You can filter objects based on their system metadata, user metadata, or object tags (S3 only). When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the advanced filter.

### Multiple filters

**Matches all of the following metadata:**

User Metadata	camera_type	equals	Brand A	+	x
---------------	-------------	--------	---------	---	---

+

+ x |

**Or matches all of the following metadata:**

User Metadata	camera_type	equals	Brand B	+	x
Object Size (MB)		less than or equals	10	+	x

+

+ x |

Cancel

Remove Filters

Save

### Related information

[Using Last Access Time in ILM rules](#)

[Configuring regions \(optional and S3 only\)](#)

### Step 2 of 3: Define placements

Step 2 (Define Placements) 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.

### 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.

This example ILM rule creates two replicated copies for the first year. Each copy is saved in a storage pool at a different site. After one year, a 2+1 erasure-coded copy is made and saved at only one site.

Configure placement instructions to specify how you want objects matched by this rule to be stored.

**Example rule**  
 Two copies for one year, then EC forever

Reference Time Ingest Time

**Placements** Sort by start day

From day 0 store for 365 days Add Remove

Type replicated Location DC1 x DC2 x Add Pool Copies 2 + x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

From day 365 store forever Add Remove

Type erasure coded Location DC1 (2 plus 1) Copies 1 + x

**Retention Diagram** Refresh

The diagram shows a timeline starting at 'Trigger' (Day 0) and ending at 'Forever'. 
 

- At Day 0, two copies are created: DC1 (blue bar) and DC2 (orange bar).
- At Year 1, the DC1 and DC2 copies are replaced by a single erasure-coded copy (DC1 (2 plus 1), orange bar with a pencil icon).
- The duration of the rule is marked as '1 years' and 'Forever'.

Cancel Back Next

### Steps

1. 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).  <b>Note:</b> To use this option, updates to Last Access Time must be enabled for the S3 bucket or Swift container.  <a href="#">Using Last Access Time in ILM rules</a>

Option	Description
Noncurrent Time	<p>The time an object version became noncurrent because a new version was ingested and replaced it as the current version.</p> <p><b>Note:</b> Noncurrent Time applies only to S3 objects in versioning-enabled buckets.</p> <p>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.”</p>
User Defined Creation Time	A time specified in user-defined metadata.



If you want to create a compliant rule, you must select **Ingest Time**.

- In the **Placements** section, select 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 (“day 0 for 365 days”). At least one instruction must start at day 0.

- If you want to create replicated copies:

- From the **Type** drop-down list, select **replicated**.
- In the **Location** field, select **Add Pool** for each storage pool 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 cannot 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, the system distributes the 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, do not specify the default All Storage Nodes storage pool and another storage pool.

**Placements** ⓘ ⌵ Sort by start day

From day  store

Type  Location  Copies

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

c. 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. If only one replicated copy of an object exists during a time period, 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.



To avoid these risks, do one or more of the following:

- Increase the number of copies for the time period.
- Click the plus sign icon **+** to create additional copies during the time period. Then, select a different storage pool or a Cloud Storage Pool.
- Select **erasure coded** for Type, instead of **replicated**.  
You can safely ignore this warning if this rule already creates multiple copies for all time periods.

d. If you specified only one storage pool, ignore the **Temporary location** field.



Temporary locations are deprecated and will be removed in a future release.

4. If you want to store objects in a Cloud Storage Pool:

- a. From the **Type** drop-down list, select **replicated**.
- b. In the **Location** field, select **Add Pool**. Then, select a Cloud Storage Pool.

When using Cloud Storage Pools, keep these rules in mind:

- You cannot select more than one Cloud Storage Pool in a single placement instruction. Similarly, you cannot select a Cloud Storage Pool and a storage pool in the same placement instruction.

If you want to use a Cloud Storage Pool, you must remove any other storage pools or Cloud Storage Pools from this 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.

Type  Location  Add Pool Copies

The number of copies cannot be more than one when a Cloud Storage Pool is selected.

- You cannot 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.

**Placements** Sort by start day

From day  store for  days Add Remove

Type  Location  Add Pool Copies  + x

Type  Location  Add Pool Copies  + x

A rule cannot store more than one object copy in any Cloud Storage Pool at the same time. You must remove one of the Cloud Storage Pools (csp1, csp2) or use multiple placement instructions with dates that do not overlap. **Overlapping days:** 0-10.  
To see the overlapping days on the Retention Diagram, click Refresh.



- 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, as this example shows, 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.

**Placements**

From day  store for  days

Type  Location   Add Pool Copies

Type  Location  Add Pool Copies

- If you want to create an erasure-coded copy:
  - From the **Type** drop-down list, select **erasure coded**.

The number of copies changes to 1. A warning appears if the rule does not have an advanced filter to ignore objects that are 200 KB or smaller.

Do not use erasure coding for objects that are 200 KB or smaller. Select **Back** to return to Step 1. Then, use **Advanced filtering** to set the Object Size (MB) filter to "greater than 0.2".



Do not use erasure coding for objects smaller than 200 KB to avoid the overhead of managing very small erasure-coded fragments.

b. If the object size warning appeared, follow these steps to clear it:

- i. Select **Back** to return to Step 1.
- ii. Select **Advanced filtering**.
- iii. Set the Object Size (MB) filter to “greater than 0.2”.

c. Select the storage location.

The storage location for an erasure-coded copy includes the name of the storage pool, followed by the name of the Erasure Coding profile.

6. Optionally, add different time periods or create additional copies at different locations:

- Click the plus icon to create additional copies at a different location during the same time period.
- Click **Add** to add a different time period to the placement instructions.




Objects are automatically deleted at the end of the final time period unless the final time period ends with **forever**.

7. Click **Refresh** to update the Retention Diagram and to confirm your placement instructions.

Each line in the diagram shows where and when object copies will be placed. The type of copy is represented by one of the following icons:

	Replicated copy
	Erasure-coded copy
	Cloud Storage Pool copy

In this example, two replicated copies will be saved to two storage pools (DC1 and DC2) for one year. Then, an erasure-coded copy will be saved for an additional 10 years, using a 6+3 erasure-coding scheme at three sites. After 11 years, the objects will be deleted from StorageGRID.



## 8. Click **Next**.

Step 3 (Define Ingest Behavior) appears.

### Related information

[What ILM rule placement instructions are](#)

[Example 4: ILM rules and policy for S3 versioned objects](#)

[Why you should not use single-copy replication](#)

[Managing objects with S3 Object Lock](#)

[Using a storage pool as a temporary location \(deprecated\)](#)

[Step 3 of 3: Define ingest behavior](#)

### Using 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 always enabled for Swift containers, but 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	<ul style="list-style-type: none"> <li>No, for the source copy</li> <li>Yes, for the destination copy</li> </ul>	<ul style="list-style-type: none"> <li>Yes, for the source copy</li> <li>Yes, for the destination copy</li> </ul>
Request to complete a multipart upload	Yes, for the assembled object	Yes, for the assembled object

### Related information

[Use S3](#)

[Use a tenant account](#)

### Step 3 of 3: Define ingest behavior

Step 3 (Define ingest behavior) 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.

**Create ILM Rule** Step 3 of 3: Define ingest behavior

Select the data protection option to use when objects are ingested:

- Strict**  
Always uses this rule's placements on ingest. Ingest fails when this rule's placements are not possible.
- Balanced**  
Optimum ILM efficiency. Attempts this rule's placements on ingest. Creates interim copies when that is not possible.
- Dual commit**  
Creates interim copies on ingest and applies this rule's placements later.

Cancel
Back
Save

### Steps

1. Select the data protection option to use when objects are ingested:

Option	Description
Strict	Always uses this rule's placements on ingest. Ingest fails when this rule's placements are not possible.

Option	Description
Balanced	Optimum ILM efficiency. Attempts this rule's placements on ingest. Creates interim copies when that is not possible.
Dual commit	Creates interim copies on ingest and applies this rule's placements later.

Balanced offers a combination of data security and efficiency that is suitable in most cases. Strict or Dual commit are generally used to meet specific requirements.

See “What the data-protection options for ingest are” and “Advantages and disadvantages of each data-protection option” for more information.



An error message appears if you select the Strict or Balanced option and the rule uses one of these placements:

- A Cloud Storage Pool at day 0
- An Archive Node at day 0
- A Cloud Storage Pool or an Archive Node when the rule uses a User Defined Creation Time as a Reference Time

2. Click **Save**.

The ILM rule is saved. The rule does not become active until it is added to an ILM policy and that policy is activated.

### Related information

[Data-protection options for ingest](#)

[Advantages, disadvantages, and limitations of the data-protection options](#)

[Example 5: ILM rules and policy for Strict ingest behavior](#)

[Creating an ILM policy](#)

### Creating a default ILM rule

Every ILM policy must have a default rule that does not filter objects. Before creating an ILM policy, you must create at least one ILM rule that can be used as the default rule for the policy.

### What you'll need

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

### About this task

The default rule is the last rule to be evaluated in an ILM policy, so it cannot use any filters. The placement instructions for the default rule are applied to any objects that are not matched by another rule in the policy.

In this example policy, the first rule applies only to objects belonging to Tenant A. The default rule, which is last, applies to objects belonging to all other tenant accounts.

+ Select Rules			
Default	Rule Name	Tenant Account	Actions
	Erasure Coding for Tenant A	Tenant A (94793396288150002349)	✘
✓	2 Copies 2 Data Centers	Ignore	✘

When you create the default rule, keep these requirements in mind:

- The default rule is automatically placed as the last rule in the policy.
- The default rule cannot use any basic or advanced filters.
- The default rule should create replicated copies.



Do not 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 for the active or proposed policy must be compliant.

## Steps

1. Select **ILM > Rules**.

The ILM Rules page appears.

2. Select **Create**.

Step 1 (Define Basics) of the Create ILM Rule wizard appears.

3. Enter a unique name for the rule in the **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** field blank.

The default rule must apply to all S3 buckets and Swift containers.

7. Do not select **Advanced filtering**

The default rule cannot specify any filters.

8. Select **Next**.

Step 2 (Define Placements) appears.

9. 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.



Do not 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 0.2** 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 cannot be saved in a Cloud Storage Pool.
  - Object copies cannot be saved on Archive Nodes.
  - 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.”

10. Click **Refresh** to update the Retention Diagram and to confirm your placement instructions.

11. Click **Next**.

Step 3 (Define Ingest Behavior) appears.

12. Select the data protection option to use when objects are ingested, and select **Save**.

### Creating an ILM policy

When you create an ILM policy, you start by selecting and arranging the ILM rules. Then, you verify the behavior of your proposed policy by simulating it against previously ingested objects. When you are satisfied that the proposed policy is functioning as intended, you can activate it to create the active policy.



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.

### Considerations for creating an ILM policy

- Use the system’s built-in policy, Baseline 2 Copies Policy, in test systems only. 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.
- 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 rule will not be evaluated by any other rule.

- The last rule in every ILM policy is the default ILM rule, which cannot 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.

### Related information

[What an ILM policy is](#)

[Example 6: Changing an ILM policy](#)

### Creating a proposed ILM policy

You can create a proposed ILM policy from scratch, or you can clone the current active policy if you want to start with the same set of rules.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have created the ILM rules you want to add to the proposed policy. As required, you can save a proposed policy, create additional rules, and then edit the proposed policy to add the new rules.
- You must have created a default ILM rule for the policy that does not contain any filters.

[Creating a default ILM rule](#)

### About this task

Typical reasons for creating a proposed ILM policy include:

- You added a new site and need to use new ILM rules to place objects at that site.
- You are decommissioning a site and you need to remove all rules that refer to the site.
- You added a new tenant that has special data protection requirements.
- You started to use a Cloud Storage Pool.



Use the system's built-in policy, Baseline 2 Copies Policy, in test systems only. 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.



If the global S3 Object Lock setting has been enabled, the steps for creating a policy are slightly different. You must ensure that the ILM policy is compliant with the requirements of buckets that have S3 Object Lock enabled.

[Creating an ILM policy after S3 Object Lock is enabled](#)

### Steps

1. Select **ILM > Policies**.

The ILM Policies page appears. From this page, you can review the list of proposed, active, and historical policies; create, edit, or remove a proposed policy; clone the active policy; or view the details for any policy.

## ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

+ Create Proposed Policy
Clone
Edit
Remove

Policy Name	Policy State	Start Date	End Date
<input checked="" type="radio"/> Baseline 2 Copies Policy	Active	2017-07-17 12:00:45 MDT	

**Viewing Active Policy - Baseline 2 Copies Policy**

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

*Rules are evaluated in order, starting from the top.*

Rule Name	Default	Tenant Account
Make 2 Copies	✓	Ignore

Simulate
Activate

## 2. Determine how you want to create the proposed ILM policy.

Option	Steps
Create a new proposed policy that has no rules already selected	<p>a. If a proposed ILM policy currently exists, select that policy, and click <b>Remove</b>.</p> <p style="margin-left: 20px;">You cannot create a new proposed policy if a proposed policy already exists.</p> <p>b. Click <b>Create Proposed Policy</b>.</p>
Create a proposed policy based on the active policy	<p>a. If a proposed ILM policy currently exists, select that policy, and click <b>Remove</b>.</p> <p style="margin-left: 20px;">You cannot clone the active policy if a proposed policy already exists.</p> <p>b. Select the active policy from the table.</p> <p>c. Click <b>Clone</b>.</p>
Edit the existing proposed policy	<p>a. Select the proposed policy from the table.</p> <p>b. Click <b>Edit</b>.</p>

The Configure ILM Policy dialog box appears.

If you are creating a new proposed policy, all fields are blank and no rules are selected.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### 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.

Default	Rule Name	Tenant Account	Actions
<i>No rules selected.</i>			

If you are cloning the active policy, the **Name** field shows the name of the active policy, appended by a version number ("v2" in the example). The rules used in the active policy are selected and shown in their current order.

Name

Reason for change

3. Enter a unique name for the proposed policy in the **Name** field.

You must enter at least 1 and no more than 64 characters. If you are cloning the active policy, you can use the current name with the appended version number or you can enter a new name.

4. Enter the reason you are creating a new proposed policy in the **Reason for change** field.

You must enter at least 1 and no more than 128 characters.

5. To add rules to the policy, select **Select Rules**.

The Select Rules for Policy dialog box appears, with all defined rules listed. 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, a button allows you to move the rule to the end of the new policy.

## Select Rules for Policy

### Select Default Rule

This list shows the rules that do not use any filters. Select one rule to be the default rule for the policy. The default rule applies to any objects that do not match another rule in the policy and is always evaluated last. The default rule should retain objects forever.

	Rule Name
<input checked="" type="radio"/>	2 copies at 2 data centers 
<input type="radio"/>	2 copies at 2 data centers for 2 years 
<input type="radio"/>	Make 2 Copies 

### Select Other Rules

The other rules in a policy are evaluated before the default rule and must use at least one filter. Each rule in this list uses at least one filter (tenant account, bucket name, or an advanced filter, such as object size).

	Rule Name	Tenant Account
<input type="checkbox"/>	1-site EC 	—
<input type="checkbox"/>	3-site EC 	—

Cancel

Apply

6. Select a rule name or the more details icon  to view the settings for that rule.

This example shows the details of an ILM rule that makes two replicated copies at two sites.

## Two-Site Replication for Other Tenants


**Description:** Two-Site Replication for Other Tenants

**Ingest Behavior:** Balanced

**Reference Time:** Ingest Time

**Filtering Criteria:** Matches all objects.

**Retention Diagram:**



Close

7. In the **Select Default Rule** section, select one default rule for the proposed policy.

The default rule applies to any objects that do not match another rule in the policy. The default rule cannot use any filter and is always evaluated last.



If no rule is listed in the Select Default Rule section, you must exit the ILM policy page and create a default rule.

[Creating a default ILM rule](#)



Do not use the Make 2 Copies stock 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.

8. In the **Select Other Rules** section, select any other rules you want to include in the policy.

The other rules are evaluated before the default rule and must use at least one filter (tenant account, bucket name, or an advanced filter, such as object size).

9. When you are done selecting rules, select **Apply**.

The rules you selected are listed. The default rule is at the end, with the other rules above it.

#### 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

	Default	Rule Name	Tenant Account	Actions
+		3-site EC	Ignore	✕
+		1-site EC	Ignore	✕
	✓	2 copies at 2 data centers	Ignore	✕

Cancel
Save

A warning appears if the default 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 longer).



	Default	Rule Name	Tenant Account	Actions
+		3-site EC	Ignore	✕
+		1-site EC	Ignore	✕
	✓	2 copies at 2 data centers for 2 years	Ignore	✕

The default ILM rule in this policy does not retain objects forever. Confirm this is the behavior you expect. Otherwise, any objects that are not matched by another rule will be deleted after 720 days.

10. Drag and drop the rows for the non-default rules to determine the order in which these rules will be evaluated.

You cannot move the default rule.



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.

11. As required, click the delete icon to delete any rules that you do not want in the policy, or select **Select Rules** to add more rules.

12. When you are done, select **Save**.

The ILM Policies page is updated:

- The policy you saved is shown as Proposed. Proposed policies do not have start and end dates.
- The **Simulate** and **Activate** buttons are enabled.

#### ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

[+ Create Proposed Policy](#) [Clone](#) [Edit](#) [Remove](#)

Policy Name	Policy State	Start Date	End Date
<input checked="" type="radio"/> Data Protection for Three Sites	Proposed		
<input type="radio"/> Data Protection for Two Sites	Active	2020-09-18 16:01:24 MDT	
<input type="radio"/> Baseline 2 Copies Policy	Historical	2020-09-17 21:32:57 MDT	2020-09-18 16:01:24 MDT

### Viewing Proposed Policy - Data Protection for Three Sites

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- 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.

See [Managing objects with information lifecycle management](#) for more information.

This policy contains a rule that makes an erasure-coded copy. Confirm that at least one rule uses the Object Size advanced filter to prevent objects that are 200 KB or smaller from being erasure coded. See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

**Reason for change:** Added a third site

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
One-Site Erasure Coding for Tenant A <a href="#">🔗</a>		Tenant A (20033011709864740158)
Three-Site Replication for Other Tenants <a href="#">🔗</a>	✓	Ignore

[Simulate](#) [Activate](#)

13. Go to [Simulating an ILM policy](#).

## Related information

[What an ILM policy is](#)

[Managing objects with S3 Object Lock](#)

### Creating an ILM policy after S3 Object Lock is enabled

If the global S3 Object Lock setting is enabled, the steps for creating a policy are slightly different. You must ensure that the ILM policy is compliant with the requirements of buckets that have S3 Object Lock enabled.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- The global S3 Object Lock setting must already be enabled for the StorageGRID system.



If the global S3 Object Lock setting has not been enabled, use the general instructions for creating a proposed policy instead.

### Creating a proposed ILM policy

- You must have created the compliant and non-compliant ILM rules you want to add to the proposed policy. As required, you can save a proposed policy, create additional rules, and then edit the proposed policy to add the new rules.

#### Example 7: Compliant ILM policy for S3 Object Lock

- You must have created a compliant default ILM rule for the policy.

### Creating a default ILM rule

## Steps

1. Select **ILM > Policies**.

The ILM Policies page appears. If the global S3 Object Lock setting is enabled, the ILM Policies page indicates which ILM rules are compliant.

#### ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

The screenshot shows the ILM Policies interface. At the top, there are buttons for '+ Create Proposed Policy', 'Clone', 'Edit', and 'Remove'. Below this is a table with the following data:

Policy Name	Policy State	Start Date	End Date
Baseline 2 Copies Policy	Active	2021-02-04 01:04:29 MST	

Below the table is a section titled 'Viewing Active Policy - Baseline 2 Copies Policy'. It contains the following text: 'Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active. Rules are evaluated in order, starting from the top. The policy's default rule must be compliant.'

Below this text is a table with the following data:

Rule Name	Default	Compliant	Tenant Account
Make 2 Copies	✓	✓	Ignore

At the bottom right of this section are buttons for 'Simulate' and 'Activate'.

2. Enter a unique name for the proposed policy in the **Name** field.

You must enter at least 1 and no more than 64 characters.

3. Enter the reason you are creating a new proposed policy in the **Reason for change** field.

You must enter at least 1 and no more than 128 characters.

4. To add rules to the policy, select **Select Rules**.

The Select Rules for Policy dialog box appears, with all defined rules listed.

- The Select Default Rule section lists the rules that can be the default for a compliant policy. It includes compliant rules that do not use filters.
- The Select Other Rules section lists the other compliant and non-compliant rules that can be selected for this policy.

## Select Rules for Policy

### Select Default Rule

This list shows the rules that are compliant and do not use any filters. Select one rule to be the default rule for the policy. The default rule applies to any objects that do not match another rule in the policy and is always evaluated last.

	Rule Name
<input type="radio"/>	Default Compliant Rule: Two Copies Two Data Centers 
<input type="radio"/>	Make 2 Copies 

### Select Other Rules

The other rules in a policy are evaluated before the default rule. If you need a different "default" rule for objects in non-compliant S3 buckets, select one non-compliant rule that does not use a filter. Any other rules in the policy must use at least one filter (tenant account, bucket name, or an advanced filter, such as object size).

	Rule Name	Compliant	Uses Filter	Is Selectable
<input type="checkbox"/>	Compliant Rule: EC for bank-records bucket - Bank of ABC 	✓	✓	Yes
<input type="checkbox"/>	Non-Compliant Rule: Use Cloud Storage Pool 			Yes

Cancel

Apply

5. Select a rule name or the more details icon  to view the settings for that rule.
6. In the **Select Default Rule** section, select one default rule for the proposed policy.

The table in this section only lists the rules that are compliant and do not use any filters.



If no rule is listed in the Select Default Rule section, you must exit the ILM policy page and create a default rule that is compliant.

[Creating a default ILM rule](#)



Do not use the Make 2 Copies stock 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.

7. In the **Select Other Rules** section, select any other rules you want to include in the policy.
  - a. If you need a different "default" rule for objects in non-compliant S3 buckets, optionally select one non-compliant rule that does not use a filter.

For example, you might want to use a Cloud Storage Pool or an Archive Node to store objects in buckets that do not have S3 Object Lock enabled.



You can only select one non-compliant rule that does not use a filter. As soon as you select one rule, the **Is Selectable** column shows **No** for any other non-compliant rules without filters.

- b. Select any other compliant or non-compliant rules you want to use in the policy.

The other rules must use at least one filter (tenant account, bucket name, or an advanced filter, such as object size).

8. When you are done selecting the rules, select **Apply**.

The rules you selected are listed. The default rule is at the end, with the other rules above it. If you also selected a non-compliant “default” rule, that rule is added as the second-to-last rule in the policy.

In this example, the last rule, 2 Copies 2 Data Centers, is the default rule: it is compliant and has no filters. The second-to-last rule, Cloud Storage Pool, also has no filters but it is not compliant.

### Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

#### 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 (and any non-compliant rule without a filter) will be automatically placed at the end of the policy and cannot be moved.

Default	Rule Name	Compliant	Tenant Account	Actions
	Compliant Rule: EC for bank-records bucket - Bank of ABC	✓	Bank of ABC (90767802913525281639)	✗
	Non-Compliant Rule: Use Cloud Storage Pool		Ignore	✗
✓	Default Compliant Rule: Two Copies Two Data Centers	✓	Ignore	✗

9. Drag and drop the rows for the non-default rules to determine the order in which these rules will be evaluated.

You cannot move the default rule or the non-compliant “default” rule.



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.

10. As required, click the delete icon to delete any rules that you do not want in the policy, or select **Select Rules** to add more rules.

11. When you are done, select **Save**.

The ILM Policies page is updated:

- The policy you saved is shown as Proposed. Proposed policies do not have start and end dates.
- The **Simulate** and **Activate** buttons are enabled.

## ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

+ Create Proposed Policy Clone Edit Remove

Policy Name	Policy State	Start Date	End Date
Compliant ILM Policy for S3 Object Lock	Proposed		
Compliant ILM Policy	Active	2021-02-05 16:22:53 MST	
Non-Compliant ILM policy	Historical	2021-02-05 15:17:05 MST	2021-02-05 16:22:53 MST
Baseline 2 Copies Policy	Historical	2021-02-04 21:35:52 MST	2021-02-05 15:17:05 MST

### Viewing Proposed Policy - Compliant ILM Policy for S3 Object Lock

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- 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.

See [Managing objects with information lifecycle management](#) for more information.

This policy contains a rule that makes an erasure-coded copy. Confirm that at least one rule uses the Object Size advanced filter to prevent objects that are 200 KB or smaller from being erasure coded. See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: Example policy

Rules are evaluated in order, starting from the top. The policy's default rule must be compliant.

Rule Name	Default	Compliant	Tenant Account
Compliant Rule: EC for bank-records bucket - Bank of ABC		✓	Bank of ABC (90767802913525281639)
Non-Compliant Rule: Use Cloud Storage Pool			Ignore
Default Compliant Rule: Two Copies Two Data Centers	✓	✓	Ignore

Simulate Activate

12. Go to [Simulating an ILM policy](#).

### Simulating an ILM policy

You should simulate a proposed policy on test objects before activating the policy and applying it to your production data. The simulation window provides a standalone environment that is safe for testing policies before they are activated and applied to data in the production environment.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must know the S3 bucket/object-key or the Swift container/object-name for each object you want to test, and you must have already ingested those objects.

#### About this task

You must carefully select the objects you want the proposed policy to test. To simulate a policy thoroughly, you should test at least one object for each filter in each rule.

For example, if a policy includes one rule to match objects in bucket A and another rule to match objects in bucket B, you must select at least one object from bucket A and one object from bucket B to test the policy thoroughly. If the policy includes a default rule to place all other objects, you must test at least one object from another bucket.

When simulating a policy, the following considerations apply:

- After you make changes to a policy, save the proposed policy. Then, simulate the behavior of the saved proposed policy.
- When you simulate a policy, the ILM rules in the policy filter the test objects, so you can see which rule was applied to each object. However, no object copies are made and no objects are placed. Running a simulation does not modify your data, rules, or the policy in any way.
- The Simulation page retains the objects you tested until you close, navigate away from, or refresh the ILM Policies page.
- Simulation returns the name of the matched rule. To determine which storage pool or Erasure Coding profile is in effect, you can view the Retention Diagram by clicking the rule name or the more details icon .
- If S3 Versioning is enabled, the policy is only simulated against the current version of the object.

## Steps

1. Select and arrange the rules, and save the proposed policy.

The policy in this example has three rules:

Rule Name	Filter	Type of Copies	Retention
X-men	<ul style="list-style-type: none"> <li>• Tenant A</li> <li>• User metadata (series=x-men)</li> </ul>	2 copies at two data centers	2 years
PNGs	Key ends with .png	2 copies at two data centers	5 years
Two Copies Two Data Centers	<i>None</i>	2 copies at two data centers	Forever

Viewing Proposed Policy - Example ILM policy

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- 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.

See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

**Reason for change:** Example policy

*Rules are evaluated in order, starting from the top.*

Rule Name	Default	Tenant Account
X-men 		Tenant A (94793396288150002349)
PNGs 		Ignore
Two Copies at Two Data Centers 	✓	Ignore

Simulate
Activate

2. Click **Simulate**.

The Simulation ILM Policy dialog box appears.

- In the **Object** field, enter the S3 bucket/object-key or the Swift container/object-name for a test object, and click **Simulate**.

A message appears if you specify an object that has not been ingested.

- Under **Simulation Results**, confirm that each object was matched by the correct rule.

In the example, the `Havok.png` and `Warpath.jpg` objects were correctly matched by the X-men rule. The `Fullsteam.png` object, which does not include `series=x-men` user metadata, was not matched by the X-men rule but was correctly matched by the PNGs rule. The default rule was not used because all three objects were matched by other rules.

### Simulate ILM Policy - Demo

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object

#### Simulation Results ?

Object	Rule Matched	Previous Match	
photos/Havok.png	X-men		✘
photos/Warpath.jpg	X-men		✘
photos/Fullsteam.png	PNGs		✘

## Examples for simulating ILM policies

These examples show how you can verify ILM rules by simulating the ILM policy before activating it.

### Example 1: Verifying rules when simulating a proposed ILM policy

This example shows how to verify rules when simulating a proposed 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 is the default rule and does not include any filters.

## Viewing Proposed Policy - Example ILM policy

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- 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.

See [Managing objects with information lifecycle management](#) for more information.

This policy contains a rule that makes an erasure-coded copy. Confirm that at least one rule uses the Object Size advanced filter to prevent objects that are 200 KB or smaller from being erasure coded. See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: Example policy

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
Two copies, two years for bucket-a		—
EC objects > 1 MB		—
Two copies, two data centers	✓	—

[Simulate](#) [Activate](#)

## Steps

1. After adding the rules and saving the policy, click **Simulate**.

The Simulate ILM Policy dialog box appears.

2. In the **Object** field, enter the S3 bucket/object-key or the Swift container/object-name for a test object, and click **Simulate**.

The Simulation Results appear, showing which rule in the policy matched each object you tested.

## Simulate ILM Policy - Example ILM policy

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object  [Simulate](#)

### Simulation Results

Object	Rule Matched	Previous Match	
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		✘

[Finish](#)

3. Confirm that each object was matched by the correct rule.

In this example:

- a. bucket-a/bucket-a object.pdf correctly matched the first rule, which filters on objects in bucket-a.
- b. 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.

c. 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: Reordering rules when simulating a proposed 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 do not match the first two rules.

### Viewing Proposed Policy - Demo

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- 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.

See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: new policy

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
PNGs		Ignore
X-men		Tenant A (24365814597594524591)
Two copies two data centers	✓	Ignore

[Simulate](#) [Activate](#)

## Steps

1. After adding the rules and saving the policy, click **Simulate**.
2. In the **Object** field, enter the S3 bucket/object-key or the Swift container/object-name for a test object, and click **Simulate**.

The Simulation Results appear, showing that the Havok.png object was matched by the **PNGs** rule.

## Simulate ILM Policy - Demo

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object

### Simulation Results

Object	Rule Matched	Previous Match	
photos/Havok.png	PNGs 		

However, the rule that the `Havok.png` object was meant to test was the **X-men** rule.

3. To resolve the issue, reorder the rules.
  - a. Click **Finish** to close the Simulate ILM Policy page.
  - b. Click **Edit** to edit the policy.
  - c. Drag the **X-men** rule to the top of the list.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### 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.

<input type="button" value="+ Select Rules"/>				
	Default	Rule Name	Tenant Account	Actions
		X-men 	Tenant A (48713995194927812566)	
		PNGs 	—	
	<input checked="" type="checkbox"/>	Two copies, two data centers 	—	

- d. Click **Save**.
4. Click **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.

## Simulate ILM Policy - Demo

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object

### Simulation Results

Object	Rule Matched	Previous Match	
photos/Havok.png	X-men 	PNGs 	



If you stay on the Configure Policies page, you can re-simulate a policy after making changes without needing to re-enter the names of the test objects.

### Example 3: Correcting a rule when simulating a proposed 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.

## Simulate ILM Policy - Demo

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object

### Simulation Results

Object	Rule Matched	Previous Match	
photos/Beast.jpg	Two copies two data centers 		

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. For each rule in the policy, view the rule settings by clicking the rule name or the more details icon  on any dialog box where the rule is displayed.
2. Review the rule's tenant account, reference time, and filtering criteria.

In this example, the metadata for the X-men rule includes an error. The metadata value was entered as "x-men1" instead of "x-men."

## X-men

Ingest Behavior:           Balanced  
Tenant Account:           06846027571548027538  
Reference Time:           Ingest Time

### Filtering Criteria:

Matches all of the following metadata:


User Metadata

series

equals

x-men1

### Retention Diagram:



Close

3. To resolve the error, correct the rule, as follows:

- If the rule is part of the proposed 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 cannot edit or remove a rule from the active policy.

Option	Description
Cloning the rule	<ol style="list-style-type: none"><li>Select <b>ILM &gt; Rules</b>.</li><li>Select the incorrect rule, and click <b>Clone</b>.</li><li>Change the incorrect information, and click <b>Save</b>.</li><li>Select <b>ILM &gt; Policies</b>.</li><li>Select the proposed policy, and click <b>Edit</b>.</li><li>Click <b>Select Rules</b>.</li><li>Select the check box for the new rule, uncheck the check box for the original rule, and click <b>Apply</b>.</li><li>Click <b>Save</b>.</li></ol>

Option	Description
Editing the rule	<ol style="list-style-type: none"> <li>Select the proposed policy, and click <b>Edit</b>.</li> <li>Click the delete icon <b>✘</b> to remove the incorrect rule, and click <b>Save</b>.</li> <li>Select <b>ILM &gt; Rules</b>.</li> <li>Select the incorrect rule, and click <b>Edit</b>.</li> <li>Change the incorrect information, and click <b>Save</b>.</li> <li>Select <b>ILM &gt; Policies</b>.</li> <li>Select the proposed policy, and click <b>Edit</b>.</li> <li>Select the corrected rule, click <b>Apply</b>, and click <b>Save</b>.</li> </ol>

4. Perform the simulation again.



Because you navigated away from the ILM Policies page to edit the rule, the objects you previously entered for simulation are no longer displayed. You must re-enter the names of the objects.

In this example, the corrected X-men rule now matches the `Beast.jpg` object based on the `series=x-men` user metadata, as expected.

**Simulate ILM Policy - Demo**

Simulates the active ILM policy or, if there is a proposed ILM policy, simulates the proposed ILM policy. Use this simulation to test the current configuration of ILM rules and determine whether ILM rules copy and place object data as intended.

Object

**Simulation Results** ?

Object	Rule Matched	Previous Match	
photos/Beast.jpg	X-men		<b>✘</b>

### Activating the ILM policy

After you add ILM rules to a proposed ILM policy, simulate the policy, and confirm it behaves as you expect, you are ready to activate the proposed policy.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- You must have saved and simulated the proposed ILM policy.



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.

### 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 are not accidentally removed.

- 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 be marked as “Active,” but it will not take effect until all nodes are online and available.

### Steps

1. When you are ready to activate a proposed policy, select the policy on the ILM Policies page and click **Activate**.

A warning message is displayed, prompting you to confirm that you want to activate the proposed policy.

#### Warning

##### Activate the proposed policy

Errors in an ILM policy can cause irreparable data loss. Review and test the policy carefully before activating. Are you sure you want to activate the proposed policy?

Cancel


OK

A prompt appears in the warning message if the default rule for the policy does not retain objects forever. In this example, the retention diagram shows that the default rule will delete objects after 2 years. You must type **2** in the text box to acknowledge that any objects not matched by another rule in the policy will be removed from StorageGRID after 2 years.

## ⚠ Activate the proposed policy

Errors in an ILM policy can cause irreparable data loss. Review and test the policy carefully before activating.

The default rule in this policy does not retain objects forever. Confirm this is the behavior you want by referring to the retention diagram for the default rule:



Now, complete the following prompt:

Any objects that are not matched by another rule in this policy will be deleted after  years.

Are you sure you want to activate the proposed policy?

Cancel

OK

2. Click **OK**.

## Result

When a new ILM policy has been activated:

- The policy is shown with a Policy State of Active in the table on the ILM Policies page. The Start Date entry indicates the date and time the policy was activated.

### ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

Policy Name	Policy State	Start Date	End Date
<input checked="" type="radio"/> New Policy	Active	2017-07-20 18:49:53 MDT	
<input type="radio"/> Baseline 2 Copies Policy	Historical	2017-07-19 21:24:30 MDT	2017-07-20 18:49:53 MDT

- The previously active policy is shown with a Policy State of Historical. The Start Date and End Date entries indicate when the policy became active and when it was no longer in effect.

## Related information

### [Example 6: Changing an ILM policy](#)

### Verifying an ILM policy with object metadata lookup

After you have activated an ILM policy, you should ingest representative test objects into the StorageGRID system. You should then perform an object metadata lookup to confirm that copies are being made as intended and placed in the correct locations.

### What you'll need

- You must have an object identifier, which can be one of:
  - **UUID**: The object's Universally Unique Identifier. Enter the UUID in all uppercase.
  - **CBID**: The object's unique identifier within StorageGRID. You can obtain an object's CBID from the

audit log. Enter the CBID in all uppercase.

- **S3 bucket and object key:** When an object is ingested through the S3 interface, the client application uses a bucket and object key combination to store and identify the object.
- **Swift container and object name:** When an object is ingested through the Swift interface, the client application uses a container and object name combination to store and identify the object.

## Steps

1. 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, S3 bucket/object-key, or Swift container/object-name.

### Object Metadata Lookup

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

Identifier

4. Click **Look Up**.

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

- System metadata, including the object ID (UUID), the object name, the name of the container, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
- Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object.
- For replicated object copies, the current storage location of each copy.
- For erasure-coded object copies, the current storage location of each fragment.
- For object copies in a Cloud Storage Pool, the location of the object, including the name of the external bucket and the object's unique identifier.
- For segmented objects and multipart objects, a list of object segments including segment identifiers and data sizes. For objects with more than 100 segments, only the first 100 segments are shown.
- All object metadata in the unprocessed, internal storage format. This raw metadata includes internal system metadata that is not guaranteed to persist from release to release.

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

## System Metadata

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

## Replicated Copies

Node	Disk Path
99-97	/var/local/rangedb/2/p/06/0B/00nM8HS TFbnQQ CV2E
99-99	/var/local/rangedb/1/p/12/0A/00nM8HS TFboW28 CXG%

## Raw Metadata

```
{
  "TYPE": "CNT",
  "CHND": "A12E96FF-B13F-4905-9E9E-45373F6E7DA8",
  "NAME": "testobject",
  "CBID": "0x8823DE7EC7C10416",
  "PHND": "FEA0AE51-534A-11EA-9FCD-31FF00C36056",
  "PPTH": "source",
  "META": {
    "BASE": {
      "PAWS": "2",

```

5. Confirm that the object is stored in the correct location or locations and that it is 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 cannot 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 the information about understanding audit messages.

## Related information

[Review audit logs](#)

[Use S3](#)

[Use Swift](#)

## Working with ILM rules and ILM policies

Once you have created ILM rules and an ILM policy, you can continue to work with them,

modifying their configuration as your storage requirements change.

### Deleting an ILM rule

To keep the list of current ILM rules manageable, delete any ILM rules that you are not likely to use.

### What you'll need

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

You cannot delete an ILM rule if it is currently used in the active policy or in the proposed policy. If you need to delete an ILM rule that is used in a policy, you must perform these steps first:



1. Clone the active policy or edit the proposed policy.
2. Remove the ILM rule from the policy.
3. Save, simulate, and activate the new policy to make sure objects are protected as expected.

### Steps

1. Select **ILM > Rules**.
2. Review the table entry for the rule you want to remove.

Confirm that the rule is not used in the active ILM policy or the proposed ILM policy.

3. If the rule you want to remove is not in use, select the radio button and select **Remove**.
4. Select **OK** to confirm that you want to delete the ILM rule.

The ILM rule is deleted.

If you delete a rule that is used in a historical policy, an  icon appears for the rule when you view the policy, which indicates that the rule has become a historical rule.

#### Viewing Historical Policy - Example ILM policy

Review the rules in this policy. If this is a proposed policy, click Simulat

Reason for change: new policy

Rules are evaluated in order, starting from the top

#### Rule Name

Erasure code larger objects

2 copies 2 sites  

This is a historical ILM rule. Historical rules are rules that were included in a policy and then edited or deleted after the policy became historical.



### Related information

[Creating an ILM policy](#)

## Editing an ILM rule

You might need to edit an ILM rule to change a filter or placement instruction.

### What you'll need

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

### About this task

You cannot edit a rule if it is being used in the proposed ILM policy or the active ILM policy. Instead, you can clone these rules and make any required changes to the cloned copy. You also cannot edit the stock ILM rule (Make 2 Copies) or ILM rules created before StorageGRID version 10.3.



Before adding an edited rule to the active ILM policy, be aware that a change to an object's placement instructions might cause an increased load on the system.

### Steps

1. Select **ILM > Rules**.

The ILM Rules page appears. This page shows all available rules and indicates which rules are being used in the active policy or the proposed policy.

#### ILM Rules

Information lifecycle management (ILM) rules determine how and where object data is stored over time. Every object ingested into the StorageGRID Webscale is evaluated against the ILM rules that make up the active ILM policy. Use this page to manage and view ILM rules. You cannot edit or remove an ILM rule that is used by an active or proposed ILM policy.

<input type="button" value="+ Create"/> <input type="button" value="Edit"/> <input type="button" value="Clone"/> <input type="button" value="Remove"/>		
Name	Used In Active Policy	Used In Proposed Policy
<input type="radio"/> Make 2 Copies	✓	✓
<input type="radio"/> PNGs		✓
<input checked="" type="radio"/> JPGs		
<input type="radio"/> X-men		✓

2. Select a rule that is not being used, and click **Edit**.

The Edit ILM Rule wizard opens.

Name:

Description:

Tenant Accounts (optional):

Bucket Name:

[Advanced filtering...](#) (0 defined)

Cancel Next

3. Complete the pages of the Edit ILM Rule wizard, following the steps for creating an ILM rule and using advanced filters, as necessary.

When editing an ILM rule, you cannot change its name.

4. Click **Save**.

If you edit a rule that is used in a historical policy, an  icon appears for the rule when you view the policy, which indicates that the rule has become a historical rule.

### Viewing Historical Policy - Example ILM policy

Review the rules in this policy. If this is a proposed policy, click Simulat

Reason for change: new policy

Rules are evaluated in order, starting from the top

Rule Name
Erasure code larger objects
2 copies 2 sites  

This is a historical ILM rule. Historical rules are rules that were included a policy and then edited or deleted after the policy became historical.



#### Related information

[Creating an ILM rule](#)

[Using advanced filters in ILM rules](#)

#### Cloning an ILM rule

You cannot edit a rule if it is being used in the proposed ILM policy or the active ILM policy. Instead, you can clone a rule and make any required changes to the cloned copy. Then, if required, you can remove the original rule from the proposed policy and replace it with the modified version. You cannot clone an ILM rule if it was created using StorageGRID version 10.2 or earlier.

## What you'll need

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

## About this task

Before adding a cloned rule to the active ILM policy, be aware that a change to an object's placement instructions might cause an increased load on the system.

## Steps

1. Select **ILM > Rules**.

The ILM Rules page appears.

### ILM Rules

Information lifecycle management (ILM) rules determine how and where object data is stored over time. Every object ingested into the StorageGRID Webscale is evaluated against the ILM rules that make up the active ILM policy. Use this page to manage and view ILM rules. You cannot edit or remove an ILM rule that is used by an active or proposed ILM policy.

<input type="button" value="+ Create"/> <input type="button" value="Edit"/> <input type="button" value="Clone"/> <input type="button" value="Remove"/>			
	Name	Used In Active Policy	Used In Proposed Policy
<input type="radio"/>	Make 2 Copies	✓	✓
<input type="radio"/>	PNGs		✓
<input checked="" type="radio"/>	JPGs		
<input type="radio"/>	X-men		✓

2. Select the ILM rule you want to clone, and click **Clone**.

The Create ILM Rule wizard opens.

3. Update the cloned rule by following the steps for editing an ILM rule and using advanced filters.

When cloning an ILM rule, you must enter a new name.

4. Click **Save**.

The new ILM rule is created.

## Related information

[Working with ILM rules and ILM policies](#)

[Using advanced filters in ILM rules](#)

### Viewing the ILM policy activity queue


You can view the number of objects that are in the queue to be evaluated against the ILM policy at any time. You might want to monitor the ILM processing queue to determine system performance. A large queue might indicate that the system is not able to keep up with the ingest rate, the load from the client applications is too great, or that some abnormal condition exists.

## What you'll need

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

## Steps

### 1. Select **Dashboard**.



### 2. Monitor the Information Lifecycle Management (ILM) section.

You can click the question mark  to see a description of the items in this section.


## Managing 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 are not 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).

As shown in the figure, 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, S3 client applications can optionally specify retention settings for any object version in that bucket. An object version must have retention settings specified to be protected by S3 Object Lock.



The StorageGRID S3 Object Lock feature provides a single retention mode that is equivalent to the Amazon S3 compliance mode. By default, a protected object version cannot be overwritten or deleted by any user. The StorageGRID S3 Object Lock feature does not support a governance mode, and it does not allow users with special permissions to bypass retention settings or to delete protected objects.

If a bucket has S3 Object Lock enabled, the S3 client application can optionally specify either or both of the following object-level retention settings when creating or updating an object:

- **Retain-until-date:** If an object version’s retain-until-date is in the future, the object can be retrieved, but it cannot be modified or deleted. As required, an object’s retain-until-date can be increased, but this date cannot be decreased.
- **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.

For details on these settings, go to “using S3 object lock” in [S3 REST API supported operations and limitations](#).

### Comparing S3 Object Lock to legacy Compliance

The S3 Object Lock feature in StorageGRID 11.5 replaces the Compliance feature that was available in previous StorageGRID versions. Because the new S3 Object Lock feature conforms to Amazon S3 requirements, it deprecates the proprietary StorageGRID Compliance feature, which is now referred to as “legacy Compliance.”

If you previously enabled the global Compliance setting, the new global S3 Object Lock setting is enabled automatically when you upgrade to StorageGRID 11.5. Tenant users will no longer be able to create new buckets with Compliance enabled in StorageGRID 11.5; however, as required, tenant users can continue to use and manage any existing legacy Compliant buckets, which includes performing the following tasks:

- Ingesting new objects into an existing bucket that has legacy Compliance enabled.
- Increasing the retention period of an existing bucket that has legacy Compliance enabled.
- Changing the auto-delete setting for an existing bucket that has legacy Compliance enabled.
- Placing a legal hold on an existing bucket that has legacy Compliance enabled.
- Lifting a legal hold.

[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.


	<b>S3 Object Lock (new)</b>	<b>Compliance (legacy)</b>
How is the feature enabled globally?	From the Grid Manager, select <b>Configuration &gt; System Settings &gt; S3 Object Lock</b> .	No longer supported.  <b>Note:</b> If you previously enabled the global Compliance setting, the global S3 Object Lock setting will be enabled automatically when you upgrade to StorageGRID 11.5.
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.	Users can no longer create new buckets with Compliance enabled; however, they can continue to add new objects to existing Compliant buckets.
Is bucket versioning supported?	Yes. Bucket versioning is required and is enabled automatically when S3 Object Lock is enabled for the bucket.	No. The legacy Compliance feature does not allow bucket versioning.
How is object retention set?	Users can set a retain-until-date for each object version.	Users must set a retention period for the entire bucket. The retention period applies to all objects in the bucket.
Can a bucket have default settings for retention and legal hold?	No. StorageGRID buckets that have S3 Object Lock enabled do not have a default retention period. Instead, you can specify a retain-until-date for each object version.	Yes
Can the retention period be changed?	The retain-until-date for an object version can be increased but never decreased.	The bucket's retention period can be increased but never decreased.

	<b>S3 Object Lock (new)</b>	<b>Compliance (legacy)</b>
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?	An object version can be deleted after the 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

### **Workflow for S3 Object Lock**

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.

The workflow diagram shows the high-level steps for using S3 Object Lock. These steps are performed by the grid administrator and by tenant users.



**Grid admin tasks**

As the workflow diagram shows, a grid administrator must perform two high-level tasks before S3 tenant users can use S3 Object Lock:

1. Create at least one compliant ILM rule and make that rule the default rule in the active ILM policy.
2. Enable the global S3 Object Lock setting for the entire StorageGRID system.

### Tenant user tasks

After the global S3 Object Lock setting has been enabled, tenants can perform these tasks:

1. Create buckets that have S3 Object Lock enabled.
2. Add objects to those buckets and specify object-level retention periods and legal hold settings.
3. As required, update a retention period or change the legal hold setting for an individual object.

### Related information

[Use a tenant account](#)

[Use S3](#)

### 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 cannot disable the setting.
- You cannot enable the global S3 Object Lock unless the default rule in the active ILM policy 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 cannot create a new proposed ILM policy or activate an existing proposed 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.

In the following example, the ILM Rules page lists three rules that are compliant with buckets with S3 Object Lock enabled.

Name	Compliant	Used In Active Policy	Used In Proposed Policy
Make 2 Copies	✓	✓	
Compliant Rule: EC for objects in bank-records bucket	✓		
2 copies 10 years, Archive forever			
2 Copies 2 Data Centers	✓		

**Compliant Rule: EC for objects in bank-records bucket**

Description: 2+1 EC at one site

Ingest Behavior: Balanced

**Compliant: Yes**

Tenant Accounts: Bank of ABC (94793396288150002349)

Bucket Name: equals 'bank-records'

Reference Time: Ingest Time

### Requirements for compliant ILM rules

If you want to enable the global S3 Object Lock setting, you must ensure that the default rule in your active ILM policy 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 cannot be saved in a Cloud Storage Pool.
- Object copies cannot be saved on Archive Nodes.
- 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.”

For example, this rule satisfies the requirements of buckets with S3 Object Lock enabled. It stores two replicated object copies from Ingest Time (day 0) to “forever.” The objects will be stored on Storage Nodes at two data centers.

**Compliant rule: 2 replicated copies at 2 sites**

Description: 2 replicated copies on Storage Nodes from Day 0 to Forever

Ingest Behavior: Balanced

Compliant: Yes

Tenant Accounts: Bank of ABC (94793396288150002349)

Reference Time: Ingest Time

Filtering Criteria: Matches all objects.

Retention Diagram:

The diagram shows two horizontal bars representing data centers DC1 and DC2. A vertical dashed line marks 'Day 0'. From Day 0, a blue bar extends to the right for DC1 and an orange bar extends to the right for DC2. Both bars end with a red arrowhead pointing to 'Forever'. A 'Trigger' icon is shown at Day 0 for each data center.

### Requirements for active and proposed ILM policies

When the global S3 Object Lock setting is enabled, active and proposed ILM policies can include both compliant and non-compliant rules.

- The default rule in the active or any proposed ILM policy must be compliant.
- Non-compliant rules only apply to objects in buckets that do not have S3 Object Lock enabled or that do not 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.

A compliant ILM policy might include these three rules:

1. A compliant rule that creates erasure-coded copies of the objects in a specific bucket with S3 Object Lock enabled. The EC copies 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 Archive Nodes and stores that copy forever. This rule only applies to buckets that do not have S3 Object Lock or legacy Compliance enabled because it stores only one object copy forever and it uses Archive Nodes.
3. A default, compliant rule that creates two replicated object copies on Storage Nodes from day 0 to forever. This rule applies to any object in any bucket that was not filtered out by the first two rules.

### 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.

This example from the Tenant Manager shows a bucket with S3 Object Lock enabled.

## Buckets

Create buckets and manage bucket settings.

1 bucket Create bucket

Actions ▾

<input type="checkbox"/>	Name ▾	S3 Object Lock <span>?</span> ▾	Region ▾	Object Count <span>?</span> ▾	Space Used <span>?</span> ▾	Date Created ▾
<input type="checkbox"/>	bank-records	✓	us-east-1	0	0 bytes	2021-01-06 16:53:19 MST

← Previous **1** Next →

- If you plan to use S3 Object Lock, you must enable S3 Object Lock when you create the bucket. You cannot enable S3 Object Lock for an existing bucket.
- Bucket versioning is required with S3 Object Lock. When S3 Object Lock is enabled for a bucket, StorageGRID automatically enables versioning for that bucket.
- After you create a bucket with S3 Object Lock enabled, you cannot disable S3 Object Lock or suspend versioning for that bucket.
- An StorageGRID bucket that has S3 Object Lock enabled does not have a default retention period. Instead, the S3 client application can optionally specify a retention date and legal hold setting for each object version that is added to that bucket.
- Bucket lifecycle configuration is supported for S3 Object Lifecycle buckets.

- CloudMirror replication is not supported for buckets with S3 Object Lock enabled.

### Requirements for objects in buckets with S3 Object Lock enabled

- The S3 client application must specify retention settings for each object that needs to be protected by S3 Object Lock.
- You can increase the retain-until-date for an object version, but you can never decrease this value.
- If you are notified of a pending legal action or regulatory investigation, you can preserve relevant information by placing a legal hold on an object version. When an object version is under a legal hold, that object cannot be deleted from StorageGRID, even if it has reached its retain-until-date. As soon as the legal hold is lifted, the object version can be deleted if the retain-until-date has been reached.
- S3 Object Lock requires the use of versioned buckets. 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 three stages:

#### 1. Object ingest

- When adding an object version to a bucket with S3 Object Lock enabled, the S3 client application can optionally specify retention settings for the object (retain-until-date, legal hold, or both). StorageGRID then generates metadata for that object, which includes a unique object identifier (UUID) and the ingest date and time.
- After an object version with retention settings is ingested, its data and S3 user-defined metadata cannot be modified.
- StorageGRID stores the object metadata independently of the object data. It maintains three copies of all object metadata at each site.

#### 2. Object retention

- Multiple copies of the object are stored by StorageGRID. The exact number and type of copies and the storage locations are determined by the compliant rules in the active ILM policy.

#### 3. Object deletion

- An object can be deleted when its retain-until-date is reached.
- An object that is under a legal hold cannot be deleted.

### Related information

[Use a tenant account](#)

[Use S3](#)

[Comparing S3 Object Lock to legacy Compliance](#)

[Example 7: Compliant ILM policy for S3 Object Lock](#)

[Review audit logs](#)

## Enabling 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.

### What you'll need

- You must have the Root Access permission.
- You must be signed in to the Grid Manager using a supported browser.
- You must have reviewed the S3 Object Lock workflow, and you must understand the considerations.
- The default rule in the active ILM policy must be compliant.

[Creating a default ILM rule](#)

[Creating an ILM policy](#)

### 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 cannot be disabled.



If you enabled the global Compliance setting using a previous version of StorageGRID, the new S3 Object Lock setting is automatically enabled when you upgrade to StorageGRID version 11.5. You can continue to use StorageGRID to manage the settings of existing compliant buckets; however, you can no longer create new compliant buckets.

[NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

### Steps

1. Select **Configuration > System Settings > S3 Object Lock**.

The S3 Object Lock Settings page appears.

#### S3 Object Lock Settings

Enable S3 Object Lock for your entire StorageGRID system if S3 tenant accounts need to satisfy regulatory compliance requirements when saving object data. After this setting is enabled, it cannot be disabled.

#### S3 Object Lock

Before enabling S3 Object Lock, you must ensure that the default rule in the active ILM policy is compliant. A compliant rule satisfies the requirements of buckets with S3 Object Lock 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 cannot be saved on Archive Nodes.
- 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".

Enable S3 Object Lock

Apply

If you had enabled the global Compliance setting using a previous version of StorageGRID, the page includes the following note:

The S3 Object Lock setting replaces the legacy Compliance setting. When this setting is enabled, tenant users can create buckets with S3 Object Lock enabled. Tenants who previously created buckets for the legacy Compliance feature can manage their existing buckets, but can no longer create new buckets with legacy Compliance enabled. See [Managing objects with information lifecycle management](#) for information.

2. Select **Enable S3 Object Lock**.
3. Select **Apply**.


A confirmation dialog box appears and reminds you that you cannot 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 cannot be disabled.
- If the default rule is not compliant, an error appears, indicating that you must create and activate a new ILM policy that includes a compliant rule as its default rule. Select **OK**, and create a new proposed policy, simulate it, and activate it.



### After you finish

After you enable the global S3 Object Lock setting, you might want to create a new ILM policy. After the setting is enabled, the ILM policy can optionally include both a compliant default rule and a non-compliant default rule. For example, you might want to use a non-compliant rule that does not have filters for objects in buckets that do not have S3 Object Lock enabled.

### Related information

[Creating an ILM policy after S3 Object Lock is enabled](#)

[Creating an ILM rule](#)

[Creating an ILM policy](#)

[Comparing S3 Object Lock to legacy Compliance](#)


## Resolving 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 displays an error message. Specifically:

- Tenant Manager users see the following error message:



- Tenant Management API users and S3 API users receive a response code of 503 `Service Unavailable` with similar message text.

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](#)

[Maintain & recover](#)

## Example ILM rules and policies

You can use the examples in this section as a starting point for your own ILM rules and policy.

- [Example 1: ILM rules and policy for object storage](#)
- [Example 2: ILM rules and policy for EC object size filtering](#)
- [Example 3: ILM rules and policy for better protection for image files](#)
- [Example 4: ILM rules and policy for S3 versioned objects](#)
- [Example 5: ILM rules and policy for Strict ingest behavior](#)
- [Example 6: Changing an ILM policy](#)
- [Example 7: Compliant ILM policy for S3 Object Lock](#)

### 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 the proposed policy to confirm it will work as intended to protect content from loss.

#### ILM rule 1 for example 1: Copy object data to two data centers

This example ILM rule copies object data to storage pools in two data centers.

Rule definition	Example value
Storage Pools	Two storage pools, each at different data centers, named Storage Pool DC1 and Storage Pool DC2.
Rule Name	Two Copies Two Data Centers
Reference Time	Ingest Time
Content Placement	On Day 0, keep two replicated copies forever—one in Storage Pool DC1 and one in Storage Pool DC2.

Configure placement instructions to specify how you want objects matched by this rule to be stored.

Two Copies Two Data Centers

Reference Time Ingest Time ▾

**Placements** Sort by start day

From day  store  Add Remove

Type  Location    Copies  + x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

**Retention Diagram** Refresh

The diagram shows a horizontal timeline starting at 'Day 0'. Two bars represent the retention duration for different storage pools. The top bar, labeled 'Storage Pool DC1', is blue and extends to the right with an arrowhead, indicating it lasts forever. The bottom bar, labeled 'Storage Pool DC2', is orange and also extends to the right with an arrowhead, indicating it lasts forever. The x-axis is labeled 'Duration' and 'Forever'.

Cancel Back Next

**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
Erasure Coding Profile	<ul style="list-style-type: none"> <li>• One storage pool across three data centers (All 3 sites)</li> <li>• Use 6+3 erasure-coding scheme</li> </ul>
Rule Name	EC for S3 bucket finance-records
Reference Time	Ingest Time
Content Placement	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.

Configure placement instructions to specify how you want objects matched by this rule to be stored.

**EC for S3 bucket finance-records**

Reference Time Ingest Time ▼

**Placements** ⓘ ↕ Sort by start day

From day  store  Add Remove

Type  Location  Copies  + ✕

**Retention Diagram** ⓘ Refresh

The diagram shows a horizontal timeline starting at 'Day 0'. A grey bar labeled 'All 3 sites (6 plus 3)' spans from the start to a vertical line at 'Day 0'. From 'Day 0', a blue bar with a right-pointing arrow extends to the right, labeled 'Forever'.

Cancel
Back
Next

### ILM policy for example 1

The StorageGRID system allows you to design sophisticated and complex ILM policies; however, in practice, most ILM policies are simple.

A typical ILM policy for a multi-site topology might include ILM rules such as the following:

- At ingest, use 6+3 erasure coding to store all objects belonging to the S3 bucket named `finance-records` across three data centers.
- If an object does not match the first ILM rule, use the policy's default ILM rule, Two Copies Two Data Centers, to store a copy of that object in two data centers, DC1 and DC2.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### 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.

Default	Rule Name	Tenant Account	Actions
	EC for S3 bucket finance-records	Ignore	
<input checked="" type="checkbox"/>	Two Copies Two Data Centers	Ignore	

### 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 the proposed policy to confirm it will work as intended to protect content from loss.

#### ILM rule 1 for example 2: Use EC for all objects larger than 200 KB

This example ILM rule erasure codes all objects larger than 200 KB (0.20 MB).

Rule definition	Example value
Rule Name	EC only objects > 200 KB
Reference Time	Ingest Time
Advanced Filtering for Object Size	Object Size (MB) greater than 0.20
Content Placement	Create a 2+1 erasure-coded copy using three sites

## Advanced Filtering

Use advanced filtering if you want a rule to apply only to specific objects. You can filter objects based on their system metadata, user metadata, or object tags (S3 only). When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the advanced filter.

### EC only objects > 200 KB

Matches all of the following metadata:

Object Size (MB)	greater than	0.2	+ x
+ x			

Cancel Remove Filters Save

The placement instructions specify that a 2+1 erasure-coded copy be created using all three sites.

### EC image files > 200 KB

Reference Time: Ingest Time

Placements Sort by start day

From day: 0 store: forever Add Remove

Type: erasure coded Location: All 3 sites (2 plus 1) Copies: 1 + x

Retention Diagram Refresh

The diagram shows a horizontal timeline starting at 'Day 0' with a trigger icon. A bar labeled 'All 3 sites (2 plus 1)' extends to the right, ending at 'Forever'.

### 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 second rule in the policy. Because ILM rule 1 for example 2 filters out all objects larger than 200 KB, ILM rule 2 for example 2 only applies to objects that are 200 KB or smaller.

Rule definition	Example value
Rule Name	Two Replicated Copies
Reference Time	Ingest Time
Advanced Filtering for Object Size	None

Rule definition	Example value
Content Placement	Create two replicated copies and save them at two data centers, DC1 and DC2

Create ILM Rule Step 2 of 3: Define Placements

Configure placement instructions to specify how you want objects matched by this rule to be stored.

**Two replicated copies**

Reference Time Ingest Time ▼

**Placements** Sort by start day ↑

From day  store forever ▼ Add Remove

Type replicated ▼ Location DC1 x DC2 x Add Pool Copies  + x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

**Retention Diagram** Refresh

The diagram shows a horizontal timeline starting at 'Day 0'. Two bars represent the retention periods for DC1 and DC2. The DC1 bar is blue and the DC2 bar is orange. Both bars start at 'Day 0' and extend to the right, labeled 'Forever' at the end. Above the bars, a 'Trigger' icon is shown at 'Day 0'. Below the bars, the label 'Duration' is shown.

Cancel Back Next

**ILM policy for example 2: Use EC for objects larger than 200 KB**

In this example policy, objects larger than 200 KB are erasure coded. Two replicated copies are made of all other objects.

This example ILM policy includes the following ILM rules:

- Erasure code all objects larger than 200 KB.
- If an object does not match the first ILM rule, use the default ILM rule to create two replicated copies of that object. Because objects larger than 200 KB have been filtered out by rule 1, rule 2 only applies to objects that are 200 KB or smaller.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### 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

Default	Rule Name	Tenant Account	Actions
	EC only objects > 200 KB	Ignore	✘
✓	Two replicated copies	Ignore	✘

Cancel
Save

### 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 larger than 200 KB are erasure coded and that three 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 the proposed policy to confirm it will work as intended to protect content from loss.

#### ILM rule 1 for example 3: Use EC for image files larger than 200 KB

This example ILM rule uses advanced filtering to erasure code all image files larger than 200 KB.

Rule definition	Example value
Rule Name	EC image files > 200 KB
Reference Time	Ingest Time
Advanced Filtering for User Metadata	User Metadata type equals image files
Advanced Filtering for Object Size	Object Size (MB) greater than 0.2
Content Placement	Create a 2+1 erasure-coded copy using three sites

## Advanced Filtering

Use advanced filtering if you want a rule to apply only to specific objects. You can filter objects based on their system metadata, user metadata, or object tags (S3 only). When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the advanced filter.

### EC image files > 200 KB

Matches all of the following metadata:

User Metadata	type	equals	image	+ x
Object Size (MB)		greater than	0.2	+ x

+ x

Cancel Remove Filters Save

Because this rule is configured as the first rule in the policy, the erasure-coding placement instruction only applies to images that are larger than 200 KB.

### EC image files > 200 KB

Reference Time: Ingest Time

Placements Sort by start day

From day: 0 store: forever Add Remove

Type: erasure coded Location: All 3 sites (2 plus 1) Copies: 1 + x

Retention Diagram Refresh

The diagram shows a horizontal bar representing the retention period. The bar starts at 'Day 0' and is labeled 'All 3 sites (2 plus 1)'. The duration of the bar is 'Forever'. A blue arrow points to the right from the end of the bar.

### ILM rule 2 for example 3: Replicate 3 copies for all remaining image files

This example ILM rule uses advanced filtering to specify that image files be replicated.

Rule definition	Example value
Rule Name	3 copies for image files
Reference Time	Ingest Time

Rule definition	Example value
Advanced Filtering for User Metadata	User Metadata type equals image files
Content Placement	Create 3 replicated copies in all Storage Nodes

## Advanced Filtering

Use advanced filtering if you want a rule to apply only to specific objects. You can filter objects based on their system metadata, user metadata, or object tags (S3 only). When objects are evaluated, the rule is applied if the object's metadata matches the criteria in the advanced filter.

### 3 copies for image files

Matches all of the following metadata:

User Metadata	▼	type	equals	▼	image	+ x
+ x						

Cancel Remove Filters Save

Because the first rule in the policy has already matched image files larger than 200 KB, these placement instructions only apply to image files 200 KB or smaller.

### 3 copies for image files

Reference Time

Sort by start day

#### Placements


From day  store  Add Remove

Type  Location  Copies  + x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

#### Retention Diagram

Refresh



Cancel Back Next

### ILM policy for example 3: Better protection for image files

In this example, the ILM policy uses three ILM rules to create a policy that erasure codes image files larger than 200 KB (0.2 MB), creates replicated copies for image files 200 KB or smaller, and makes two replicated copies for any non-image files.

This example ILM policy includes rules that perform the following:

- Erasure code all image files larger than 200 KB.
- Create three copies of any remaining image files (that is, images that are 200 KB or smaller).
- Apply the default rule to any remaining objects (that is, all non-image files).

#### Viewing Active Policy - Better protection for image files

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

**Reason for change:** ILM policy for example 3

*Rules are evaluated in order, starting from the top.*

Rule Name	Default	Tenant Account
EC only objects > 200 KB		Ignore
3 copies for image files		Ignore
Make 2 Copies	✓	Ignore

[Simulate](#) [Activate](#)

### 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.

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 the proposed policy to confirm it will work as intended to protect content from loss.



If you create ILM policies to manage noncurrent object versions, be aware that you must know the object version's UUID or CBID to simulate the policy. To find an object's UUID and CBID, use Object Metadata Lookup while the object is still current.

#### Related information

[How S3 versioned objects are deleted](#)

[Verifying an ILM policy with object metadata lookup](#)

#### ILM rule 1 for example 4: Save three copies for 10 years

This example ILM rule stores a copy of each object at three data centers 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 at different data centers, named DC1, DC2, and DC3.
Rule Name	Three Copies Ten Years
Reference Time	Ingest Time
Content Placement	On Day 0, keep three replicated copies for 10 years (3,652 days), one in DC1, one in DC2, and one in DC3. At the end of 10 years, delete all copies of the object.

### Create ILM Rule Step 2 of 3: Define Placements

Configure placement instructions to specify how you want objects matched by this rule to be stored.

**Three Copies Ten Years**  
 Save three copies for ten years

Reference Time:

**Placements** Sort by start day

From day  store   days Add Remove

Type:  Location:  Copies:  + x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

---

**Retention Diagram** Refresh

Trigger: Day 0, Day 3652

Duration: 3652 days, Forever

Cancel Back Next

### 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. This rule uses the **Noncurrent Time** option for Reference Time.

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, named DC1 and DC2.
Rule Name	Noncurrent Versions: Two Copies Two Years
Reference Time	Noncurrent Time
Content Placement	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 DC1 and one in DC2. At the end of 2 years, delete the noncurrent versions.

**Noncurrent Versions: Two Copies Two Years**  
 Save two copies of noncurrent versions for two years

Reference Time: Noncurrent Time

**Placements** Sort by start day

From day  store for  days Add Remove

Type replicated Location DC1 x DC2 x Add Pool Copies  + x

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

**Retention Diagram** Refresh

#### 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 data centers. Keep copies of the current object version for 10 years.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### 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.

Default	Rule Name	Tenant Account	Actions
	Noncurrent Versions: Two Copies Two Years	Ignore	✘
✓	Three Copies Ten Years	Ignore	✘

The default ILM rule in this policy does not retain objects forever. Confirm this is the behavior you expect. Otherwise, any objects that are not matched by another rule will be deleted after 3652 days.

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).



To simulate noncurrent object versions, you must use that version's UUID or CBID. While the object is still current, you can use Object Metadata Lookup to find its UUID and CBID.

- 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 DC3 is purged and the two copies at DC1 and DC2 are stored for 2 more years.

### Related information

[Verifying an ILM policy with object metadata lookup](#)

### 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 the proposed policy to confirm it will work as intended to protect content from loss.

## Related information

[How objects are ingested](#)

[Step 3 of 3: Define 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 Filtering	Location Constraint equals eu-west-3
Storage Pools	DC1 (Paris)
Rule Name	Strict ingest to guarantee Paris data center
Reference Time	Ingest Time
Content Placement	On Day 0, keep two replicated copies forever in DC1 (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

**Description:** Strict ingest to guarantee Paris data center  
**Ingest Behavior:** Strict  
**Tenant Account:** Paris tenant (25580610012441844135)  
**Reference Time:** Ingest Time  
**Filtering Criteria:**

Matches all of the following metadata:


System Metadata

Location Constraint (S3 only)

equals

eu-west-3

### Retention Diagram:



### 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 cannot 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 Filtering	<i>Not specified</i>
Storage Pools	DC1 (Paris) and DC2 (US)
Rule Name	2 Copies 2 Data Centers
Reference Time	Ingest Time
Content Placement	On Day 0, keep two replicated copies forever at two data centers
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.

## 2 Copies 2 Data Centers

Description: 2 Copies 2 Data Centers


Ingest Behavior: Balanced

Reference Time: Ingest Time

Filtering Criteria:

Matches all objects.

Retention Diagram:



### 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 cannot be satisfied.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### 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				
Default	Rule Name	Tenant Account	Actions	
	Strict ingest to guarantee Paris data center	Paris tenant (25580610012441844135)	✘	
✓	2 Copies 2 Data Centers	Ignore	✘	

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 are not available, ingest fails.
- All other objects are matched by the second rule, including objects that belong to the Paris tenant and that do not 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: Changing an ILM policy

You might need to create and activate a new ILM policy if your data protection needs change or you add new sites.

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 the active ILM policy needs to be revised to store data at the new site.



The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Before activating a new policy, simulate the proposed policy to confirm it will work as intended to protect content from loss.

### How does changing an ILM policy affect 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.

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.

## Related information

[Creating an ILM policy](#)

### 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.

#### ILM Policies

Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.

<a href="#">+ Create Proposed Policy</a> <a href="#">Clone</a> <a href="#">Edit</a> <a href="#">Remove</a>			
Policy Name	Policy State	Start Date	End Date
<input checked="" type="radio"/> Data Protection for Two Sites	Active	2020-06-10 16:42:09 MDT	
<input type="radio"/> Baseline 2 Copies Policy	Historical	2020-06-09 21:48:34 MDT	2020-06-10 16:42:09 MDT

#### Viewing Active Policy - Data Protection for Two Sites

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: Data Protection for Two Sites

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
One-Site Erasure Coding for Tenant A <a href="#">🔗</a>		Tenant A (49752734300032812036)
Two-Site Replication for Other Tenants <a href="#">🔗</a>	✓	Ignore

[Simulate](#) [Activate](#)

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.



The first rule in this example uses an advanced filter to ensure that erasure coding is not used for small objects. Any of Tenant A's objects that are smaller than 200 KB will be protected by the second rule, which uses replication.

### Rule 1: One-site erasure coding for Tenant A

Rule definition	Example value
Rule Name	One-Site Erasure Coding for Tenant A
Tenant Account	Tenant A
Storage Pool	Data Center 1
Content Placement	2+1 erasure coding in Data Center 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	Data Center 1 and Data Center 2
Content Placement	Two replicated copies from day 0 to forever: one copy at Data Center 1 and one copy at Data Center 2.

### Proposed ILM policy for example 6: Data protection at three sites

In this example, the ILM policy is being updated 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 Data Center 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 proposed ILM policy, which is designed to protect data at all three sites.

Before activating a new ILM policy:

- Review and carefully simulate the policy. Errors in an ILM policy can cause irreparable data loss.
- 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.

See [Managing objects with information lifecycle management](#) for more information.

This policy contains a rule that makes an erasure-coded copy. Confirm that at least one rule uses the Object Size advanced filter to prevent objects that are 200 KB or smaller from being erasure coded. See [Managing objects with information lifecycle management](#) for more information.

Review the rules in this policy. If this is a proposed policy, click Simulate to verify the policy and then click Activate to make the policy active.

Reason for change: Data Protection for Three Sites

Rules are evaluated in order, starting from the top.

Rule Name	Default	Tenant Account
Three-Site Erasure Coding for Tenant A 		Tenant A (49752734300032812036)
Three-Site Replication for Other Tenants 	✓	Ignore

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 Data Centers (includes Data Center 1, Data Center 2, and Data Center 3)
Content Placement	2+1 erasure coding in All 3 Data Centers 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	Data Center 1, Data Center 2, and Data Center 3
Content Placement	Three replicated copies from day 0 to forever: one copy at Data Center 1, one copy at Data Center 2, and one copy at Data Center 3.

## Activating the proposed ILM policy for example 6

When you activate a new proposed 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 Data Center 1. In the new proposed ILM policy, objects belonging to Tenant A will be protected using 2+1 erasure coding at Data Centers 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 data center.
- 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 data centers.



The existing 2+1 fragments at Data Center 1 are not 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 other tenants are protected using two replicated copies in storage pools at Data Centers 1 and 2. In the new proposed ILM policy, objects belonging to other tenants will be protected using three replicated copies in storage pools at Data Centers 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 data center.
- Existing objects belonging to these other tenants are re-evaluated during the ongoing ILM scanning process. Because the existing object copies at Data Center 1 and Data Center 2 continue to satisfy the replication requirements of the new ILM rule, StorageGRID only needs to create one new copy of the object for Data Center 3.

## Performance impact of activating this policy

When the proposed 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 Data Center 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 an 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 are not 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 the proposed policy to confirm it will work as intended to protect content from loss.

### Related information

[Managing objects with S3 Object Lock](#)

[Creating 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)

# Buckets

Create buckets and manage bucket settings.

1 bucket

Create bucket

Actions ▾

<input type="checkbox"/>	Name	S3 Object Lock	Region	Object Count	Space Used	Date Created
<input type="checkbox"/>	bank-records	✓	us-east-1	0	0 bytes	2021-01-06 16:53:19 MST

← Previous 1 Next →

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 cannot 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: an erasure-coded 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	bank-records

Rule definition	Example value
Advanced filtering	Object Size (MB) greater than 0.20  <b>Note:</b> This filter ensures that erasure coding is not used for objects 200 KB or smaller.

Create ILM Rule Step 1 of 3: Define Basics

Name

Description

Tenant Accounts (optional)

Bucket Name

[Advanced filtering...](#) (0 defined)

Cancel Next

Rule definition	Example value
Reference Time	Ingest Time
Placements	From day 0 store forever
Erasure Coding Profile	<ul style="list-style-type: none"> <li>• Create an erasure-coded copy on Storage Nodes at three data center sites</li> <li>• Uses 6+3 erasure-coding scheme</li> </ul>

Configure placement instructions to specify how you want objects matched by this rule to be stored.

Compliant Rule: EC objects in bank-record bucket - Bank of ABC

Reference Time Ingest Time

**Placements** Sort by start day

From day  store  Add Remove

Type  Location  Copies  + x

**Retention Diagram** Refresh

The diagram shows a horizontal timeline starting at 'Day 0'. A grey bar labeled 'Three Data Centers (6 plus 3)' extends to the right. A blue arrow labeled 'Forever' starts at the end of the grey bar and points further to the right. A vertical line marks 'Day 0' at the start of the grey bar.

Cancel Back Save

### 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 do not have S3 Object Lock (or the legacy Compliance feature) enabled.
Advanced filtering	Not specified

Name

Description

Tenant Accounts (optional)

Bucket Name  Value

Advanced filtering... (0 defined)

Cancel Next

Rule definition	Example value
Reference Time	Ingest Time
Placements	<ul style="list-style-type: none"> <li>• On Day 0, keep two replicated copies on Storage Nodes in Data Center 1 and Data Center 2 for 365 days</li> <li>• After 1 year, keep one replicated copy in a Cloud Storage Pool forever</li> </ul>

### 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 and it 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
Bucket Name	Not specified
Advanced filtering	Not specified

### Create ILM Rule Step 1 of 3: Define Basics

Name

Description

Tenant Accounts (optional)

Bucket Name

[Advanced filtering...](#) (0 defined)

Cancel Next

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 Rule: Two Copies Two Data Centers**

Reference Time:

**Placements** Sort by start day

From day:  store:

Type:  Location:  Copies:

Specifying multiple storage pools might cause data to be stored at the same site if the pools overlap. See [Managing objects with information lifecycle management](#) for more information.

**Retention Diagram** Refresh

The diagram shows a timeline starting at Day 0. Two horizontal bars represent the retention periods for Data Center 1 and Data Center 2. Both bars start at Day 0 and extend to the right, labeled 'Forever'. Data Center 1 is represented by a blue bar, and Data Center 2 is represented by an orange bar. A vertical line marks Day 0, and a vertical line marks the end of the retention period, labeled 'Forever'.

### 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 proposed policy.

### Adding 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 larger than 200 KB 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.

## Configure ILM Policy

Create a proposed policy by selecting and arranging rules. Then, save the policy and edit it later as required. Click Simulate to verify a saved policy using test objects. When you are ready, click Activate to make this policy the active ILM policy for the grid.

Name

Reason for change

### 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 (and any non-compliant rule without a filter) will be automatically placed at the end of the policy and cannot be moved.

+ Select Rules

Default	Rule Name	Compliant	Tenant Account	Actions
	Compliant Rule: EC for bank-records bucket - Bank of ABC	✓	Bank of ABC (90767802913525281639)	✕
	Non-Compliant Rule: Use Cloud Storage Pool		Ignore	✕
✓	Default Compliant Rule: Two Copies Two Data Centers	✓	Ignore	✕

Cancel

Save

## Simulating the proposed policy

After you have added rules in your proposed 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 larger than 200 KB 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 200 KB 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.

## Activating the policy

When you are completely satisfied that the new policy protects object data as expected, you can activate it.

## System hardening

Learn the system settings, best practices, and recommendations for protecting a StorageGRID system from security threats.

- [Hardening a StorageGRID system](#)
- [Hardening guidelines for software upgrades](#)
- [Hardening guidelines for StorageGRID networks](#)

- [Hardening guidelines for StorageGRID nodes](#)
- [Hardening guidelines for server certificates](#)
- [Other hardening guidelines](#)]

## Hardening a StorageGRID system

System hardening is the process of eliminating as many security risks as possible from a StorageGRID system.

This document provides an overview of the hardening guidelines that are specific to StorageGRID. These guidelines are a supplement to industry-standard best practices for system hardening. For example, these guidelines assume that you use strong passwords for StorageGRID, use HTTPS instead of HTTP, and enable certificate-based authentication where available.

As you install and configure StorageGRID, you can use these guidelines to help you meet any prescribed security objectives for information system confidentiality, integrity, and availability.

StorageGRID follows the *NetApp Vulnerability Handling Policy*. Reported vulnerabilities are verified and addressed according to the product security incident response process.

### General considerations for hardening a StorageGRID system

When hardening a StorageGRID system, you must 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 Docker container 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 are followed by your organization. You might need to comply with specific regulatory or corporate requirements.

### Related information

[Vulnerability Handling Policy](#)

## 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 are not included in earlier releases.

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 the StorageGRID software download page. For step-by-step instructions for upgrading StorageGRID software, see the instructions for upgrading StorageGRID. For instructions on applying a hotfix, see the recovery and maintenance instructions.

### **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.

Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

### **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.

Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

### **Upgrade 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.

Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

### **Related information**

[NetApp Downloads: StorageGRID](#)

[Upgrade software](#)

[Maintain & recover](#)

[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.

### **Guidelines for the 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

Networks 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 the 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 in the installation guide for your platform.
- If untrusted clients can access the Admin Network, block access to StorageGRID on the Admin Network with an external firewall.

### **Guidelines for the 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 in the installation guide for your platform.
- Accept inbound client traffic only on explicitly configured endpoints. See the information about managing untrusted Client Networks in the instructions for administering StorageGRID.

### **Related information**

[Network guidelines](#)

[Grid primer](#)

[Administer StorageGRID](#)

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

[Install Ubuntu or Debian](#)

[Install VMware](#)

### **Hardening guidelines for StorageGRID nodes**

StorageGRID nodes can be deployed on VMware virtual machines, within a Docker container 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.

## 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.

Nodes running on VMware platforms and StorageGRID appliances use an internal firewall that is managed automatically. 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 the installation guide for your platform.

## Virtualization, containers, and shared hardware

For all StorageGRID nodes, avoid running StorageGRID on the same physical hardware as untrusted software. Do not 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.

## Disable unused services

For all StorageGRID nodes, you should disable or block access to unused services. For example, if you are not planning to configure client access to the audit shares for CIFS or NFS, block or disable access to these services.

## Protect nodes during installation

Do not allow untrusted users to access StorageGRID nodes over the network when the nodes are being installed. Nodes are not 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, do not 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 in the instructions for administering StorageGRID.
- 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.

- Do not enable outbound services for untrusted tenants. For example, when creating the account for an untrusted tenant, do not allow the tenant to use its own identity source and do not allow the use of platform services. See the steps for creating a tenant account in the instructions for administering StorageGRID.
- 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 in the instructions for administering StorageGRID.
- Optionally, connect to external services using the Client Network. Then, select **Configuration > Network Settings > Untrusted Client Network** 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 instead of using the CLB service on Gateway Nodes. See the steps for managing load balancing in the instructions for administering StorageGRID.



The CLB service is deprecated.

- 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 > Network Settings > Untrusted Client Network** 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.
- If the storage controller in the appliance includes FDE or FIPS drives and the Drive Security feature is enabled, use SANtricity to configure Drive Security keys.
- For appliances without FDE or FIPS drives, enable node encryption using a Key Management Server (KMS).

See the installation and maintenance instructions for your StorageGRID hardware appliance.

### Related information

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

[Install Ubuntu or Debian](#)

[Install VMware](#)

[Administer StorageGRID](#)

[Use a tenant account](#)

[SG100 & SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

## Hardening guidelines for server 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 Server Certificate:** Used to secure access to the Grid Manager, the Tenant Manager, the Grid Management API, and the Tenant Management API.
- **Object Storage API Service Endpoints Server Certificate:** Used to secure access to Storage Nodes and Gateway Nodes, which S3 and Swift client applications use to upload and download object data.



StorageGRID manages the certificates used for load balancer endpoints separately. To configure load balancer certificates, see the steps for configuring load balancer endpoints in the instructions for administering StorageGRID.

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 are not 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 do not 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.

## 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.

### 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.

See the instructions for monitoring and troubleshooting for more information about StorageGRID logs. See the instructions for 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 messages and details to NetApp technical support, your organization’s internal support team, or a support partner. By default, AutoSupport messages to NetApp technical support are 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 messages, NetApp strongly recommends using HTTPS as the default transport protocol for sending AutoSupport messages to NetApp support.

Optionally, you can configure an Admin proxy for more control over AutoSupport communication from Admin Nodes to NetApp technical support. See the steps for creating an Admin proxy in the instructions for administering StorageGRID.

### 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, do not 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) in the instructions for using tenant accounts.

## 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.

## Related information

[Monitor & troubleshoot](#)

[Review audit logs](#)

[Use a tenant account](#)

[Administer StorageGRID](#)

# Configure StorageGRID for FabricPool

Learn how to configure StorageGRID as a NetApp FabricPool cloud tier.


- [Configuring StorageGRID for FabricPool](#)
- [Information needed to attach StorageGRID as a cloud tier](#)
- [Using StorageGRID information lifecycle management with FabricPool data](#)
- [Creating a traffic classification policy for FabricPool](#)
- [Other best practices for StorageGRID and FabricPool](#)

## Configuring StorageGRID for FabricPool

If you use NetApp ONTAP software, you can use NetApp FabricPool to tier inactive, or cold, data to a NetApp StorageGRID object storage system.

Use these instructions to:

- Get an overview of configuring a StorageGRID object storage system for use with FabricPool.
- Learn how to obtain the information you provide to ONTAP when you attach StorageGRID as a FabricPool cloud tier.
- Learn the best practices for configuring the StorageGRID information lifecycle management (ILM) policy, a StorageGRID traffic classification policy, and other StorageGRID options for a FabricPool workload.



### What you'll need

Before using these instructions:

- 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.

## Related information

- [TR-4598: FabricPool Best Practices for ONTAP 9.8](#)
- [ONTAP 9 Documentation Center](#)

## What FabricPool is

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. Data in a FabricPool is stored in a tier based on whether it is frequently accessed or not. Using a FabricPool helps you reduce storage cost without compromising performance, efficiency, or protection.

No architectural changes are required, and you can continue managing your database and application environment from the central ONTAP storage system.

## What object storage is

Object storage 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 are not 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.

## Using StorageGRID as a FabricPool cloud tier

FabricPool can tier ONTAP data to a number of object store 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.

## Using multiple ONTAP clusters with StorageGRID

These instructions describe how to connect StorageGRID to a single ONTAP cluster. However, you might want to connect the same StorageGRID system to multiple ONTAP clusters.

The only requirement for tiering data from multiple ONTAP clusters to a single StorageGRID system is that you must use a different S3 bucket for each cluster. Based on your requirements, you can use the same high availability (HA) group, load balancer endpoint, and tenant account for all clusters, or you can configure each of these items for each cluster.

## Information needed to attach StorageGRID as a cloud tier

Before you can attach StorageGRID as an cloud tier for FabricPool, you must perform some configuration steps in StorageGRID and obtain certain values.

### About this task

The following table lists the information you must provide to ONTAP when you attach StorageGRID as a cloud tier for FabricPool. The topics in this section explain how to use the StorageGRID Grid Manager and Tenant Manager to obtain the information you need.



The exact field names listed and the process you use to enter the required values in ONTAP depend on whether you are using the ONTAP CLI (storage aggregate object-store config create) or ONTAP System Manager (**Storage > Aggregates & Disks > Cloud Tier**).

For more information, refer to the following:

- [TR-4598: FabricPool Best Practices for ONTAP 9.8](#)
- [ONTAP 9 Documentation Center](#)

ONTAP field	Description
Object store name	Any unique and descriptive name. For example, <code>StorageGRID_Cloud_Tier</code> .
Provider type	StorageGRID (System Manager) or SGWS (CLI).
Port	The port that FabricPool will use when it connects to StorageGRID. You determine which port number to use when you define the StorageGRID load balancer endpoint.  <a href="#">Creating a load balancer endpoint for FabricPool</a>
Server name	The fully qualified domain name (FQDN) for the StorageGRID load balancer endpoint. For example, <code>s3.storagegrid.company.com</code> .  Note the following: <ul style="list-style-type: none"><li>• The domain name that you specify here must match the domain name on the CA certificate you upload for the StorageGRID load balancer endpoint.</li><li>• The DNS record for this domain name must map to each IP address you will use to connect to StorageGRID.</li></ul> <a href="#">Configuring the DNS server for StorageGRID IP addresses</a>
Container name	The name of the StorageGRID bucket you will use with this ONTAP cluster. For example, <code>fabricpool-bucket</code> . You create this bucket in the Tenant Manager.  Note the following: <ul style="list-style-type: none"><li>• The bucket name cannot be changed once the configuration is created.</li><li>• The bucket cannot have versioning enabled.</li><li>• You must use a different bucket for each ONTAP cluster that will tier data to StorageGRID.</li></ul> <a href="#">Creating an S3 bucket and obtaining an access key</a>

ONTAP field	Description
Access key and secret password	<p>The access key and secret access key for the StorageGRID tenant account.</p> <p>You generate these values in the Tenant Manager.</p> <p><a href="#">Creating an S3 bucket and obtaining an access key</a></p>
SSL	Must be enabled.
Object store certificate	<p>The CA certificate you uploaded when you created the StorageGRID load balancer endpoint.</p> <p><b>Note:</b> 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.</p> <p><a href="#">Creating a load balancer endpoint for FabricPool</a></p>

### After you finish

After you have obtained the required StorageGRID information, you can go to ONTAP to add StorageGRID as a cloud tier, add the cloud tier as an aggregate, and set volume tiering policies.

### Best practices for load balancing

Before attaching StorageGRID as a FabricPool cloud tier, you use the StorageGRID Grid Manager to configure at least one load balancer endpoint.

#### What load balancing is

When data is tiered from FabricPool to 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 FabricPool 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 the following technical report:

#### [StorageGRID Load Balancer Options](#)



The separate Connection Load Balancer (CLB) service on Gateway Nodes is deprecated and no longer recommended for use with FabricPool.

## Best practices for StorageGRID load balancing

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 both an Admin Node and a Gateway Node or even two Admin Nodes. Make sure that there is adequate networking, hardware, or virtualization infrastructure for each load-balancing node, whether you are using SG100 or SG1000 services appliances, bare metal nodes, or virtual machine (VM) based nodes.

You must configure a StorageGRID load balancer endpoint to define the port that Gateway Nodes and Admin Nodes will use for incoming and outgoing FabricPool requests.

### Best practices for the load balancer endpoint certificate

When creating a load balancer endpoint for use with FabricPool, you must use HTTPS as the protocol. You can then either upload a certificate that is signed by either a publicly trusted or a private Certificate Authority (CA), or you can generate a self-signed certificate. The certificate allows ONTAP to authenticate with StorageGRID.

As a best practice, you should use a CA server certificate to secure the connection. Certificates signed by a CA can be rotated nondisruptively.

When requesting a CA certificate for use with the load balancer endpoint, ensure that the domain name on the certificate matches the server name you enter in ONTAP for that load balancer endpoint. If possible, use a wildcard (\*) to allow for virtual-host-style URLs. For example:

```
*.s3.storagegrid.company.com
```

When you add StorageGRID as a FabricPool cloud tier, you must install the same certificate to the ONTAP cluster, as well as the root and any subordinate certificate authority (CA) certificates.



StorageGRID uses server certificates for a number of purposes. If you are connecting to the Load Balancer service, you do not need to upload the Object Storage API Service Endpoints Server Certificate.

To learn more about the server certificate for a load balancing endpoint:

- [Managing load balancing](#)
- [Hardening guidelines for server certificates](#)

## Best practices for high availability groups

Before attaching StorageGRID as a FabricPool cloud tier, you use the StorageGRID Grid Manager to configure a high availability (HA) group.

### What a high availability (HA) group is

To ensure that the Load Balancer service is always available to manage FabricPool data, you can group the network interfaces of multiple Admin and Gateway Nodes into a single entity, known as a high availability (HA) group. If the active node in the HA group fails, another node in the group can continue to manage the workload.

Each HA group provides highly available access to the shared services on the associated nodes. For example,

an HA group consisting of all Admin Nodes provides highly available access to some Admin Node management services and to the Load Balancer service. An HA group that consists of only Gateway Nodes or of both Admin Nodes and Gateway Nodes provides highly available access to the shared Load Balancer service.

When creating an HA group, you select network interfaces belonging to the Grid Network (eth0) or the Client Network (eth2). All interfaces in an HA group must be within the same network subnet.

An HA group maintains one or more virtual IP addresses that are added to the active interface in the group. If the active interface becomes unavailable, the virtual IP addresses are moved to another interface. This 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.

If you configure an HA group of load-balancing nodes, FabricPool connects to the virtual IP addresses of that HA group.

### **Best practices for high availability (HA) groups**

The best practices for creating a StorageGRID HA group for FabricPool depend on the workload, as follows:

- If you plan to use FabricPool with primary workload data, you must create a 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](#).

### **Configuring the DNS server for StorageGRID IP addresses**

After configuring high availability groups and load balancer endpoints, you must ensure that the domain name system (DNS) for the ONTAP system includes a record to associate the StorageGRID server name (fully qualified domain name) to the IP address that FabricPool 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 a HA group, FabricPool will connect to the virtual IP addresses of that HA group.
- If you are not using a HA group, FabricPool can connect to the StorageGRID Load Balancer service using the IP address of any Gateway Node or Admin Node.

You must also ensure that the DNS record references all required endpoint domain names, including any wildcard names.

### **Creating 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 consists of one or more network interfaces on Admin Nodes, Gateway Nodes, or both.

## What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

## About this task

Each HA group uses virtual IP addresses (VIPs) to provide highly available access to the shared services on the associated nodes.

For details about this task, see [Managing high availability groups](#).

## Steps

1. Select **Configuration > Network Settings > High Availability Groups**.
2. Select one or more of the network interfaces. The network interfaces must belong to the same subnet on either the Grid Network (eth0) or the Client Network (eth2).
3. Assign one node to be the Preferred Master.

The preferred Master is the active interface unless a failure occurs that causes the VIP addresses to be reassigned to a Backup interface.

4. Enter up to ten IPv4 addresses for the HA group.

The addresses must be within the IPv4 subnet shared by all of the member interfaces.

## Create High Availability Group

### High Availability Group

Name

Description

### Interfaces

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

Select Interfaces

Node Name	Interface	IPv4 Subnet	Preferred Master
DC1-ADM1	eth0	10.96.98.0/23	<input checked="" type="radio"/>
DC1-G1	eth0	10.96.98.0/23	<input type="radio"/>

Displaying 2 interfaces.

### Virtual IP Addresses

Virtual IP Subnet: 10.96.98.0/23. All virtual IP addresses must be within this subnet. There must be at least 1 and no more than 10 virtual IP addresses.

Virtual IP Address 1

+

Cancel

Save

## Creating a load balancer endpoint for FabricPool

When configuring StorageGRID for use with FabricPool, you configure a load balancer endpoint and upload the load balancer endpoint certificate, which is used to secure the connection between ONTAP and StorageGRID.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.
- You have the following files:
  - Server Certificate: The custom server certificate file.
  - Server Certificate Private Key: The custom server certificate private key file.
  - CA Bundle: A single 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.

### About this task

For details about this task, see [Configuring load balancer endpoints](#).

### Steps

1. Select **Configuration > Network Settings > Load Balancer Endpoints**.

#### Create Endpoint

Display Name

Port

Protocol  HTTP  HTTPS

Endpoint Binding Mode  Global  HA Group VIPs  Node Interfaces

2. Select **Add endpoint**.
3. Enter the following information.

Field	Description
Display name	A descriptive name for the endpoint
Port	<p>The StorageGRID port you want to use for load balancing. This field defaults to 10433, but you can enter any unused external port. If you enter 80 or 443, the endpoint is configured only on Gateway Nodes, since these ports are reserved on Admin Nodes.</p> <p><b>Note:</b> Ports used by other grid services are not permitted. See the list of ports used for internal and external communications:</p> <p><a href="#">Network port reference</a></p> <p>You must provide this same port number to ONTAP when you attach StorageGRID as a FabricPool cloud tier.</p>
Protocol	Must be <b>HTTPS</b> .

Field	Description
Endpoint Binding Mode	<p>Use the <b>Global</b> setting (recommended) or restrict the accessibility of this endpoint to one of the following:</p> <ul style="list-style-type: none"> <li>• Specific high availability (HA) virtual IP addresses (VIPs). Use this selection only if you require much higher levels of isolation of workloads.</li> <li>• Specific network interfaces of specific nodes.</li> </ul>

4. Select **Save**.

The Edit Endpoint dialog box appears.

5. For **Endpoint Service Type**, select **S3**.

6. Select **Upload Certificate** (recommended) and then browse to your server certificate, certificate private key, and CA bundle.

**Load Certificate**

Upload the PEM-encoded custom certificate, private key, and CA bundle files.

Server Certificate

Certificate Private Key

CA Bundle

7. Select **Save**.

### Creating a tenant account for FabricPool

You must create a tenant account in the Grid Manager for FabricPool use.

#### What you'll need

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

#### About this task

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.

You can use the same tenant account for multiple ONTAP clusters. Or, you can create a dedicated tenant account for each ONTAP cluster as required.



These instructions assume that you have configured single sign-on (SSO) for the Grid Manager. If you are not using SSO, use the instructions for [creating a tenant account if StorageGRID is not using SSO](#).

## Steps

1. Select **Tenants**.
2. Select **Create**.
3. Enter a display name for the FabricPool tenant account.
4. Select **S3**.
5. Leave the **Allow Platform Services** check box selected to enable the use of platform services.

If platform services are enabled, a tenant can use features, such as CloudMirror replication, that access external services.

6. Leave the **Storage Quota** field blank.
7. In the **Root Access Group** field, select an existing federated group from the Grid Manager to have the initial Root Access permission for the tenant.
8. Select **Save**.

## Creating an S3 bucket and obtaining an access key

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.

### What you'll need

- You must have created a tenant account for FabricPool use.

### About this task

These instructions describe how to use the StorageGRID Tenant Manager to create a bucket and obtain access keys. You can also perform these tasks using the Tenant Management API or the StorageGRID S3 REST API.

To learn more:

- [Use a tenant account](#)
- [Use S3](#)

## 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.

- a. Select **STORAGE (S3) > Buckets**.
- b. Select **Create bucket**.
- c. Enter the name of the StorageGRID bucket you will use with FabricPool. For example, `fabricpool-bucket`.



You cannot change the bucket name after creating the bucket.

Bucket names must comply 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.
  - 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.
- d. Select the region for this bucket.

By default, all buckets are created in the `us-east-1` region.

## Create bucket ✕

### Enter bucket details

Enter the bucket's name and select the bucket's region.

Bucket name

Region

us-east-1▼

Cancel Create bucket

- e. Select **Create 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 create a new access key and secret access key in the future, remember to update the corresponding values in ONTAP immediately to ensure that ONTAP can store and retrieve data in StorageGRID without interruption.

## Using StorageGRID information lifecycle management with FabricPool data

If you are using FabricPool to tier data to StorageGRID, you must understand the requirements for creating StorageGRID information lifecycle management (ILM) rules and an ILM policy to manage FabricPool data. You must ensure the ILM rules that apply to FabricPool data are not disruptive.



FabricPool has no knowledge of StorageGRID ILM rules or policies. Data loss can occur if the StorageGRID ILM policy is misconfigured.

To learn more: [Manage objects with ILM](#)

### ILM guidelines for FabricPool data

Review these guidelines to ensure that your ILM rules and ILM policy are suitable for FabricPool data and your business requirements. If you are already using StorageGRID ILM, you might need to update your active ILM policy 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 significantly less storage capacity, than replication. The 4+1 and 6+1 schemes use less capacity than 2+1, but at the cost of lower throughput and less flexibility when you add Storage Nodes during grid expansion.

- 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.

- 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 are not deleted by StorageGRID ILM.
- Do not create rules that will move FabricPool cloud tier data out of the bucket to another location. You cannot use ILM rules to archive FabricPool data to tape using an Archive Node or use a Cloud Storage Pool to move FabricPool data to Glacier.



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.

### Example ILM policy for FabricPool data

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 the proposed policy to confirm it will work as intended to protect content from loss.

To learn more: [Manage objects with ILM](#)

### Steps

- Create a storage pool named **DEN**. Select the Denver site.
- Create an Erasure Coding profile named **2 plus 1**. Select the 2+1 erasure-coding scheme and the **DEN** storage pool.
- 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 plus 1 erasure coding for FabricPool data
Bucket Name	<code>fabricpool-bucket</code>  You could also filter on the FabricPool tenant account.
Advanced Filtering	Object Size (MB) greater than 0.2 MB.  <b>Note:</b> FabricPool only writes 4 MB objects, but you must add an Object Size filter because this rule uses erasure coding.
Reference Time	Ingest Time
Placement	From day 0 store forever
Type	Erasure coded

Rule definition	Example value
Location	DEN (2 plus 1)
Ingest Behavior	Balanced

4. Create an ILM rule that will create two replicated copies of any objects not matched by the first rule. Do not 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 Filtering	<i>none</i>
Reference Time	Ingest Time
Placement	From day 0 store forever
Type	Replicated
Location	DEN
Copies	2
Ingest Behavior	Balanced

5. Create a proposed 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.
6. Ingest test objects into the grid.
7. Simulate the policy with the test objects to verify the behavior.
8. 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 erasure-coded and replicated copies will be maintained in StorageGRID until they are deleted by the S3 client. StorageGRID ILM will never delete these items.

## Creating a traffic classification policy for FabricPool

You can optionally design a StorageGRID traffic classification policy to optimize quality of service for the FabricPool workload.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must 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 the majority of 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.

- As a general rule, 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 might need to be broad to account for the unknown 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.

To learn more: [Managing traffic classification policies](#)

### Steps

1. Select **Configuration > Network Settings > Traffic Classification**.
2. Enter a name and a description.
3. In the Matching Rules section, create at least one rule.
  - a. Select **Create**.
  - b. Select **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**.
4. Optionally, create one or more limits.



Even if no limits are set for a traffic classification policy, metrics are collected so you can understand traffic trends.

- a. Select **Create**.
- b. Select the type of traffic you want to limit and the limit to apply.

This example FabricPool traffic classification lists the types of network traffic you can limit and the types of values you can select. The traffic types and values for an actual policy would be based on your

specific requirements.

### Edit Traffic Classification Policy "FabricPool"

---

#### Policy

Name 

Description (optional)

---

#### Matching Rules

Traffic that matches any rule is included in the policy.

Type	Inverse Match	Match Value
<input checked="" type="radio"/> Endpoint	<input checked="" type="checkbox"/>	FabricPool (https 10443)

Displaying 1 matching rule.

---

#### Limits (Optional)

Type	Value	Units
<input checked="" type="radio"/> Concurrent Read Requests	50	Concurrent Requests
<input checked="" type="radio"/> Concurrent Write Requests	15	Concurrent Requests
<input checked="" type="radio"/> Read Request Rate	100	Requests/Second
<input checked="" type="radio"/> Write Request Rate	25	Requests/Second
<input checked="" type="radio"/> Per-Request Bandwidth In	2000000	Bytes/Second
<input checked="" type="radio"/> Per-Request Bandwidth Out	10000000	Bytes/Second

Displaying 6 limits.

5. After creating the traffic classification policy, select the policy and then select **Metrics** to determine if the policy is limiting traffic as expected.

## Traffic Classification Policies

Traffic classification policies can be used to identify network traffic for metrics reporting and optional traffic limiting.

Name	Description	ID
<input checked="" type="radio"/> FabricPool	Limit traffic other than FabricPool	587f53b2-7cf2-44b9-af5c-694ebbd4a2c5

Displaying 1 traffic classification policy.

## Other best practices for StorageGRID and FabricPool

When configuring a StorageGRID system for use with FabricPool, you should avoid setting global options that might affect how your data is saved.

### Object encryption

When configuring StorageGRID, you can optionally enable the global **Stored Object Encryption** setting if data encryption is required for other StorageGRID clients (**Configuration > System Settings > Grid Options**). 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, do not enable the global **Compress Stored Objects** setting (**Configuration > System Settings > Grid Options**). The data that is tiered from FabricPool to StorageGRID is already compressed. Enabling **Compress Stored Objects** will not further reduce an object's size.

### Consistency level

For FabricPool buckets, the recommended bucket consistency level is **Read-after-new-write**, which is the default setting for a new bucket. Do not edit FabricPool buckets to use **Available** or any other consistency level.

### FabricPool tiering

If the StorageGRID node uses storage assigned from a NetApp AFF 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.

# Use StorageGRID

## Use a tenant account

Learn how to use a StorageGRID tenant account.

- [Using the Tenant Manager](#)
- [Managing system access for tenant users](#)
- [Managing S3 tenant accounts](#)
- [Managing S3 platform services](#)

### Using the Tenant Manager

The Tenant Manager allows you to manage all aspects of a StorageGRID tenant account.

You can use the Tenant Manager to monitor a tenant account's storage usage and to manage users with identity federation or by creating local groups and users. For S3 tenant accounts, you can also manage S3 keys, manage S3 buckets, and configure platform services.

### Using a StorageGRID 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.

Each tenant account has its own federated or local groups, users, S3 buckets or Swift containers, and objects.

Optionally, 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 do not need to create separate tenant accounts. See instructions for implementing S3 client applications.

- **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.

### Creating tenant accounts

Tenant accounts are created by a StorageGRID grid administrator using the Grid Manager. When creating a tenant account, the grid administrator specifies the following information:

- Display name for the tenant (the tenant's account ID is assigned automatically and cannot be changed).
- Whether the tenant account will use the S3 or Swift.
- For S3 tenant accounts: Whether the tenant account is allowed to use platform services. If the use of

platform services is allowed, the grid must be configured to support their use.

- Optionally, a storage quota for the tenant account—the maximum number of gigabytes, terabytes, or petabytes available for the tenant’s objects. A tenant’s storage quota represents a logical amount (object size), not a physical amount (size on disk).
- If identity federation is enabled for the StorageGRID system, which federated group has Root Access permission to configure the tenant account.
- If single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid’s identity source, and the initial password for the tenant’s local root user.

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.

### Configuring S3 tenants

After an S3 tenant account is created, you can access the Tenant Manager to perform tasks such as the following:

- Setting up identity federation (unless the identity source is shared with the grid), or creating local groups and users
- Managing S3 access keys
- Creating and managing S3 buckets, including compliant buckets
- Using platform services (if enabled)
- Monitoring storage usage



While you can create and manage S3 buckets with the Tenant Manager, you must have S3 access keys and use the S3 REST API to ingest and manage objects.

### Configuring Swift tenants

After a Swift tenant account is created, users with the Root Access permission can access the Tenant Manager to perform tasks such as the following:

- Setting up identity federation (unless the identity source is shared with the grid), and creating local groups and users
- Monitoring storage usage



Swift users must have the Root Access permission to access the Tenant Manager. However, the Root Access permission does not allow users to authenticate into the Swift REST API to create containers and ingest objects. Users must have the Swift Administrator permission to authenticate into the Swift REST API.

### Related information

[Administer StorageGRID](#)

[Use S3](#)

[Use Swift](#)

## Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## Signing in to the Tenant Manager

You access the Tenant Manager by entering the URL for the tenant into the address bar of a supported web browser.

### What you'll need

- You must have your login credentials.
- You must 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 contains either the fully qualified domain name (FQDN) or the IP address used to access an Admin Node, and could optionally 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 must have this account ID.


- You must be using a supported web browser.
- Cookies must be enabled in your web browser.
- You must have 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 you see depends on the URL you entered and whether your organization is using single sign-on (SSO). You will see one of the following screens:

- The Grid Manager sign-in page. Click the **Tenant Login** link in the upper right.



- The Tenant Manager sign-in page. The **Account ID** field might already be completed, as shown below.

- i. 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.
- ii. Enter your username and password.
- iii. Click **Sign in**.

The Tenant Manager Dashboard appears.

- Your organization's SSO page, if SSO is enabled on the grid. For example:

Sign in with your organizational account


someone@example.com

Password

Sign in

Enter your standard SSO credentials, and click **Sign in**.

- The Tenant Manager SSO sign-in page.



StorageGRID® Sign in

Recent S3 tenant

Account ID 27469746059057031822

For Grid Manager, leave this field blank.

Sign in

- 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.
- Click **Sign in**.
- Sign in with your standard SSO credentials on your organization's SSO sign-in page.

The Tenant Manager Dashboard appears.

5. If you received an initial password from someone else, change your password to secure your account. Select **username** > **Change Password**.



If SSO is enabled for the StorageGRID system, you cannot change your password from the Tenant Manager.

#### Related information

[Administer StorageGRID](#)


[Web browser requirements](#)

## Signing out of the Tenant Manager

When you are done working with the Tenant Manager, you must sign out to ensure that unauthorized users cannot 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**.

Option	Description
SSO not in use	<p>You are signed out of the Admin Node. The Tenant Manager sign in page is displayed.</p> <p><b>Note:</b> If you signed into more than one Admin Node, you must sign out of each node.</p>
SSO enabled	<p>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 <b>Recent Accounts</b> drop-down, and the tenant's <b>Account ID</b> is shown.</p> <p><b>Note:</b> 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.</p>

## Understanding the 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

## Total objects

8,418,886  
objects

## Tenant details

Name Human Resources  
ID 4955 9096 9804 4285 4354

View the instructions for Tenant Manager.

[Go to documentation](#) [↗](#)

## Tenant account summary

The top of the Dashboard contains the following information:

- The number of configured buckets or containers, groups, and users
- The number of platform services endpoints, if any have been configured

You can select the links to view the details.

The right side of the Dashboard contains the following information:

- The total number of objects for the tenant.

For an S3 account, if no objects have been ingested and you have the Root Access permission, getting started guidelines appear instead of the total number of objects.

- The tenant account name and ID.
- A link to the StorageGRID documentation.

## 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 utilization 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 utilization is recalculated. Quota utilization 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.

## Quota usage alerts

If quota usage alerts have been enabled in the Grid Manager, they 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. For more information, see the alerts reference in the instructions for monitoring and troubleshooting StorageGRID.

 Only 0.6% of the quota is remaining. If the quota is exceeded, you can no longer upload new objects.

If you exceed your quota, you cannot upload new objects.

 The quota has been met. You cannot upload new objects.



To view additional details and manage rules and notifications for alerts, see the instructions for monitoring and troubleshooting StorageGRID.

## 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 an endpoint error, select Endpoints to display the Endpoints page.

## Related information

[Troubleshooting platform services endpoint errors](#)

[Monitor & troubleshoot](#)

## Understanding the 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.

To access the Swagger documentation for the Tenant Management API:

### Steps

1. Sign in to the Tenant Manager.
2. Select **Help > API Documentation** from the Tenant Manager header.

## 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").

See "Protecting against Cross-Site Request Forgery" for information on improving authentication security.



If single sign-on (SSO) 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" in the instructions for administering StorageGRID.

- **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, as follows:

Protocol	Permission allows
S3	<ul style="list-style-type: none"><li>• Creating compliant and non-compliant buckets</li><li>• Modifying legacy compliance settings</li><li>• Setting the consistency control for operations performed on objects</li><li>• Creating, updating, and deleting a bucket's CORS configuration</li><li>• Enabling and disabling last access time updates for objects</li><li>• Managing the configuration settings for platform services, including CloudMirror replication, notifications, and search integration (metadata-notification)</li><li>• Deleting empty buckets</li></ul>
Swift	Setting the consistency level used for 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.
- **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.
- **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 to determine how global S3 Object Lock (compliance) is configured for the StorageGRID system.

- **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
<b>type</b> string <small>(query)</small>	filter by group type
<b>limit</b> integer <small>(query)</small>	maximum number of results
<b>marker</b> string <small>(query)</small>	marker-style pagination offset (value is Group's URN)
<b>includeMarker</b> boolean <small>(query)</small>	if set, the marker element is also returned
<b>order</b> string <small>(query)</small>	pagination order (desc requires marker)

Responses

Response content type application/json ▼

Code	Description
200	<div style="display: flex; justify-content: space-between; align-items: center; margin-bottom: 5px;"> <span>Example Value</span> <span>Model</span> </div> <pre style="background-color: #2d3748; color: white; padding: 5px; border-radius: 3px; font-family: monospace; font-size: 0.9em;"> {   "responseTime": "2018-02-01T16:22:31.066Z",   "status": "success",   "apiVersion": "2.1" } </pre>

## Issuing API requests



Any API operations you perform using the API Docs webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

## Steps

1. Click 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 click **Model** to learn the requirements for each field.
4. Click **Try it out**.
5. Provide any required parameters, or modify the request body as required.
6. Click **Execute**.
7. Review the response code to determine if the request was successful.

### Related information

[Protecting against Cross-Site Request Forgery \(CSRF\)](#)

[Administer StorageGRID](#)

### Tenant Management API versioning

The Tenant Management API uses versioning to support non-disruptive upgrades.

For example, this Request URL specifies version 3 of the API.

```
https://hostname_or_ip_address/api/v3/authorize
```

The major version of the Tenant Management API is bumped when changes are made that are **not compatible** with older versions. The minor version of the Tenant Management 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

When StorageGRID software is installed for the first time, only the most recent version of the Tenant Management API is enabled. However, when StorageGRID is upgraded to a new feature release, you continue to have access to the older API version for at least one StorageGRID feature release.

Outdated requests are marked as deprecated in the following ways:

- The response header is "Deprecated: true"
- The JSON response body includes "deprecated": true

### Determining which API versions are supported in the current release

Use the following API request to return a list of the supported API major versions:

```
GET https://{{IP-Address}}/api/versions
{
  "responseTime": "2019-01-10T20:41:00.845Z",
  "status": "success",
  "apiVersion": "3.0",
  "data": [
    2,
    3
  ]
}
```

### Specifying an API version for a request

You can specify the API version using a path parameter (`/api/v3`) or a header (`Api-Version: 3`). If you provide both values, the header value overrides the path value.

```
curl https://[IP-Address]/api/v3/grid/accounts

curl -H "Api-Version: 3" https://[IP-Address]/api/grid/accounts
```

### Protecting 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.

## Managing system access for tenant users

You grant users access to a tenant account by importing groups from a federated identity source and assigning management permissions. You can also create local tenant groups and users, unless single sign-on (SSO) is in effect for the entire StorageGRID system.

- [Using identity federation](#)
- [Managing groups](#)
- [Managing local users](#)

### Using 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.

- [Configuring a federated identity source](#)
- [Forcing synchronization with the identity source](#)
- [Disabling identity federation](#)

### Configuring a federated identity source

You can configure identity federation if you want tenant groups and users to be managed in another system such as Active Directory, OpenLDAP, or Oracle Directory Server.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must have specific access permissions.
- You must be using Active Directory, OpenLDAP, or Oracle Directory Server as the identity provider. If you want to use an LDAP v3 service that is not listed, you must contact technical support.
- 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.

#### 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 cannot 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.

## Steps

1. Select **ACCESS MANAGEMENT > Identity federation**.
2. Select **Enable identity federation**.
3. In the LDAP service type section, select **Active Directory**, **OpenLDAP**, or **Other**.

If you select **OpenLDAP**, configure the OpenLDAP server. See the guidelines for configuring an OpenLDAP server.

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.
  - **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. In the Configure LDAP server section, enter the required LDAP server and network connection information.
  - **Hostname:** The server hostname 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`
- **Password:** The password associated with the username.
- **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.

6. In the **Transport Layer Security (TLS)** section, select a security setting.

- **Use STARTTLS (recommended):** Use STARTTLS to secure communications with the LDAP server. This is the recommended option.
- **Use LDAPS:** The LDAPS (LDAP over SSL) option uses TLS to establish a connection to the LDAP server. This option is supported for compatibility reasons.
- **Do not use TLS:** The network traffic between the StorageGRID system and the LDAP server will not be secured.

This 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 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.

8. Select **Test connection** to validate your connection settings for the LDAP server.

A confirmation message appears in the upper right corner of the page if the connection is valid.

9. If the connection is valid, select **Save**.

The following screenshot shows example configuration values for an LDAP server that uses Active Directory.

## LDAP service type

Select the type of LDAP service you want to configure.

Active Directory

OpenLDAP

Other

## Configure LDAP server (All fields are required)

Hostname

my-active-directory.example.com

Port

389

Username

MyDomain\Administrator

Password

••••••••

Group Base DN

DC=storagegrid,DC=example,DC=com

User Base DN

DC=storagegrid,DC=example,DC=com

### Related information

[Tenant management permissions](#)

[Guidelines for configuring an OpenLDAP server](#)

### 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.

### 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 Administrator's Guide for OpenLDAP.

## 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 Administrator's Guide for OpenLDAP.

### Forcing 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.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must have specific access permissions.
- The saved identity source must be enabled.

#### Steps

1. Select **ACCESS MANAGEMENT > Identity federation**.

The Identity federation page appears. The **Sync server** button is at the top right of the page.



If the saved identity source is not enabled, the **Sync server** button will not be active.

2. Select **Sync server**.

A confirmation message is displayed indicating that synchronization started successfully.

#### Related information

[Tenant management permissions](#)

#### Disabling identity federation

If you configured an identity federation service for this tenant, you can temporarily or permanently disable identity federation for tenant groups and users. When identity federation is disabled, there is no communication between the StorageGRID system and the identity source. However, any settings you have configured are retained, allowing you to easily re-enable identity federation in the future.

## What you'll need

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

## 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 tenant account 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.

## Steps

1. Select **ACCESS MANAGEMENT > Identity federation**.
2. Deselect the **Enable identity federation** check box.
3. Select **Save**.

## Related information

[Tenant management permissions](#)

## Managing groups

You assign permissions to user groups to control which tasks tenant users can perform. You can import federated groups from an identity source, such as Active Directory or OpenLDAP, or you can create local groups.



If single sign-on (SSO) is enabled for your StorageGRID system, local users will not be able to sign in to the Tenant Manager, although they can access S3 and Swift resources, based on group permissions.

## 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. Changes might take up to 15 minutes to take effect because of caching.

Permission	Description
Root Access	Provides full access to the Tenant Manager and the Tenant Management API.  <b>Note:</b> 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  <b>Note:</b> Swift users must have the Swift Administrator permission to perform any operations with the Swift REST API.
Manage Your Own S3 Credentials	S3 tenants only. Allows users to create and remove their own S3 access keys. Users who do not have this permission do not see the <b>STORAGE (S3) &gt; My S3 access keys</b> menu option.
Manage All Buckets	<ul style="list-style-type: none"> <li>• 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.</li> </ul> <p>Users who do not have this permission do not see the <b>Buckets</b> menu option.</p> <ul style="list-style-type: none"> <li>• Swift tenants: Allows Swift users to control the consistency level for Swift containers using the Tenant Management API.</li> </ul> <p><b>Note:</b> You can only assign the Manage All Buckets permission to Swift groups from the Tenant Management API. You cannot assign this permission to Swift groups using the Tenant Manager.</p>
Manage Endpoints	S3 tenants only. Allows users to use the Tenant Manager or the Tenant Management API to create or edit endpoints, which are used as the destination for StorageGRID platform services.  Users who do not have this permission do not see the <b>Platform services endpoints</b> menu option.

#### Related information

[Use S3](#)

[Use Swift](#)

#### Creating groups for an S3 tenant

You can manage permissions for S3 user groups by importing federated groups or creating local groups.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must 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.

## Steps

1. Select **ACCESS MANAGEMENT > Groups**.



2. Select **Create group**.
3. 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.

4. 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.
5. Select **Continue**.
6. Select an Access mode. 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.
  - **Read-write** (default): Users can log into Tenant Manager and manage the tenant configuration.
  - **Read-only**: Users can only view settings and features. They cannot make any changes or perform any operations in the Tenant Manager or Tenant Management API. Local read-only users can change their own passwords.
7. Select the Group permissions for this group.

See the information about tenant management permissions.

8. Select **Continue**.
9. Select a group policy to determine which S3 access permissions the members of this group will have.
  - **No S3 Access:** Default. Users in this group do not 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 cannot 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 cannot edit this string.
  - **Custom:** Users in the group are granted the permissions you specify in the text box. See the instructions for implementing an S3 client application for detailed information about group policies, including language syntax and examples.
10. 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.

In this example, members of the group are only permitted to list and access a folder matching their username (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.



The screenshot shows the AWS IAM console interface for creating a group policy. On the left, there are four radio button options: "No S3 Access", "Read Only Access", "Full Access", and "Custom". The "Custom" option is selected, and a note below it says "(Must be a valid JSON formatted string.)". On the right, a text area contains a JSON policy string. The JSON string defines two statements: one for listing buckets in a specific folder and another for allowing object actions in a specific folder.

```
{
  "Statement": [
    {
      "Sid": "AllowListBucketOfASpecificUserPrefix",
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::department-bucket",
      "Condition": {
        "StringLike": {
          "s3:prefix": "${aws:username}/*"
        }
      }
    },
    {
      "Sid": "AllowUserSpecificActionsOnlyInTheSpecificFolder",
      "Effect": "Allow",
      "Action": "s3:*Object",
      "Resource": "arn:aws:s3:::department-bucket/${aws:username}/*"
    }
  ]
}
```

11. Select the button that appears, depending on whether you are creating a federated group or a local group:
  - Federated group: **Create group**

- Local group: **Continue**

If you are creating a local group, step 4 (Add users) appears after you select **Continue**. This step does not appear for federated groups.

12. Select the check box for each user you want to add to the group, then select **Create group**.

Optionally, you can save the group without adding users. You can add users to the group later, or select the group when you add new users.

13. Select **Finish**.

The group you created appears in the list of groups. Changes might take up to 15 minutes to take effect because of caching.

## Related information

[Tenant management permissions](#)

[Use S3](#)

## Creating 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.

## What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must 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.

## Steps

1. Select **ACCESS MANAGEMENT > Groups**.



2. Select **Create group**.
3. 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.

4. 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.
5. Select **Continue**.
6. Select an Access mode. 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.
  - **Read-write** (default): Users can log into Tenant Manager and manage the tenant configuration.
  - **Read-only**: Users can only view settings and features. They cannot make any changes or perform any operations in the Tenant Manager or Tenant Management API. Local read-only users can change their own passwords.
7. Set the Group permission.
  - Select the **Root Access** check box if users need to sign in to the Tenant Manager or Tenant Management API. (Default)
  - Unselect the **Root Access** check box if users do not need access to the Tenant Manager or Tenant Management API. For example, unselect the check box for applications that do not need to access the tenant. Then, assign the **Swift Administrator** permission to allow these users to manage containers and objects.
8. Select **Continue**.

9. Select the **Swift administrator** check box if the user needs to be able to use the Swift REST API.

Swift users must have the Root Access permission to access the Tenant Manager. However, the Root Access permission does not allow users to authenticate into the Swift REST API to create containers and ingest objects. Users must have the Swift Administrator permission to authenticate into the Swift REST API.

10. Select the button that appears, depending on whether you are creating a federated group or a local group:

- Federated group: **Create group**
- Local group: **Continue**

If you are creating a local group, step 4 (Add users) appears after you select **Continue**. This step does not appear for federated groups.

11. Select the check box for each user you want to add to the group, then select **Create group**.

Optionally, you can save the group without adding users. You can add users to the group later, or select the group when you create new users.

12. Select **Finish**.

The group you created appears in the list of groups. Changes might take up to 15 minutes to take effect because of caching.

## Related information

[Tenant management permissions](#)

[Use Swift](#)

## Viewing and editing group details

When you view the details for a group, you can change the group's display name, permissions, policies, and the users that belong to the group.

## What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Root Access permission.

## Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select the name of the group whose details you want to view or edit.

Alternatively, you can select **Actions > View group details**.

The group details page appears. The following example shows the S3 group details page.

## Overview

Display name:	<b>Applications</b> 
Unique name:	<b>group/Applications</b>
Type:	<b>Local</b>
Access mode:	<b>Read-write</b>
Permissions:	<b>Root Access</b>
S3 Policy:	<b>None</b>
Number of users in this group:	<b>0</b>

Group permissions

S3 group policy

Users

## Manage group permissions

Select an access mode for this group and select one or more permissions.

### Access mode

Select whether users can change settings and perform operations or whether they can only view settings and features.

Read-write  Read-only

### Group permissions

Select the tenant account permissions you want to assign to this group.

**Root Access**

Allows users to access all Tenant Manager features. Root Access permission supersedes all other permissions.

**Manage All Buckets**

Allows users to change settings of all S3 buckets (or Swift containers) in this account.

**Manage Endpoints**

Allows users to configure endpoints for platform services.

**Manage Your Own S3 Credentials**

Allows users to create and delete their own S3 access keys.

Save changes

### 3. Make changes to the group settings as needed.



To ensure your changes are saved, select **Save changes** after you make changes in each section. When your changes are saved, a confirmation message appears in the upper right corner of the page.

- a. Optionally, select the display name or edit icon  to update the display name.

You cannot change a group's unique name. You cannot edit the display name for a federated group.

- b. Optionally, update the permissions.

- c. For group policy, make the appropriate changes for your S3 or Swift tenant.

- If you are editing a group for an S3 tenant, optionally select a different S3 group policy. If you select a custom S3 policy, update the JSON string as required.
- If you are editing a group for a Swift tenant, optionally select or unselect the **Swift Administrator** check box.

For more information about the Swift Administrator permission, see the instructions for creating groups for a Swift tenant.

- d. Optionally, add or remove users.

### 4. Confirm that you have selected **Save changes** for each section you changed.

Changes might take up to 15 minutes to take effect because of caching.

#### Related information

[Creating groups for an S3 tenant](#)

[Creating groups for a Swift tenant](#)

#### Adding users to a local group

You can add users to a local group as needed.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Root Access permission.

#### Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select the name of the local group you want to add users to.

Alternatively, you can select **Actions > View group details**.

The group details page appears.

## Overview

Display name:	<b>Applications</b> 
Unique name:	<b>group/Applications</b>
Type:	<b>Local</b>
Access mode:	<b>Read-write</b>
Permissions:	<b>Root Access</b>
S3 Policy:	<b>None</b>
Number of users in this group:	<b>0</b>

Group permissions

S3 group policy

Users

## Manage group permissions

Select an access mode for this group and select one or more permissions.

### Access mode

Select whether users can change settings and perform operations or whether they can only view settings and features.

Read-write  Read-only

### Group permissions

Select the tenant account permissions you want to assign to this group.

**Root Access**

Allows users to access all Tenant Manager features. Root Access permission supersedes all other permissions.

**Manage All Buckets**

Allows users to change settings of all S3 buckets (or Swift containers) in this account.

**Manage Endpoints**

Allows users to configure endpoints for platform services.

**Manage Your Own S3 Credentials**

Allows users to create and delete their own S3 access keys.

Save changes

3. Select **Manage Users**, and then select **Add users**.

Username	Full Name	Denied
User_02	User_02_Managers	

4. Select the users you want to add to the group, and then select **Add users**.

<input checked="" type="checkbox"/>	Username	Full Name	Denied
<input checked="" type="checkbox"/>	User_01	User_01_Applications	

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

### Editing a group name

You can edit the display name for a group. You cannot edit the unique name for a group.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Root Access permission.

### Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select the check box for the group whose display name you want to edit.
3. Select **Actions > Edit group name**.

The Edit group name dialog box appears.

**Edit group name** ✕

Specify a new name for the group **Applications**.

Must contain at least 1 and no more than 32 characters

Applications

Cancel Save changes

4. If you are editing a local group, update the display name as needed.

You cannot change a group's unique name. You cannot edit the display name for a federated group.

5. Select **Save changes**.

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

### Related information

[Tenant management permissions](#)

### Duplicating a group

You can create new groups more quickly by duplicating an existing group.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Root Access permission.

### Steps

1. Select **ACCESS MANAGEMENT > Groups**.
2. Select the check box for the group you want to duplicate.
3. Select **Duplicate group**. For additional details on creating a group, see the instructions for creating groups for an S3 tenant or for a Swift tenant.
4. 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.

5. 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.

6. Select **Continue**.
7. As needed, modify the permissions for this group.
8. Select **Continue**.
9. As needed, if you are duplicating a group for an S3 tenant, optionally select a different policy from the **Add S3 policy** radio buttons. If you selected a custom policy, update the JSON string as required.
10. Select **Create group**.

### Related information

[Creating groups for an S3 tenant](#)

[Creating groups for a Swift tenant](#)

[Tenant management permissions](#)

### Deleting a group

You can delete a group from the system. Any users who belong only to that group will no longer be able to sign in to the Tenant Manager or use the tenant account.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Root Access permission.

### Steps

1. Select **ACCESS MANAGEMENT > Groups**.



2 groups Create group

Actions ▾

<input type="checkbox"/>	Name ▾	ID ▾	Type ▾	Access mode ▾
<input type="checkbox"/>	Applications	22cc2e27-88ee-4461-a8c6-30b550beec0	Local	Read-write
<input type="checkbox"/>	Managers	8b15b131-1d21-4539-93ad-f2298347c4d8	Local	Read-write

← Previous **1** Next →

2. Select the check boxes for the groups you want to delete.

3. Select **Actions** > **Delete group**.

A confirmation message appears.

4. Select **Delete group** to confirm you want to delete the groups indicated in the confirmation message.

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

### Related information

[Tenant management permissions](#)

### Managing 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 cannot remove the root user.

### What you’ll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a read-write user group that has the Root Access permission.



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 S3 or Swift client applications to access the tenant’s resources, based on group permissions.

### Accessing the Users page

Select **ACCESS MANAGEMENT** > **Users**.

# Users

View local and federated users. Edit properties and group membership of local users.

3 users

Create user

Actions

<input type="checkbox"/>	Username	Full Name	Denied	Type
<input type="checkbox"/>	root	Root		Local
<input type="checkbox"/>	User_01	User_01		Local
<input type="checkbox"/>	User_02	User_02		Local

## Creating local users

You can create local users and assign them to one or more local groups to control their access permissions.

S3 users who do not belong to any groups do not 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 do not belong to any groups do not have management permissions or Swift container access.

## Steps

1. Select **Create user**.
2. Complete the following fields.
  - **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 cannot be changed.
  - **Password**: A password, which is used when the user signs in.
  - **Confirm password**: Type the same password you typed in the Password field.
  - **Deny access**: If you select **Yes**, this user cannot sign in to the tenant account, even though the user might still belong to one or more groups.

As an example, you can use this feature to temporarily suspend a user's ability to sign in.

3. Select **Continue**.
4. Assign the user to one or more local groups.

Users who do not belong to any groups will have no management permissions. Permissions are cumulative. Users will have all permissions for all groups they belong to.

#### 5. Select **Create user**.

Changes might take up to 15 minutes to take effect because of caching.

### Editing user details

When you edit the details for a user, you can change the user's full name and password, add the user to different groups, and prevent the user from accessing the tenant.

### Steps

1. In the Users list, select the name of the user whose details you want to view or edit.

Alternatively, you can select the check box for the user, and then select **Actions > View user details**.

2. Make changes to the user settings as needed.

- a. Change the user's full name as needed by selecting the full name or the edit icon  in the Overview section.

You cannot change the username.

- b. On the **Password** tab, change the user's password as needed.
- c. On the **Access** tab, allow the user to sign in (select **No**), or prevent the user from signing in (select **Yes**) as needed.
- d. On the **Groups** tab, add the user to groups or remove the user from groups as needed.
- e. As necessary for each section, select **Save changes**.

Changes might take up to 15 minutes to take effect because of caching.

### Duplicating local users

You can duplicate a local user to create a new user more quickly.

### Steps

1. In the Users list, select the user you want to duplicate.
2. Select **Duplicate user**.
3. Modify the following fields for the new user.
  - **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 cannot be changed.
  - **Password**: A password, which is used when the user signs in.
  - **Confirm password**: Type the same password you typed in the Password field.
  - **Deny access**: If you select **Yes**, this user cannot sign in to the tenant account, even though the user might still belong to one or more groups.

As an example, you can use this feature to temporarily suspend a user's ability to sign in.

4. Select **Continue**.
5. Select one or more local groups.

Users who do not belong to any groups will have no management permissions. Permissions are cumulative. Users will have all permissions for all groups they belong to.

6. Select **Create user**.

Changes might take up to 15 minutes to take effect because of caching.

### Deleting local users

You can permanently delete local users who no longer need to access the StorageGRID tenant account.

Using the Tenant Manager, you can delete local users, but not federated users. You must use the federated identity source to delete federated users.

### Steps

1. In the Users list, select the check box for the local user you want to delete.
2. Select **Actions > Delete user**.
3. In the confirmation dialog box, select **Delete user** to confirm you want to delete the user from the system.

Changes might take up to 15 minutes to take effect because of caching.

### Related information

[Tenant management permissions](#)

## Managing S3 tenant accounts

You can use the Tenant Manager to manage S3 access keys and to create and manage S3 buckets.

- [Managing S3 access keys](#)
- [Managing S3 buckets](#)

### Managing 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.

### About this task

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.

### Creating 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 in the S3 tenant account.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must have the Manage Your Own S3 Credentials 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, do not create more keys than you need, and delete the keys you are not 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 do not need to periodically create new keys, you do not 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.

1 Choose expiration time ————— 2 Download access key

### Choose expiration time

Do not set an expiration time  
This access key will never expire.

Set an expiration time

MM/DD/YYYY HH : MM AM

Cancel **Create access key**


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.



Do not close this dialog box until you have copied or downloaded this information.



6. Select **Finish**.

The new key is listed on the My access keys page. Changes might take up to 15 minutes to take effect because of caching.

### Related information

[Tenant management permissions](#)

### Viewing 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.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must 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**.

The My access keys page appears and lists any existing access keys.

<input type="checkbox"/>	Access key ID	Expiration time
<input type="checkbox"/>	*****OTLS	2020-11-23 12:00:00 MST
<input type="checkbox"/>	*****0M45	2020-12-01 19:00:00 MST
<input type="checkbox"/>	*****69QJ	None
<input type="checkbox"/>	*****3R8P	None

2. Sort the keys by **Expiration time** or **Access key ID**.

3. As needed, create new keys and manually delete 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.

**Related information**

[Creating your own S3 access keys](#)

[Deleting your own S3 access keys](#)

**Deleting 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.

## What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must 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**.

The My access keys page appears and lists any existing access keys.

2. Select the check box for each access key you want to remove.
3. Select **Delete key**.

A confirmation dialog box appears.



4. Select **Delete key**.

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

## Related information

[Tenant management permissions](#)

### Creating 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.

## What you'll need

- You must be signed in to the Tenant Manager using a supported browser.

- You must have 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, do not create more keys than the user needs, and delete the keys that are not 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 do not need to periodically create new keys, you do not 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 **Do not 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.

# Create access key

1 Choose expiration time ————— 2 Download access key

## Choose expiration time

Do not set an expiration time  
This access key will never expire.

Set an expiration time

MM/DD/YYYY  HH : MM AM

Cancel **Create access key**

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.



Do not close this dialog box until you have copied or downloaded this information.

Create access key

1 Choose expiration time ————— 2 Download access key

### Download access key

To save the keys for future reference, select **Download .csv**, or copy and paste the values to another location.

**i** You will not be able to view the Access key ID or Secret access key after you close this dialog.

Access key ID

SHTBQKDRVHQ34YKKUAWX

Secret access key

UGu9+XeACtnOWQYFdbzmgmgVXXDvCkSOzT1Osz9K

Download .csv Finish

7. Select **Finish**.

The new key is listed on the Access keys tab of the user details page. Changes might take up to 15 minutes to take effect because of caching.

### Related information

[Tenant management permissions](#)

### Viewing 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.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must 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**.

The Users page appears and lists the existing users.

2. Select the user whose S3 access keys you want to view.

The User details page appears.

3. Select **Access keys**.

Manage access keys

Add or delete access keys for this user.

Create key Actions

Displaying 4 results

<input type="checkbox"/>	Access key ID	Expiration time
<input type="checkbox"/>	*****WX5J	2020-11-21 12:00:00 MST
<input type="checkbox"/>	*****6OHM	2020-11-23 13:00:00 MST
<input type="checkbox"/>	*****J505	None
<input type="checkbox"/>	*****4MTF	None

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

[Creating another user's S3 access keys](#)

[Deleting another user's S3 access keys](#)

## Deleting 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.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must 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**.

The Users page appears and lists the existing users.


2. Select the user whose S3 access keys you want to manage.

The User details page appears.

3. Select **Access keys**, and then select the check box for each access key you want to delete.

4. Select **Actions > Delete selected key**.

A confirmation dialog box appears.



5. Select **Delete key**.

A confirmation message appears in the upper right corner of the page. Changes might take up to 15 minutes to take effect because of caching.

### Related information

[Tenant management permissions](#)

## Managing S3 buckets

If you are using an S3 tenant with the appropriate permissions, you can create, view, and delete S3 buckets, update consistency level settings, configure Cross-Origin Resource Sharing (CORS), enable and disable last access time update settings, and manage S3 platform services.


### Using S3 Object Lock

You can use the S3 Object Lock feature in StorageGRID if your objects must comply with regulatory requirements for retention.

### 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).

As shown in the figure, 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, S3 client applications can optionally specify retention settings for any object version in that bucket. An object version must have retention settings specified to be protected by S3 Object Lock.



The StorageGRID S3 Object Lock feature provides a single retention mode that is equivalent to the Amazon S3 compliance mode. By default, a protected object version cannot be overwritten or deleted by any user. The StorageGRID S3 Object Lock feature does not support a governance mode, and it does not allow users with special permissions to bypass retention settings or to delete protected objects.

If a bucket has S3 Object Lock enabled, the S3 client application can optionally specify either or both of the following object-level retention settings when creating or updating an object:

- **Retain-until-date:** If an object version's retain-until-date is in the future, the object can be retrieved, but it cannot be modified or deleted. As required, an object's retain-until-date can be increased, but this date cannot be decreased.
- **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.

For details on these settings, go to “using S3 object lock” in [S3 REST API supported operations and limitations](#).

### Managing 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 the NetApp Knowledge Base article.


[NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

### S3 Object Lock workflow

The workflow diagram shows the high-level steps for using the S3 Object Lock feature in StorageGRID.

Before you can create buckets with S3 Object Lock enabled, the grid administrator must enable the global S3 Object Lock setting for the entire StorageGRID system. The grid administrator must also ensure that the information lifecycle management (ILM) policy is “compliant”; it must meet the requirements of buckets with S3 Object Lock enabled. For details, contact your grid administrator or see the instructions for managing objects with information lifecycle management.

After the global S3 Object Lock setting has been enabled, you can create buckets with S3 Object Lock enabled. You can then use the S3 client application to optionally specify retention settings for each object version.



**Related information**

[Manage objects with ILM](#)

**Requirements for S3 Object Lock**

Before enabling S3 Object Lock for a bucket, review the requirements for S3 Object Lock buckets and objects and the lifecycle of objects in buckets with S3 Object Lock enabled.

**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.

This example from the Tenant Manager shows a bucket with S3 Object Lock enabled.

# Buckets

Create buckets and manage bucket settings.

1 bucket

Create bucket

Actions ▾

<input type="checkbox"/>	Name ▾	S3 Object Lock  ▾	Region ▾	Object Count  ▾	Space Used  ▾	Date Created ▾
<input type="checkbox"/>	bank-records	✓	us-east-1	0	0 bytes	2021-01-06 16:53:19 MST

← Previous 1 Next →

- If you plan to use S3 Object Lock, you must enable S3 Object Lock when you create the bucket. You cannot enable S3 Object Lock for an existing bucket.
- Bucket versioning is required with S3 Object Lock. When S3 Object Lock is enabled for a bucket, StorageGRID automatically enables versioning for that bucket.
- After you create a bucket with S3 Object Lock enabled, you cannot disable S3 Object Lock or suspend versioning for that bucket.
- An StorageGRID bucket that has S3 Object Lock enabled does not have a default retention period. Instead, the S3 client application can optionally specify a retention date and legal hold setting for each object version that is added to that bucket.
- Bucket lifecycle configuration is supported for S3 Object Lifecycle buckets.
- CloudMirror replication is not supported for buckets with S3 Object Lock enabled.

## Requirements for objects in buckets with S3 Object Lock enabled

- The S3 client application must specify retention settings for each object that needs to be protected by S3 Object Lock.
- You can increase the retain-until-date for an object version, but you can never decrease this value.
- If you are notified of a pending legal action or regulatory investigation, you can preserve relevant information by placing a legal hold on an object version. When an object version is under a legal hold, that object cannot be deleted from StorageGRID, even if it has reached its retain-until-date. As soon as the legal hold is lifted, the object version can be deleted if the retain-until-date has been reached.
- S3 Object Lock requires the use of versioned buckets. 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 three stages:

### 1. Object ingest

- When adding an object version to a bucket with S3 Object Lock enabled, the S3 client application can optionally specify retention settings for the object (retain-until-date, legal hold, or both). StorageGRID

then generates metadata for that object, which includes a unique object identifier (UUID) and the ingest date and time.

- After an object version with retention settings is ingested, its data and S3 user-defined metadata cannot be modified.
- StorageGRID stores the object metadata independently of the object data. It maintains three copies of all object metadata at each site.

## 2. Object retention

- Multiple copies of the object are stored by StorageGRID. The exact number and type of copies and the storage locations are determined by the compliant rules in the active ILM policy.

## 3. Object deletion

- An object can be deleted when its retain-until-date is reached.
- An object that is under a legal hold cannot be deleted.

### Creating an S3 bucket

You can use the Tenant Manager to create S3 buckets for object data. When you create a bucket, you must specify the bucket's name and region. If the global S3 Object Lock setting is enabled for the StorageGRID system, you can optionally enable S3 Object Lock for the bucket.

### What you'll need


- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies.
- If you plan to create a bucket with S3 Object Lock, the global S3 Object Lock setting must have been enabled for the StorageGRID system and you must have reviewed the requirements for S3 Object Lock buckets and objects.

### [Using S3 Object Lock](#)

### Steps


1. Select **STORAGE (S3) > Buckets**.

The Buckets page appears and lists any buckets that have already been created.



2. Select **Create bucket**.

The Create bucket wizard appears.



If the global S3 Object Lock setting is enabled, Create bucket includes a second step for managing S3 Object Lock for the bucket.

3. Enter a unique name for the bucket.



You cannot change the bucket name after creating the bucket.

Bucket names must comply 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.
- 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.



See the Amazon Web Services (AWS) Documentation for more information.

#### 4. Select the region for this bucket.

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 `us-east-1` region.



You cannot change the region after creating the bucket.

#### 5. Select **Create bucket** or **Continue**.

- If the global S3 Object Lock setting is not enabled, select **Create bucket**. The bucket is created and added to the table on the Buckets page.
- If the global S3 Object Lock setting is enabled, select **Continue**. Step 2, Manage S3 Object Lock, appears.

**Create bucket**

Enter details ————— 2 Manage S3 Object Lock  
Optional

**Manage S3 Object Lock** (This step is optional)

S3 Object Lock allows you to specify retention and legal hold settings for the objects ingested into a bucket. If you want to use S3 Object Lock, you must enable this setting when you create the bucket. You cannot add or disable S3 Object Lock after a bucket is created.

If S3 Object Lock is enabled, bucket versioning is required and will be enabled automatically.

Enable S3 Object Lock

Previous **Create bucket**

#### 6. Optionally, select the check box to enable S3 Object Lock for this bucket.

S3 Object Lock must be enabled for the bucket before an S3 client application can specify retain-until-date and legal hold settings for the objects added to the bucket.



You cannot enable or disable S3 Object Lock after creating the bucket.



If you enable S3 Object Lock for a bucket, bucket versioning is enabled automatically.

## 7. Select **Create bucket**.

The bucket is created and added to the table on the Buckets page.

### Related information

[Manage objects with ILM](#)

[Understanding the Tenant Management API](#)

[Use S3](#)

### Viewing S3 bucket details

You can view a list of the buckets and bucket settings in your tenant account.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.

### Steps

#### 1. Select **STORAGE (S3) > Buckets**.

The Buckets page appears and lists all buckets for the tenant account.

<input type="checkbox"/>	Name	S3 Object Lock	Region	Object Count	Space Used	Date Created
<input type="checkbox"/>	bucket-01	✓	us-east-1	0	0 bytes	2020-11-04 14:16:59 MST
<input type="checkbox"/>	bucket-02		us-east-1	0	0 bytes	2020-11-04 14:17:14 MST

#### 2. Review the information for each bucket.

As required, you can sort the information by any column, or you can page forward and back through the list.

- Name: The bucket's unique name, which cannot be changed.
- S3 Object Lock: Whether S3 Object Lock is enabled for this bucket.

This column is not displayed if the global S3 Object lock setting is disabled. This column also shows information for any legacy Compliant buckets.

- Region: The bucket's region, which cannot be changed.
- Object Count: The number of objects in this bucket.
- Space Used: The logical size of all objects in this bucket. The logical size does not include the actual space required for replicated or erasure-coded copies or for object metadata.
- Date Created: The date and time the bucket was created.



The Object Count and Space Used values displayed are estimates. These estimates are affected by the timing of ingests, network connectivity, and node status.

3. To view and manage the settings for a bucket, select the bucket name.

The bucket details page appears.

This page allows you to view and edit the settings for bucket options, bucket access, and platform services.

See the instructions for configuring each setting or platform service.

Buckets > bucket-02

### Overview

Name:	<b>bucket-02</b>
Region:	<b>us-east-1</b>
S3 Object Lock:	<b>Disabled</b>
Date created:	<b>2020-11-04 14:51:59 MST</b>

**Bucket options**    Bucket access    Platform services

Consistency level	Read-after-new-write	▼
Last access time updates	Disabled	▼

#### Related information

[Changing the consistency level](#)

[Enabling or disabling last access time updates](#)

[Configuring Cross-Origin Resource Sharing \(CORS\)](#)

[Configuring CloudMirror replication](#)

[Configuring event notifications](#)

[Configuring the search integration service](#)

### Changing the consistency level

If you are using an S3 tenant, you can use the Tenant Manager or the Tenant Management API to change the consistency control for operations performed on the objects in S3 buckets.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies.

#### About this task

Consistency level makes a trade-off between the availability of the objects and the consistency of those objects across different Storage Nodes and sites. In general, you should use the **Read-after-new-write** consistency level for your buckets. If the **Read-after-new-write** consistency level does not meet the client application's requirements, you can change the consistency level by setting the bucket consistency level or by using the `Consistency-Control` header. The `Consistency-Control` header overrides the bucket consistency level.



When you change a bucket's consistency level, only those objects that are ingested after the change are guaranteed to meet the revised level.

#### Steps

1. Select **STORAGE (S3) > Buckets**.
2. Select the bucket name from the list.

The bucket details page appears.

3. Select **Bucket options > Consistency level**.

Bucket options
Bucket access
Platform services

Consistency level
Read-after-new-write (default)
⤴

Change the consistency control for operations performed on the objects in the bucket. Consistency level makes a trade-off between the availability of the objects and the consistency of those objects across different Storage Nodes and sites.

In general, use the **Read-after-new-write** consistency level for your buckets. Then, if objects do not meet availability or consistency requirements, change the client application's behavior, or set the Consistency-Control header for an individual API request, which overrides the bucket setting.

- All**  
Provides the highest guarantee 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.

Note: If your application attempts HEAD operations on keys that do not exist, set the Consistency Level to **Available**, unless you require AWS S3 consistency guarantees. Otherwise, a high number of 500 Internal Server errors can result if one or more Storage Nodes are unavailable.

- Available**  
Behaves the same as the **Read-after-new-write** consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than **Read-after-new-write** if Storage Nodes are unavailable. Differs from AWS S3 consistency guarantees for HEAD operations only.

Save changes

4. Select a consistency level for operations performed on the objects in this bucket.

Consistency level	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.

Consistency level	Description
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. Matches Amazon S3 consistency guarantees.  <b>Note:</b> If your application attempts HEAD operations on keys that do not exist, set the Consistency Level to <b>Available</b> , unless you require Amazon S3 consistency guarantees. Otherwise, a high number of 500 Internal Server errors can result if one or more Storage Nodes are unavailable.
Available (eventual consistency for HEAD operations)	Behaves the same as the <b>Read-after-new-write</b> consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than <b>Read-after-new-write</b> if Storage Nodes are unavailable. Differs from Amazon S3 consistency guarantees for HEAD operations only.

5. Select **Save changes**.

#### Related information

[Tenant management permissions](#)

#### Enabling or disabling 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 in its placement instructions. You can ignore these instructions if your StorageGRID system does not include such a rule.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies.

**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.



See the instructions for managing objects with information lifecycle management.

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 copy an object from one bucket to another	<ul style="list-style-type: none"> <li>• No, for the source copy</li> <li>• Yes, for the destination copy</li> </ul>	<ul style="list-style-type: none"> <li>• No, for the source copy</li> <li>• Yes, for the destination copy</li> </ul>	<ul style="list-style-type: none"> <li>• Yes, for the source copy</li> <li>• Yes, for the destination copy</li> </ul>	<ul style="list-style-type: none"> <li>• Yes, for the source copy</li> <li>• Yes, for the destination copy</li> </ul>
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 **STORAGE (S3) > Buckets**.
2. Select the bucket name from the list.

The bucket details page appears.

3. Select **Bucket options > Last access time updates**.
4. Select the appropriate radio button to enable or disable last access time updates.

Bucket options
Bucket access
Platform services

Consistency level Read-after-new-write ▼

Last access time updates Disabled ▲

Enable or disable last access time updates for the objects in this bucket.

When last access time updates are disabled, the following behavior applies to objects in the bucket:

- Requests to retrieve an object, its access control list, or its metadata do not update the object's last access time. The object is not added to ILM evaluation queues.
- Requests to update an object's metadata update the object's last access time. The object is added to ILM evaluation queues.
- Requests to copy an object from one bucket to another do not update the last access time for the source copy and do not add the source object to the ILM evaluation queue. However, the last access time is updated for the destination copy, and the destination object is added to ILM evaluation queues.
- A request to complete a multipart upload causes the last access time for the assembled object to be updated. The new object is added to ILM evaluation queues.

i Updating the last access time when an object is retrieved can reduce performance, especially for small objects.

Enable last access time updates when retrieving an object

Disable last access time updates when retrieving an object

Save changes

5. Select **Save changes**.

## Related information

[Tenant management permissions](#)

[Manage objects with ILM](#)

## Configuring 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.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies.

### 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>.

### Steps

1. Use a text editor to create the XML required to enable CORS.

This example shows the XML used to enable CORS for an S3 bucket. This XML allows any domain to send GET requests to the bucket, but it only allows the `http://www.example.com` domain to send 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 Developer Guide](#).

2. In the Tenant Manager, select **STORAGE (S3) > Buckets**.
3. Select the bucket name from the list.

The bucket details page appears.

4. Select **Bucket access > Cross-Origin Resource Sharing (CORS)**.
5. Select the **Enable CORS** check box.
6. Paste the CORS configuration XML into the text box, and select **Save changes**.

**Bucket options** | **Bucket access** | Platform services

**Cross-Origin Resource Sharing (CORS)** Disabled ▲

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.

Enable CORS

Clear

```
<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>
```

Save changes

7. To modify the CORS setting for the bucket, update the CORS configuration XML in the text box or select **Clear** to start over. Then select **Save changes**.
8. To disable CORS for the bucket, unselect the **Enable CORS** check box, and then select **Save changes**.

### Deleting an S3 bucket

You can use the Tenant Manager to delete an S3 bucket that is empty.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission. These permissions override the permissions settings in group or bucket policies.

#### 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 cannot delete an S3 bucket if it contains objects or noncurrent object versions. For information about how S3 versioned objects are deleted, see the instructions for managing objects with information lifecycle management.

## Steps

1. Select **STORAGE (S3) > Buckets**.

The Buckets page appears and shows all existing S3 buckets.



The screenshot shows the AWS S3 Buckets page. At the top, it says "Buckets" and "Create buckets and manage bucket settings." Below that, it indicates "2 buckets" and has a "Create bucket" button. An "Actions" dropdown menu is visible. The main content is a table with the following columns: Name, S3 Object Lock, Region, Object Count, Space Used, and Date Created. There are two rows of buckets:


<input type="checkbox"/>	Name	S3 Object Lock	Region	Object Count	Space Used	Date Created
<input type="checkbox"/>	bucket-01	✓	us-east-1	0	0 bytes	2020-11-04 14:16:59 MST
<input type="checkbox"/>	bucket-02		us-east-1	0	0 bytes	2020-11-04 14:17:14 MST

At the bottom right of the table, there are navigation arrows: "← Previous 1 Next →".

2. Select the check box for the empty bucket you want to delete.

The Actions menu is enabled.


3. From the Actions menu, select **Delete empty bucket**.



The screenshot shows the AWS S3 Buckets page with the "Actions" dropdown menu open. The "Delete empty bucket" option is selected. The table below shows that the checkbox for "bucket-02" is now checked:

<input type="checkbox"/>	Name	S3 Object Lock	Region	Object Count	Space Used	Date Created
<input type="checkbox"/>	bucket-01	✓	us-east-1	0	0 bytes	2020-11-04 14:16:59 MST
<input checked="" type="checkbox"/>	bucket-02		us-east-1	0	0 bytes	2020-11-04 14:17:14 MST

A confirmation message appears.



The screenshot shows a confirmation dialog box titled "Delete empty bucket". The text inside asks: "Are you sure you want to delete empty bucket bucket-02?". At the bottom, there are two buttons: "Cancel" and "Delete bucket".

4. If you are sure you want to delete the bucket, select **Delete bucket**.

StorageGRID confirms that the bucket is empty and then deletes the bucket. This operation might take a few minutes.

If the bucket is not empty, an error message appears. You must delete all objects before you can delete the bucket.

 Unable to delete the bucket because it is not empty. You must delete all objects before you can delete this bucket.

## Related information

[Manage objects with ILM](#)

## Managing S3 platform services

If the use of platform services is allowed for your S3 tenant account, you can use platform services to leverage external services and configure CloudMirror replication, notifications, and search integration for S3 buckets.

- [What platform services are](#)
- [Considerations for using platform services](#)
- [Configuring platform services endpoints](#)
- [Configuring CloudMirror replication](#)
- [Configuring event notifications](#)
- [Using the search integration service](#)

## What platform services are

StorageGRID platform services can help you implement a hybrid cloud strategy.

If the use of platform services is allowed for your tenant account, you can configure the following services for any S3 bucket:

- **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 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 Amazon Simple Notification Service™ (SNS).

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.

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, 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, a Simple Notification Service (SNS) topic, or an Elasticsearch cluster hosted locally, on AWS, or elsewhere.

After you create an 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:

1. 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.
2. If you also want to send notifications when these objects are stored to the bucket, you must add a notifications configuration.
3. Finally, 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
CloudMirror replication	<ul style="list-style-type: none"> <li>• GET Bucket replication</li> <li>• PUT Bucket replication</li> </ul>
Notifications	<ul style="list-style-type: none"> <li>• GET Bucket notification</li> <li>• PUT Bucket notification</li> </ul>
Search integration	<ul style="list-style-type: none"> <li>• GET Bucket metadata notification configuration</li> <li>• PUT Bucket metadata notification configuration</li> </ul> <p>These operations are custom to StorageGRID.</p>

See the instructions for implementing S3 client applications for details on how StorageGRID implements these APIs.

### Related information

[Use S3](#)

[Understanding the CloudMirror replication service](#)

[Understanding notifications for buckets](#)

[Understanding the search integration service](#)

[Considerations for using platform services](#)

### Understanding the 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 destination buckets.

CloudMirror replication operates independently of the grid's active ILM policy. 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.

If you enable CloudMirror replication for an existing bucket, only the new objects added to that bucket are replicated. Any existing objects in the bucket are not 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 AWS 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.

In StorageGRID, you can replicate the objects in a single bucket to multiple destination buckets. To do so, specify the destination for each rule in the replication configuration XML. You cannot replicate an object to more than one bucket at the same time.

Additionally, you can configure CloudMirror replication on versioned or unversioned buckets, and you can specify a versioned or unversioned bucket as the destination. 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 behavior for the CloudMirror replication service is the same as the deletion behavior of the Cross Region Replication (CRR) service provided by Amazon S3 — 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 does not replicate the delete marker to the destination bucket or delete the destination object.

As objects are replicated to the destination bucket, StorageGRID marks them as “replicas.” A destination StorageGRID bucket will not replicate objects marked as replicas again, protecting you from accidental replication loops. This replica marking is internal to StorageGRID and does not 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 bi-directional CloudMirror between two grids.

The uniqueness and ordering of events in the destination bucket are not 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.

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.

## Related information

[Configuring CloudMirror replication](#)

## Understanding 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 Amazon Simple Notification Service (SNS).

You can 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 SNS topic specified as the URN of an endpoint.

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 are not 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.

## Supported notifications and messages

StorageGRID event notification follows the Amazon S3 API with the following limitations:

- You cannot configure a notification for the following event types. These event types are **not** supported.
  - `s3:ReducedRedundancyLostObject`
  - `s3:ObjectRestore:Completed`
- Event notifications sent from StorageGRID use the standard JSON format except that they do not include some keys and use specific values for others, as shown in the table:

Key name	StorageGRID value
eventSource	sgws:s3
awsRegion	not included
x-amz-id-2	not included
arn	urn:sgws:s3:::bucket_name

### Related information

[Configuring event notifications](#)

### Understanding the 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 or its metadata is 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.

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.

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 cannot edit the document's field types in the index.

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 are not 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.

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 *Interoperability Matrix Tool* to determine the supported versions of Elasticsearch.

### Related information

[NetApp Interoperability Matrix Tool](#)

[Configuration XML for search integration](#)

[Object metadata included in metadata notifications](#)

[JSON generated by the search integration service](#)

[Configuring the search integration service](#)

### Considerations for using platform services

Before implementing platform services, review the recommendations and considerations for using these services.

#### 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.

Consideration	Details
Destination endpoint throttling	<p>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.</p> <p>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.</p> <p>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.</p>
Ordering guarantees	<p>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.</p> <p>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.</p>
ILM-driven object deletions	<p>To match the deletion behavior of the AWS CRR and SNS services, CloudMirror and event notification requests are not 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.</p> <p>In contrast, search integration requests are sent when objects are deleted because of ILM.</p>

#### Considerations for using the 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 TB, which is the same as the maximum object size supported by StorageGRID.

<p>Bucket versioning and version IDs</p>	<p>If the source S3 bucket in StorageGRID has versioning enabled, you should also enable versioning for the destination bucket.</p> <p>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.</p> <p><b>Note:</b> Version IDs for the source bucket in StorageGRID are not related to the version IDs for the destination bucket.</p>
<p>Tagging for object versions</p>	<p>The CloudMirror service does not replicate any PUT Object tagging or DELETE Object tagging requests that supply a version ID, due to limitations in the S3 protocol. Because version IDs for the source and destination are not related, there is no way to ensure that a tag update to a specific version ID will be replicated.</p> <p>In contrast, the CloudMirror service does replicate PUT Object tagging requests or DELETE Object tagging requests that do not 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.</p>
<p>Multipart uploads and ETag values</p>	<p>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.</p>
<p>Objects encrypted with SSE-C (server-side encryption with customer-provided keys)</p>	<p>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.</p>
<p>Bucket with S3 Object Lock enabled</p>	<p>If the destination S3 bucket for CloudMirror replication has S3 Object Lock enabled, the replication operation will fail with an AccessDenied error.</p>

**Related information**

[Use S3](#)

**Configuring 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. Contact your StorageGRID administrator for more information.

### **What a platform services endpoint is**

When you create a platform services endpoint, you specify the information that StorageGRID needs to access the external destination.

For example, if you want to replicate objects from a StorageGRID bucket to an S3 bucket, you create a platform services endpoint that includes the information and credentials StorageGRID needs to access the destination bucket on AWS.

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 SNS topic and notifications about object deletion to a second 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 Simple Notification Service (SNS) endpoints. Simple Queue Service (SQS) or AWS Lambda endpoints are not supported.

### **Endpoints for the search integration service**

StorageGRID supports search integration endpoints that represent Elasticsearch clusters. These Elasticsearch clusters can be in a local datacenter 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 do not 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](#)

### **Specifying the URN for a 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 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:mystore`, as follows:

- If the service is hosted on AWS, use `arn:aws`.
- If the service is hosted locally, use `urn:mystore`

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
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
Search integration	domain-name/index-name/type-name  <b>Note:</b> If the Elasticsearch cluster is <b>not</b> configured to create indexes automatically, you must create the index manually before you create the endpoint.

## URNs for services hosted on AWS

For AWS 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:mystore: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:

```
urn:mystore:sns:optional:optional:sns-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.

### Creating a platform services endpoint

You must create at least one endpoint of the correct type before you can enable a

platform service.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must belong to a user group that has the Manage Endpoints permission.
- The resource referenced by the platform services endpoint must have been created:
  - CloudMirror replication: S3 bucket
  - Event notification: SNS topic
  - Search notification: Elasticsearch index, if the destination cluster is not configured to automatically create indexes.
- You must 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)

#### [Specifying the URN for a platform services endpoint](#)

- Authentication credentials (if required):
  - Access Key: Access key ID and secret access key
  - Basic HTTP: Username and password
- Security certificate (if using a custom CA certificate)

### Steps

1. Select **STORAGE (S3) > Platform services endpoints**.

The Platform services endpoints page appears.

# 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.

0 endpoints

Create endpoint

Delete endpoint

	Display name ?	Last error ?	Type ?	URI ?	URN ?
No endpoints found					
<p>Create endpoint</p>					

2. Select **Create endpoint**.

# Create endpoint ✕

1 Enter details — 2 Select authentication type Optional — 3 Verify server Optional

## Enter endpoint details

Enter the endpoint's display name, URI, and URN.

Display name ?

URI ?

URN ?

[Cancel](#) [Continue](#)

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 do not 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 do not specify a port, port 443 is used for HTTPS URIs and port 80 is used for HTTP URIs.

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 a 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, do not 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 cannot change an endpoint's URN after the endpoint has been created.

6. Select **Continue**.

7. Select a value for **Authentication type**, and then enter the required credentials.

The screenshot shows a 'Create endpoint' wizard with three steps: 1. Enter details (checked), 2. Select authentication type (Optional), and 3. Verify server (Optional). The current step is 'Select authentication type', which asks the user to 'Select the method used to authenticate connections to the endpoint.' A dropdown menu is open, showing three options: 'Anonymous' (selected), 'Access Key', and 'Basic HTTP'. At the bottom right, there are 'Previous' and 'Continue' buttons.

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.	<ul style="list-style-type: none"> <li>• Access key ID</li> <li>• Secret access key</li> </ul>
Basic HTTP	Uses a username and password to authenticate connections to the destination.	<ul style="list-style-type: none"> <li>• Username</li> <li>• Password</li> </ul>

8. Select **Continue**.

9. Select a radio button for **Verify server** to choose how TLS connection to the endpoint is verified.

**Create endpoint**

1 Enter details ———— 2 Select authentication type Optional ———— **3** Verify server Optional

### Verify server

Use this method to validate the certificate for TLS connections to the endpoint resource. If you select "Use custom CA certificate," copy and paste the custom security certificate in the text box.

Use custom CA certificate
  Use operating system CA certificate
  Do not verify certificate

```

-----BEGIN CERTIFICATE-----
abcdefghijklmnopqrstuvwxyz123456780ABCDEFGHIJKL
123456/7890ABCDEFabcdefghijklmnopqrstuvwxyzABCD
-----END CERTIFICATE-----
  
```

Previous Test and create endpoint

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 <b>CA Certificate</b> text box.

Type of certificate verification	Description
Use operating system CA certificate	Use the default 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 are not 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

[Specifying the URN for a platform services endpoint](#)

[Configuring CloudMirror replication](#)

[Configuring event notifications](#)

[Configuring the search integration service](#)

#### Testing the connection for a 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.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Manage Endpoints 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.

# 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.

4 endpoints

Create endpoint

Delete endpoint

<input type="checkbox"/>	Display name <a href="#">?</a> <a href="#">↕</a>	Last error <a href="#">?</a> <a href="#">↕</a>	Type <a href="#">?</a> <a href="#">↕</a>	URI <a href="#">?</a> <a href="#">↕</a>	URN <a href="#">?</a> <a href="#">↕</a>
<input type="checkbox"/>	my-endpoint-1		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket1
<input type="checkbox"/>	my-endpoint-2	<span>✖</span> 2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es:::mydomain/sveloso/_doc
<input type="checkbox"/>	my-endpoint-3		Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1
<input type="checkbox"/>	my-endpoint-4		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket2

2. Select the endpoint whose connection you want to test.

The endpoint details page appears.

## Overview [^](#)

Display name: **my-endpoint-1** [✎](#)

Type: **S3 Bucket**

URI: **http://10.96.104.167:10443**

URN: **urn:sgws:s3:::bucket1**

---

**Connection** **Configuration**

### Verify connection [?](#)

Some errors might continue to appear after they are resolved. To see if an error is current or to force the removal of a resolved error, select **Test connection**.

**Test connection**

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**.

### Editing a 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 cannot change the URN for a platform services endpoint.

### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the Manage Endpoints 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.

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.

4 endpoints Create endpoint

Delete endpoint

<input type="checkbox"/>	Display name <sup>?</sup> <span>⬇</span>	Last error <sup>?</sup> <span>⬇</span>	Type <sup>?</sup> <span>⬇</span>	URI <sup>?</sup> <span>⬇</span>	URN <sup>?</sup> <span>⬇</span>
<input type="checkbox"/>	my-endpoint-1		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket1
<input type="checkbox"/>	my-endpoint-2	<span style="color: red;">✖</span> 2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es:::mydomain/sveloso/_doc
<input type="checkbox"/>	my-endpoint-3		Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1
<input type="checkbox"/>	my-endpoint-4		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket2

2. Select the endpoint you want to edit.

The endpoint details page appears.

3. Select **Configuration**.

## Overview

Display name: **my-endpoint-3** 

Type: **Notifications**

URI: **http://10.96.104.202:8080/**

URN: **arn:aws:sns:us-west-2::example1**

Connection

Configuration

## Edit configuration

### Endpoint details

URI 

http://10.96.104.202:8080/

URN 

arn:aws:sns:us-west-2::example1

Authentication type 

Basic HTTP 

Username 

testme

Password 

••••••••

Edit password

Verify server 

- Use custom CA certificate
- Use operating system CA certificate
- Do not verify certificate

```
-----BEGIN CERTIFICATE-----  
abcdefghijklmnop123456780ABCDEFGHIJKL  
123456/7890ABCDEFabcdefghijklmnop1ABCD  
-----END CERTIFICATE-----
```

Test and save changes

4. As needed, change the configuration of the endpoint.



You cannot 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 Basic HTTP authentication, change the username as needed. Change the password as needed by selecting **Edit password** and entering the new password. If you need to cancel your changes, select **Revert password edit**.
    - 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**.
  - 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**.

#### Related information

[Creating a platform services endpoint](#)

#### Deleting a platform services endpoint

You can delete an endpoint if you no longer want to use the associated platform service.

#### What you'll need

- You must be signed in to the Tenant Manager using a supported browser.
- You must belong to a user group that has the **Manage Endpoints** 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.

# 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.

4 endpoints

Create endpoint

Delete endpoint

<input type="checkbox"/>	Display name	Last error	Type	URI	URN
<input type="checkbox"/>	my-endpoint-1		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket1
<input type="checkbox"/>	my-endpoint-2	2 hours ago	Search	http://10.96.104.30:9200	urn:sgws:es:::mydomain/sveloso/_doc
<input type="checkbox"/>	my-endpoint-3		Notifications	http://10.96.104.202:8080/	arn:aws:sns:us-west-2::example1
<input type="checkbox"/>	my-endpoint-4		S3 Bucket	http://10.96.104.167:10443	urn:sgws:s3:::bucket2

2. Select the check box 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.

## Delete endpoint

**Are you sure you want to delete endpoint my-endpoint-10?**

This might take a few minutes.

When you delete an endpoint, you can no longer use it to access external resources.

#### 4. Select **Delete endpoint**.

##### Troubleshooting 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.

##### Determining if an 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.

## Overview ^

Display name:	<a href="#">my-endpoint-2</a>
Type:	Search
URI:	http://10.96.104.30:9200
URN:	urn:sgws:es:::mydomain/sveloso/_doc

Connection

Configuration

### Verify connection

Some errors might continue to appear after they are resolved. To see if an error is current or to force the removal of a resolved error, select **Test connection**.

Test connection

#### Last error details

2 hours ago

Endpoint failure: Endpoint has an AWS failure: RequestError: send request failed; caused by: url.Error; caused by: net:OpError; caused by: os.SyscallError (logID: 143H5UDUUKMGDRWJ)

### Checking if an 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.

### Resolving 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

1378

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 “InvalidAccessKeyld.”

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.

### Additional platform services troubleshooting

For additional information about troubleshooting platform services, see the instructions for administering StorageGRID.

[Administer StorageGRID](#)

### Related information

[Creating a platform services endpoint](#)

[Testing the connection for a platform services endpoint](#)

[Editing a platform services endpoint](#)

### Configuring CloudMirror replication

The CloudMirror replication service is one of the three StorageGRID platform services. You can use CloudMirror replication to automatically replicate objects to an external S3 bucket.

### What you’ll need

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already created a bucket to act as the replication source.
- The endpoint that you intend to use as a destination for CloudMirror replication must already exist, and you must have its URN.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission, which allows you to manage the settings for all S3 buckets in your tenant account. 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. To enable CloudMirror replication for a bucket, you must create and apply valid bucket replication configuration XML. The replication configuration XML 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.

For general information on bucket replication and how to configure it, see the Amazon documentation on cross-region replication (CRR). For information on how StorageGRID implements the S3 bucket replication configuration API, see the instructions for implementing S3 client applications.

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 are not. 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 do not 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. 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 > Replication**.

5. Select the **Enable replication** check box.

6. Paste the replication configuration XML into the text box, and select **Save changes**.

Bucket options
Bucket access
Platform services

**Replication**
Disabled
▲

Enable the CloudMirror replication service to copy objects from a source bucket to a destination bucket that is specified in an endpoint.

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already configured an endpoint for each destination bucket.
- You must specify the URN of each endpoint in the replication configuration XML for the source bucket.

Enable replication

Clear

```

<ReplicationConfiguration>
  <Role></Role>
  <Rule>
    <Status>Enabled</Status>
    <Prefix>2020</Prefix>
    <Destination>
      <Bucket>urn:sgws:s3:::2017-records</Bucket>
      <StorageClass>STANDARD</StorageClass>
    </Destination>
  </Rule>
</ReplicationConfiguration>

```

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

[Understanding the CloudMirror replication service](#)

[Use S3](#)

[Creating a platform services endpoint](#)

## Configuring event notifications

The notifications service is one of the three StorageGRID platform services. You can enable notifications for a bucket to send information about specified events to a destination service that supports the AWS Simple Notification Service™ (SNS).

### What you'll need

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must 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 must already exist, and you must have its URN.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission, which allows you to manage the settings for all S3 buckets in your tenant account. 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 event notifications, whenever a specified event occurs for an object in the source bucket, a notification is generated and sent to the Simple Notification Service (SNS) topic used as the destination endpoint. To enable notifications for a bucket, you must create and apply valid notification configuration XML. The notification configuration XML must use the URN of an event notifications endpoint for each destination.

For general information on event notifications and how to configure them, see Amazon documentation. For information on how StorageGRID implements the S3 bucket notification configuration API, see the instructions for implementing S3 client applications.

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** check box.
6. Paste the notification configuration XML into the text box, and select **Save changes**.

Bucket options
Bucket access
Platform services

Replication
Disabled
▼

Event notifications
Disabled
▲

Enable the event notification service for an S3 bucket if you want StorageGRID to send notifications about specified events to a destination Amazon Simple Notification Service (SNS).

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already configured an endpoint for the destination of event notifications.
- You must specify the URN of that endpoint in the notification configuration XML for the source bucket.

Enable event notifications

Clear

```

<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>
      
```

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 SNS topic.

For example, if your destination topic is hosted on the AWS Simple Notification Service (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

[Understanding notifications for buckets](#)

[Use S3](#)

[Creating a platform services endpoint](#)

## Using the search integration service

The search integration service is one of the three StorageGRID platform services. You can enable this service to send object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

You can configure search integration by using the Tenant Manager to apply custom StorageGRID configuration XML to a bucket.



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 than the *notification configuration XML* used to enable event notifications.

See the instructions for implementing S3 client applications for details about the following custom StorageGRID S3 REST API operations:

- DELETE Bucket metadata notification configuration request
- GET Bucket metadata notification configuration request
- PUT Bucket metadata notification configuration request

## Related information

[Configuration XML for search integration](#)

[Object metadata included in metadata notifications](#)

[JSON generated by the search integration service](#)

[Configuring the search integration service](#)

[Use S3](#)

## Configuration XML for search integration

The search integration service is configured using a set of rules contained within `<MetadataNotificationConfiguration>` and `</MetadataNotificationConfiguration>` tags. Each rule specifies the objects that the rule applies to, and the destination where StorageGRID should send those objects' 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 metadata for objects with the prefix `/videos` to another. Configurations that have overlapping prefixes are not valid, and are rejected when they are 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.

Destinations must be specified using the URN of a StorageGRID endpoint that has been created for the search integration service. These endpoints refer to an index and type defined on an Elasticsearch cluster.

```

<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
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

Name	Description	Required
Prefix	<p>Objects that match the prefix are affected by the rule, and their metadata is sent to the specified destination.</p> <p>To match all objects, specify an empty prefix.</p> <p>Included in the Rule element.</p>	Yes
Destination	<p>Container tag for the destination of a rule.</p> <p>Included in the Rule element.</p>	Yes
Urn	<p>URN of the destination where object metadata is sent. Must be the URN of a StorageGRID endpoint with the following properties:</p> <ul style="list-style-type: none"> <li>• es must be the third element.</li> <li>• The URN must end with the index and type where the metadata is stored, in the form domain-name/myindex/mytype.</li> </ul> <p>Endpoints are configured using the Tenant Manager or Tenant Management API. They take the following form:</p> <ul style="list-style-type: none"> <li>• arn:aws:es:region:account-ID:domain/mydomain/myindex/mytype</li> <li>• urn:mysite:es:::mydomain/myindex/mytype</li> </ul> <p>The endpoint must be configured before the configuration XML is submitted, or configuration will fail with a 404 error.</p> <p>Urn is included in the Destination element.</p>	Yes

Use the sample metadata notification configuration XML to learn how to construct your own XML.

### 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>

```

### 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-1:3333333: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-1:2222222:domain/es-
domain/graphics/videotype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

### Related information

[Use S3](#)

[JSON generated by the search integration service](#)

[Configuring the search integration service](#)

## Configuring the search integration service

The search integration service sends object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

### What you'll need

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must 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 must already exist, and you must have its URN.
- You must belong to a user group that has the Manage All Buckets or the Root Access permission, which allows you to manage the settings for all S3 buckets in your tenant account. 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 are not automatically sent for existing objects. You must update these existing objects to ensure that their metadata is added to the destination search index.

### Steps

1. Use a text editor to create the metadata notification XML required to enable search integration.
  - See the information about configuration XML for search integration.
  - When configuring the XML, use the URN of a search integration endpoint as the destination.

```
<MetadataNotificationConfiguration>
  <Rule>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-
1:111111111111:domain/mydomain/myindex/mytype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

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** check box.
6. Paste the metadata notification configuration into the text box, and select **Save changes**.

Bucket options
Bucket access
Platform services

Replication
Disabled
▼

Event notifications
Disabled
▼

Search integration
Disabled
▲

Enable the search integration service to send object metadata to a destination search index whenever an object is created, deleted, or its metadata or tags are updated.

- Platform services must be enabled for your tenant account by a StorageGRID administrator.
- You must have already configured an endpoint for the search integration service.
- You must specify the URN of that endpoint in the search integration configuration XML for the bucket you want to index.

Enable search integration

Clear

```

<MetadataNotificationConfiguration>
  <Rule>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-1:111111111111:domain/mydomain/myindex/mytype</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

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:
  - a. 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.

- b. 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 unselect the **Enable search integration** check box.
- If you are using the S3 API directly, use a DELETE Bucket metadata notification request. See the instructions for implementing S3 client applications.

## Related information

[Understanding the search integration service](#)

[Configuration XML for search integration](#)

[Use S3](#)

[Creating a platform services endpoint](#)

## 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 and description
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



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 cannot edit the document's field types in the index.

## Use S3

Learn how client applications can use the S3 API to interface with the StorageGRID system.

- [Support for the S3 REST API](#)
- [Configuring tenant accounts and connections](#)
- [How StorageGRID implements the S3 REST API](#)
- [S3 REST API supported operations and limitations](#)
- [StorageGRID S3 REST API operations](#)
- [Bucket and group access policies](#)
- [Configuring security for the REST API](#)
- [Monitoring and auditing operations](#)
- [Benefits of active, idle, and concurrent HTTP connections](#)

## Support for the S3 REST API

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-premises object storage that uses the StorageGRID system. This requires minimal changes to a client application's current use of S3 REST API calls.

- [Changes to S3 REST API support](#)
- [Supported versions](#)
- [Support for StorageGRID platform services](#)

### Changes to S3 REST API support

You should be aware of changes to the StorageGRID system's support for the S3 REST API.

Release	Comments
11.5	<ul style="list-style-type: none"><li>• Added support for managing bucket encryption.</li><li>• Added support for S3 Object Lock and deprecated legacy Compliance requests.</li><li>• Added support for using DELETE Multiple Objects on versioned buckets.</li><li>• The <code>Content-MD5</code> request header is now correctly supported.</li></ul>
11.4	<ul style="list-style-type: none"><li>• Added support for DELETE Bucket tagging, GET Bucket tagging, and PUT Bucket tagging. Cost allocation tags are not supported.</li><li>• For buckets created in StorageGRID 11.4, restricting object key names to meet performance best practices is no longer required.</li><li>• Added support for bucket notifications on the <code>s3:ObjectRestore:Post</code> event type.</li><li>• 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.</li><li>• Added support for TLS 1.3, and updated list of supported TLS cipher suites.</li><li>• The CLB service is deprecated.</li></ul>

Release	Comments
11.3	<ul style="list-style-type: none"> <li>• Added support for server-side encryption of object data with customer-provided keys (SSE-C).</li> <li>• Added support for DELETE, GET, and PUT Bucket lifecycle operations (Expiration action only) and for the <code>x-amz-expiration</code> response header.</li> <li>• Updated PUT Object, PUT Object - Copy, and Multipart Upload to describe the impact of ILM rules that use synchronous placement at ingest.</li> <li>• Updated list of supported TLS cipher suites. TLS 1.1 ciphers are no longer supported.</li> </ul>
11.2	<p>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.</p> <p><b>Note:</b> Uses of ARN/URN in other configuration JSON/XML, including those used in custom StorageGRID features, have not changed.</p>
11.1	<p>Added support for Cross-Origin Resource Sharing (CORS), HTTP for S3 client connections to grid nodes, and compliance settings on buckets.</p>
11.0	<p>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 control setting.</p>
10.4	<p>Added support for ILM scanning changes to versioning, Endpoint Domain Names page updates, conditions and variables in policies, policy examples, and the PutOverwriteObject permission.</p>
10.3	<p>Added support for versioning.</p>
10.2	<p>Added support for group and bucket access policies, and for multipart copy (Upload Part - Copy).</p>
10.1	<p>Added support for multipart upload, virtual hosted-style requests, and v4 authentication.</p>

Release	Comments
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.

### Supported versions

StorageGRID supports the following specific versions of S3 and HTTP.

Item	Version
S3 specification	<i>Simple Storage Service API Reference</i> 2006-03-01
HTTP	1.1  For more information about HTTP, see HTTP/1.1 (RFCs 7230-35).  <b>Note:</b> StorageGRID does not support HTTP/1.1 pipelining.

### Related information

[IETF RFC 2616: Hypertext Transfer Protocol \(HTTP/1.1\)](#)

[Amazon Web Services \(AWS\) Documentation: Amazon Simple Storage Service API Reference](#)

### Support for StorageGRID platform services

StorageGRID platform services enable StorageGRID tenant accounts to leverage external services such as a remote S3 bucket, a Simple Notification Service (SNS) endpoint, or an Elasticsearch cluster to extend the services provided by a grid.

The following table summarizes the available platform services and the S3 APIs used to configure them.

Platform service	Purpose	S3 API used to configure the service
CloudMirror replication	Replicates objects from a source StorageGRID bucket to the configured remote S3 bucket.	PUT Bucket replication
Notifications	Sends notifications about events in a source StorageGRID bucket to a configured Simple Notification Service (SNS) endpoint.	PUT Bucket notification

Platform service	Purpose	S3 API used to configure the service
Search integration	Sends object metadata for objects stored in a StorageGRID bucket to a configured Elasticsearch index.	PUT Bucket metadata notification  <b>Note:</b> This is a StorageGRID custom S3 API.

A grid administrator must enable the use of platform services for a tenant account before they can be used. Then, a tenant administrator must create an endpoint that represents the remote service in the tenant account. This step is required before a service can be configured.

### Recommendations for using platform services

Before using platform services, you must be aware of the following recommendations:

- NetApp recommends that you allow no 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.
- If an S3 bucket in the StorageGRID system has both versioning and CloudMirror replication enabled, NetApp recommends that the destination endpoint also have S3 bucket versioning enabled. This allows CloudMirror replication to generate similar object versions on the endpoint.
- CloudMirror replication is not supported if the source bucket has S3 Object Lock enabled.
- CloudMirror replication will fail with an AccessDenied error if the destination bucket has legacy Compliance enabled.

### Related information

[Use a tenant account](#)

[Administer StorageGRID](#)

[Operations on buckets](#)

[PUT Bucket metadata notification configuration request](#)

## Configuring tenant accounts and connections

Configuring StorageGRID to accept connections from client applications requires creating one or more tenant accounts and setting up the connections.

### Creating and configuring S3 tenant accounts

An S3 tenant account is required before S3 API clients can store and retrieve objects on StorageGRID. Each tenant account has its own account ID, groups and users, and containers and objects.

S3 tenant accounts are created by a StorageGRID grid administrator using the Grid Manager or the Grid Management API. When creating an S3 tenant account, the grid administrator specifies the following information:

- Display name for the tenant (the tenant's account ID is assigned automatically and cannot be changed).
- Whether the tenant account is allowed to use platform services. If the use of platform services is allowed,

the grid must be configured to support their use.

- Optionally, a storage quota for the tenant account—the maximum number of gigabytes, terabytes, or petabytes available for the tenant's objects. A tenant's storage quota represents a logical amount (object size), not a physical amount (size on disk).
- If identity federation is enabled for the StorageGRID system, which federated group has Root Access permission to configure the tenant account.
- If single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid's identity source, and the initial password for the tenant's local root user.

After an S3 tenant account is created, tenant users 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), and create local groups and users
- Manage S3 access keys
- Create and manage S3 buckets, including buckets that have S3 Object Lock enabled
- Use platform services (if enabled)
- Monitor storage usage



S3 tenant users can create and manage S3 buckets with the Tenant Manager, but they must have S3 access keys and use the S3 REST API to ingest and manage objects.

#### Related information

[Administer StorageGRID](#)

[Use a tenant account](#)

#### How client connections can be configured

A grid administrator makes configuration choices that affect how S3 clients connect to StorageGRID to store and retrieve data. The specific information you need to make a connection depends upon the configuration that was chosen.

Client applications can store or retrieve objects by connecting to any of the following:

- The Load Balancer service on Admin Nodes or Gateway Nodes, or optionally, the virtual IP address of a high availability (HA) group of Admin Nodes or Gateway Nodes
- The CLB service on Gateway Nodes, or optionally, the virtual IP address of a high availability group of Gateway Nodes



The CLB service is deprecated. Clients configured before the StorageGRID 11.3 release can continue to use the CLB service on Gateway Nodes. All other client applications that depend on StorageGRID to provide load balancing should connect using the Load Balancer service.

- Storage Nodes, with or without an external load balancer

When configuring StorageGRID, a grid administrator can use the Grid Manager or the Grid Management API to perform the following steps, all of which are optional:

## 1. Configure endpoints for the Load Balancer service.

You must configure endpoints to use the Load Balancer service. The Load Balancer service on Admin Nodes or Gateway Nodes distributes incoming network connections from client applications to Storage Nodes. When creating a load balancer endpoint, the StorageGRID administrator specifies a port number, whether the endpoint accepts HTTP or HTTPS connections, the type of client (S3 or Swift) that will use the endpoint, and the certificate to be used for HTTPS connections (if applicable).

## 2. Configure Untrusted Client Networks.

If a StorageGRID administrator configures a node's Client Network to be untrusted, the node only accepts inbound connections on the Client Network on ports that are explicitly configured as load balancer endpoints.

## 3. Configure high availability groups.

If an administrator creates an HA group, the network interfaces of multiple Admin Nodes or Gateway Nodes are placed into an active-backup configuration. Client connections are made using the virtual IP address of the HA group.

For more information about each option, see the instructions for administering StorageGRID.

### Related information

[Administer StorageGRID](#)

### Summary: IP addresses and ports for client connections

Client applications connect to StorageGRID using the IP address of a grid node and the port number of a service on that node. If high availability (HA) groups are configured, client applications can connect using the virtual IP address of the HA group.

### Information required to make client connections

The 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. Contact your StorageGRID administrator for more information, or see the instructions for administering StorageGRID for a description of how to find this information 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	<ul style="list-style-type: none"><li>• Load balancer endpoint port</li></ul>
HA group	CLB <b>Note:</b> The CLB service is deprecated.	Virtual IP address of an HA group	Default S3 ports: <ul style="list-style-type: none"><li>• HTTPS: 8082</li><li>• HTTP: 8084</li></ul>
Admin Node	Load Balancer	IP address of the Admin Node	<ul style="list-style-type: none"><li>• Load balancer endpoint port</li></ul>

Where connection is made	Service that client connects to	IP address	Port
Gateway Node	Load Balancer	IP address of the Gateway Node	<ul style="list-style-type: none"> <li>• Load balancer endpoint port</li> </ul>
Gateway Node	CLB  <b>Note:</b> The CLB service is deprecated.	IP address of the Gateway Node  <b>Note:</b> By default, HTTP ports for CLB and LDR are not enabled.	Default S3 ports: <ul style="list-style-type: none"> <li>• HTTPS: 8082</li> <li>• HTTP: 8084</li> </ul>
Storage Node	LDR	IP address of Storage Node	Default S3 ports: <ul style="list-style-type: none"> <li>• HTTPS: 18082</li> <li>• HTTP: 18084</li> </ul>

### Example

To connect an S3 client 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 an S3 Load Balancer endpoint is 10443, then an S3 client could use the following URL to connect to StorageGRID:

- `https://192.0.2.5:10443`

It is possible to configure a DNS name for the IP address that clients use to connect to StorageGRID. Contact your local network administrator.

### Related information

[Administer StorageGRID](#)

### Deciding to use HTTPS or HTTP connections

When client connections are made using a Load Balancer endpoint, connections must be made using the protocol (HTTP or HTTPS) that was specified for that endpoint. To use HTTP for client connections to Storage Nodes or to the CLB service on Gateway Nodes, you must enable its use.

By default, when client applications connect to Storage Nodes or the CLB service on Gateway Nodes, they must use encrypted HTTPS for all connections. Optionally, you can enable less-secure HTTP connections by selecting the **Enable HTTP Connection** grid option in the Grid Manager. For example, a client application might use HTTP when testing the connection to a Storage Node in a non-production environment.



Be careful when enabling HTTP for a production grid since requests will be sent unencrypted.



The CLB service is deprecated.

If the **Enable HTTP Connection** option is selected, clients must use different ports for HTTP than they use for HTTPS. See the instructions for administering StorageGRID.

### Related information

[Administer StorageGRID](#)

[Benefits of active, idle, and concurrent HTTP connections](#)

### Endpoint domain names for S3 requests

Before you can use S3 domain names for client requests, a StorageGRID administrator must configure the system to accept connections that use S3 domain names in S3 path-style and S3 virtual hosted-style requests.

#### About this task

To enable you to use S3 virtual hosted style-requests, a grid administrator must perform the following tasks:

- 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`, the grid administrator 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 to include DNS records that match the endpoint domain names, including any required wildcard records.

If the client connects using the Load Balancer service, the certificate that the grid administrator configures is the certificate for the load balancer endpoint that the client uses.



Each load balancer endpoint has its own certificate, and each endpoint can be configured to recognize different endpoint domain names.

If the client connects Storage Nodes or to the CLB service on Gateway Nodes, the certificate that the grid administrator configures is the single custom server certificate used for the grid.



The CLB service is deprecated.

See the instructions for administering StorageGRID for more information.

After these steps have been completed, you can use virtual hosted-style requests (for example, `bucket.s3.company.com`).

### Related information

[Administer StorageGRID](#)

[Configuring security for the REST API](#)

### Testing your 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 to the system.

## What you'll need

- You must have downloaded and installed the AWS CLI from [aws.amazon.com/cli](https://aws.amazon.com/cli).
- You must have created an S3 tenant account in the StorageGRID system.

## Steps

1. Configure the Amazon Web Services settings to use the account you created in the StorageGRID system:
  - a. Enter configuration mode: `aws configure`
  - b. Enter the AWS Access Key ID for the account you created.
  - c. Enter the AWS 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.

```
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 the S3 REST API

A client application can use S3 REST API calls to connect to StorageGRID to create, delete, and modify buckets, as well as storing and retrieving objects.

- [Conflicting client requests](#)
- [Consistency controls](#)
- [How StorageGRID ILM rules manage objects](#)
- [Object versioning](#)
- [Recommendations for implementing the 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 controls

Consistency controls provide a trade-off between the availability of the objects and the consistency of those objects across different Storage Nodes and sites, 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 level, you can specify a consistency control for each bucket or for each API operation.

### Consistency controls

The consistency control 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 control for a bucket or an API operation to one of the following values:

Consistency control	Description
all	All nodes receive the data immediately, or the request will fail.

Consistency control	Description
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	<p>(Default) Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Matches Amazon S3 consistency guarantees.</p> <p><b>Note:</b> If your application uses HEAD requests on objects that do not exist, you might receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable. To prevent these errors, set the consistency control to “available” unless you require consistency guarantees similar to Amazon S3.</p>
available (eventual consistency for HEAD operations)	Behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than “read-after-new-write” if Storage Nodes are unavailable. Differs from Amazon S3 consistency guarantees for HEAD operations only.

### Using the “read-after-new-write” and “available” consistency controls

When a HEAD or GET operation uses the “read-after-new-write” consistency control or a GET operation uses the “available” consistency control, 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 level until it reaches the highest consistency level, “all,” which requires all copies of the object metadata to be available.

If a HEAD or GET operation uses the “read-after-new-write” consistency control but the object does not exist, the object lookup will always reach the “all” consistency level. Because this consistency level requires all copies of the object metadata to be available, you can receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable.

Unless you require consistency guarantees similar to Amazon S3, you can prevent these errors for HEAD operations by setting the consistency control to “available.” When a HEAD operation uses the “available” consistency control, StorageGRID provides eventual consistency only. It does not retry a failed operation until it reaches the “all” consistency level, so it does not require that all copies of the object metadata be available.

### Specifying the consistency control for an API operation

To set the consistency control for an individual API operation, consistency controls must be supported for the

operation, and you must specify the consistency control in the request header. This example sets the consistency control to “strong-site” for a GET Object operation.

```
GET /bucket/object HTTP/1.1
Date: date
Authorization: <em>authorization name</em>
Host: <em>host</em>
Consistency-Control: strong-site
```



You must use the same consistency control for both the PUT Object and GET Object operations.

### Specifying the consistency control for a bucket

To set the consistency control for bucket, you can use the StorageGRID PUT Bucket consistency request and the GET Bucket consistency request. Or you can use the Tenant Manager or the Tenant Management API.

When setting the consistency controls for a bucket, be aware of the following:

- Setting the consistency control for a bucket determines which consistency control 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 control for an individual API operation overrides the consistency control for the bucket.
- In general, buckets should use the default consistency control, “read-after-new-write.” If requests are not working correctly, change the application client behavior if possible. Or, configure the client to specify the consistency control for each API request. Set the consistency control at the bucket level only as a last resort.

### How consistency controls and ILM rules interact to affect data protection

Both your choice of consistency control and your ILM rule affect how objects are protected. These settings can interact.

For example, the consistency control 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 level and ingest behavior can provide better initial data protection and more predictable system responses.

The following ingest behaviors are available for ILM rules:

- **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.
- **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.



Before selecting the ingest behavior for an ILM rule, read the full description of these settings in the instructions for managing objects with information lifecycle management.

## Example of how the consistency control and ILM rule can interact

Suppose you have a two-site grid with the following ILM rule and the following consistency level setting:

- **ILM rule:** Create two object copies, one at the local site and one at a remote site. The Strict ingest behavior is selected.
- **Consistency level:** “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 level, 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 cannot be retrieved.

The inter-relationship between consistency levels and ILM rules can be complex. Contact NetApp if you require assistance.

### Related information

[Manage objects with ILM](#)

[GET Bucket consistency request](#)

[PUT Bucket consistency request](#)

## How StorageGRID ILM rules manage objects

The grid administrator creates information lifecycle management (ILM) rules to manage object data ingested into the StorageGRID system from S3 REST API client applications. These rules are then added to the ILM policy to determine how and where object data is stored over time.

ILM settings determine the following aspects of an object:

- **Geography**

The location of an object’s data, either within the StorageGRID system (storage pool) or in a Cloud Storage Pool.

- **Storage grade**

The type of storage used to store object data: for example flash or spinning disk.

- **Loss protection**

How many copies are made and the types of copies that are created: replication, erasure coding, or both.

- **Retention**

The changes over time to how an object's data is managed, where it is stored, and how it is protected from loss.

- **Protection during ingest**

The method used to protect object data during ingest: synchronous placement (using the Balanced or Strict options for Ingest Behavior), or making interim copies (using the Dual commit option).

ILM rules can filter and select objects. For objects ingested using S3, ILM rules can filter objects based on the following metadata:

- Tenant Account
- Bucket Name
- Ingest Time
- Key
- Last Access Time



By default, updates to last access time are disabled for all S3 buckets. If your StorageGRID system includes an ILM rule that uses the Last Access Time option, you must enable updates to last access time for the S3 buckets specified in that rule. You can enable last access time updates using the PUT Bucket last access time request, the **S3 > Buckets > Configure Last Access Time** check box in the Tenant Manager, or using the Tenant Management API. When enabling last access time updates, be aware that StorageGRID performance might be reduced, especially in systems with small objects.

- Location Constraint
- Object Size
- User Metadata
- Object Tag

For more information about ILM, see the instructions for managing objects with information lifecycle management.

#### **Related information**

[Use a tenant account](#)

[Manage objects with ILM](#)

[PUT Bucket last access time request](#)

#### **Object versioning**

You can use versioning to retain multiple versions of an object, which protects 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 1,000 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 to turn on this functionality for the bucket. Each object in your 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. 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.

See the instructions for managing objects with information lifecycle management for an example ILM policy for S3 versioned objects.

### Related information

[Manage objects with ILM](#)

### Recommendations for implementing the 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 do not expect the object to actually exist, you should use the “Available” consistency control. For example, you should use the “Available” consistency control 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 one or more Storage Nodes are unavailable.

You can set the “Available” consistency control for each bucket using the PUT Bucket consistency request, or you can specify the consistency control in the request header for an individual API operation.

#### Recommendations for object keys

For buckets that are created in StorageGRID 11.4 or later, restricting object key names to meet performance best practices is no longer required. For example, you can now use random values for the first four characters of object key names.

For buckets that were created in releases earlier than StorageGRID 11.4, continue to follow these recommendations for object key names:

- You should not use random values as the first four characters of object keys. This is in contrast to the former AWS recommendation for key prefixes. Instead, you should 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, you should 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
```

### Recommendations for “range reads”

If the **Compress Stored Objects** option is selected (**Configuration > Grid Options**), S3 client applications should avoid performing GET Object operations that specify a range of bytes be returned. These “range read” operations are inefficient because StorageGRID must effectively uncompress the objects to access the requested bytes. GET Object operations that request a small range of bytes from a very large object are especially inefficient; for example, it is very 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.

### Related information

[Consistency controls](#)

[PUT Bucket consistency request](#)

[Administer StorageGRID](#)

## S3 REST API supported operations and limitations

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.

- [Authenticating requests](#)
- [Operations on the service](#)
- [Operations on buckets](#)

- [Custom operations on buckets](#)
- [Operations on objects](#)
- [Operations for multipart uploads](#)
- [Error responses](#)

## 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 common request headers defined by the *Simple Storage Service API Reference*, with one exception.

Request header	Implementation
Authorization	<p>Full support for AWS Signature Version 2</p> <p>Support for AWS Signature Version 4, with the following exceptions:</p> <ul style="list-style-type: none"> <li>• The SHA256 value is not calculated for the body of the request. The user-submitted value is accepted without validation, as if the value <code>UNSIGNED-PAYLOAD</code> had been provided for the <code>x-amz-content-sha256</code> header.</li> </ul>
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

## Related information

[Amazon Web Services \(AWS\) Documentation: Amazon Simple Storage Service API Reference](#)

## Authenticating 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.

### Using 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.

### Using 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 do not need to know the secret access key in order 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
GET Service	Implemented with all Amazon S3 REST API behavior.
GET Storage Usage	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. This is an operation on the service with a path of <code>/</code> and a custom query parameter ( <code>?x-ntap-sg-usage</code> ) added.
OPTIONS /	Client applications can issue <code>OPTIONS /</code> 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.

### Related information

[GET Storage Usage request](#)

## Operations on buckets

The StorageGRID system supports a maximum of 1,000 buckets for each S3 tenant

account.

Bucket name restrictions follow the AWS US Standard region restrictions, but you should further restrict them to DNS naming conventions in order to support S3 virtual hosted-style requests.

[Amazon Web Services \(AWS\) Documentation: Bucket Restrictions and Limitations](#)

### Endpoint domain names for S3 request

The GET Bucket (List Objects) and GET Bucket versions operations support StorageGRID consistency controls.

You can check whether updates to last access time are enabled or disabled for individual buckets.

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
DELETE Bucket	Implemented with all Amazon S3 REST API behavior.
DELETE Bucket cors	This operation deletes the CORS configuration for the bucket.
DELETE Bucket encryption	This operation deletes the default encryption from the bucket. Existing encrypted objects remain encrypted, but any new objects added to the bucket are not encrypted.
DELETE Bucket lifecycle	This operation deletes the lifecycle configuration from the bucket.
DELETE Bucket policy	This operation deletes the policy attached to the bucket.
DELETE Bucket replication	This operation deletes the replication configuration attached to the bucket.
DELETE Bucket tagging	This operation uses the <code>tagging</code> subresource to remove all tags from a bucket.

Operation	Implementation
GET Bucket (List Objects), version 1 and version 2	<p>This operation 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:</p> <ul style="list-style-type: none"> <li>• <code>STANDARD</code>, which indicates the object is stored in a storage pool consisting of Storage Nodes.</li> <li>• <code>GLACIER</code>, which indicates that the object has been moved to the external bucket specified by the Cloud Storage Pool.</li> </ul> <p>If the bucket contains large numbers of deleted keys that have the same prefix, the response might include some <code>CommonPrefixes</code> that do not contain keys.</p>
GET Bucket acl	This operation returns a positive response and the ID, <code>DisplayName</code> , and <code>Permission</code> of the bucket owner, indicating that the owner has full access to the bucket.
GET Bucket cors	This operation returns the <code>cors</code> configuration for the bucket.
GET Bucket encryption	This operation returns the default encryption configuration for the bucket.
GET Bucket lifecycle	This operation returns the lifecycle configuration for the bucket.
GET Bucket location	This operation returns the region that was set using the <code>LocationConstraint</code> element in the <code>PUT Bucket</code> request. If the bucket's region is <code>us-east-1</code> , an empty string is returned for the region.
GET Bucket notification	This operation returns the notification configuration attached to the bucket.
GET Bucket Object versions	With <code>READ</code> access on a bucket, this operation with the <code>versions</code> subresource lists metadata of all of the versions of objects in the bucket.
GET Bucket policy	This operation returns the policy attached to the bucket.
GET Bucket replication	This operation returns the replication configuration attached to the bucket.

Operation	Implementation
GET Bucket tagging	This operation uses the <code>tagging</code> subresource to return all tags for a bucket.
GET Bucket versioning	This implementation uses the <code>versioning</code> subresource to return the versioning state of a bucket. The versioning state returned indicates if the bucket is “Unversioned” or if the bucket is version “Enabled” or “Suspended.”
GET Object Lock Configuration	This operation determines if S3 Object Lock is enabled for a bucket. <a href="#">Using S3 Object Lock</a>
HEAD Bucket	This operation determines if a bucket exists and you have permission to access it.

Operation	Implementation
PUT Bucket	<p>This operation creates a new bucket. By creating the bucket, you become the bucket owner.</p> <ul style="list-style-type: none"> <li>• Bucket names must comply with the following rules: <ul style="list-style-type: none"> <li>◦ Must be unique across each StorageGRID system (not just unique within the tenant account).</li> <li>◦ Must be DNS compliant.</li> <li>◦ Must contain at least 3 and no more than 63 characters.</li> <li>◦ 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.</li> <li>◦ Must not look like a text-formatted IP address.</li> <li>◦ Should not use periods in virtual hosted style requests. Periods will cause problems with server wildcard certificate verification.</li> </ul> </li> <li>• 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 do not know the region name you should use. <p><b>Note:</b> An error will occur if your PUT Bucket request uses a region that has not been defined in StorageGRID.</p> </li> <li>• You can include the <code>x-amz-bucket-object-lock-enabled</code> request header to create a bucket with S3 Object Lock enabled.</li> </ul> <p>You must enable S3 Object Lock when you create the bucket. You cannot 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.</p> <p><a href="#">Using S3 Object Lock</a></p>

Operation	Implementation
PUT Bucket cors	<p>This operation 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>.</p>
PUT Bucket encryption	<p>This operation 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 do not use the <code>KMSMasterKeyID</code> parameter.</p> <p>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).</p>

Operation	Implementation
PUT Bucket lifecycle	<p>This operation 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:</p> <ul style="list-style-type: none"> <li>• Expiration (Days, Date)</li> <li>• NoncurrentVersionExpiration (NoncurrentDays)</li> <li>• Filter (Prefix, Tag)</li> <li>• Status</li> <li>• ID</li> </ul> <p>StorageGRID does not support these actions:</p> <ul style="list-style-type: none"> <li>• AbortIncompleteMultipartUpload</li> <li>• ExpiredObjectDeleteMarker</li> <li>• Transition</li> </ul> <p>To understand how the Expiration action in a bucket lifecycle interacts with ILM placement instructions, see “How ILM operates throughout an object’s life” in the instructions for managing objects with information lifecycle management.</p> <p><b>Note:</b> 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.</p>

Operation	Implementation
PUT Bucket notification	<p>This operation configures notifications for the bucket using the notification configuration XML included in the request body. You should be aware of the following implementation details:</p> <ul style="list-style-type: none"> <li>• StorageGRID supports Simple Notification Service (SNS) topics as destinations. Simple Queue Service (SQS) or Amazon Lambda endpoints are not supported.</li> <li>• 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.</li> </ul> <p>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>.</p> <ul style="list-style-type: none"> <li>• You cannot configure a notification for the following event types. These event types are <b>not</b> supported. <ul style="list-style-type: none"> <li>◦ <code>s3:ReducedRedundancyLostObject</code></li> <li>◦ <code>s3:ObjectRestore:Completed</code></li> </ul> </li> <li>• Event notifications sent from StorageGRID use the standard JSON format except that they do not include some keys and use specific values for others, as shown in the following listing:</li> <li>• <b>eventSource</b> <pre>sgws:s3</pre> </li> <li>• <b>awsRegion</b> <pre>not included</pre> </li> <li>• <b>x-amz-id-2</b> <pre>not included</pre> </li> <li>• <b>arn</b> <pre>urn:sgws:s3:::bucket_name</pre> </li> </ul>
PUT Bucket policy	This operation sets the policy attached to the bucket.

Operation	Implementation
PUT Bucket replication	<p>This operation 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:</p> <ul style="list-style-type: none"> <li>• 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. See the Amazon documentation on replication configuration for details.</li> <li>• Bucket replication can be configured on versioned or unversioned buckets.</li> <li>• 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.</li> <li>• Destination buckets must be specified as the URN of StorageGRID endpoints as specified in the Tenant Manager or the Tenant Management API.</li> </ul> <p>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></p> <ul style="list-style-type: none"> <li>• You do not need to specify a <code>Role</code> in the configuration XML. This value is not used by StorageGRID and will be ignored if submitted.</li> <li>• If you omit the storage class from the configuration XML, StorageGRID uses the <code>STANDARD</code> storage class by default.</li> <li>• If you delete an object from the source bucket or you delete the source bucket itself, the cross-region replication behavior is as follows: <ul style="list-style-type: none"> <li>◦ If you delete the object or bucket before it has been replicated, the object/bucket is not replicated and you are not notified.</li> <li>◦ 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.</li> </ul> </li> </ul>

Operation	Implementation
PUT Bucket tagging	<p>This operation 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:</p> <ul style="list-style-type: none"> <li>• Both StorageGRID and Amazon S3 support up to 50 tags for each bucket.</li> <li>• Tags associated with a bucket must have unique tag keys. A tag key can be up to 128 Unicode characters in length.</li> <li>• Tag values can be up to 256 Unicode characters in length.</li> <li>• Key and values are case sensitive.</li> </ul>
PUT Bucket versioning	<p>This implementation 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:</p> <ul style="list-style-type: none"> <li>• <code>Enabled</code>: Enables versioning for the objects in the bucket. All objects added to the bucket receive a unique version ID.</li> <li>• <code>Suspended</code>: Disables versioning for the objects in the bucket. All objects added to the bucket receive the version ID <code>null</code>.</li> </ul>

### Related information

[Amazon Web Services \(AWS\) Documentation: Cross-Region Replication](#)

[Consistency controls](#)

[GET Bucket last access time request](#)

[Bucket and group access policies](#)

[Using S3 Object Lock](#)

[S3 operations tracked in the audit logs](#)

[Manage objects with ILM](#)

[Use a tenant account](#)

### Creating an 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 the section on object lifecycle management in the *Amazon Simple Storage Service Developer Guide*. Note that StorageGRID only supports Expiration actions; it does not support Transition actions.

[Amazon Simple Storage Service Developer Guide: Object lifecycle management](#)

### What a 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

If you apply a lifecycle configuration to a bucket, the lifecycle settings for the bucket always override StorageGRID ILM settings. StorageGRID uses the Expiration settings for the bucket, not ILM, to determine whether to delete or retain specific objects.

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” in the instructions for managing objects with information lifecycle management.



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:

- DELETE Bucket lifecycle
- GET Bucket lifecycle
- PUT Bucket lifecycle

### Creating the 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"
    }
  ]
}

```

## Applying a lifecycle configuration to a bucket

After you have created the lifecycle configuration file, you apply it to a bucket by issuing a PUT Bucket lifecycle request.

This request applies the lifecycle configuration in the example file to objects in a bucket named `testbucket:bucket`

```
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 GET Bucket lifecycle 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.

## Validating that bucket lifecycle expiration applies to an object

You can determine if an expiration rule in the lifecycle configuration applies to a specific object when issuing a PUT Object, HEAD Object, or GET Object 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” in the instructions for performing StorageGRID administration.

For example, this PUT Object 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 HEAD Object 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": {}
}
```

### Related information

[Operations on buckets](#)

[Manage objects with ILM](#)

### Custom operations on buckets

The StorageGRID system supports custom bucket operations that are added on to the S3 REST API and are specific to the system.

The following table lists the custom bucket operations supported by StorageGRID.

Operation	Description	For more information
GET Bucket consistency	Returns the consistency level being applied to a particular bucket.	<a href="#">GET Bucket consistency request</a>
PUT Bucket consistency	Sets the consistency level applied to a particular bucket.	<a href="#">PUT Bucket consistency request</a>
GET Bucket last access time	Returns whether last access time updates are enabled or disabled for a particular bucket.	<a href="#">GET Bucket last access time request</a>
PUT Bucket last access time	Allows you to enable or disable last access time updates for a particular bucket.	<a href="#">PUT Bucket last access time request</a>

Operation	Description	For more information
DELETE Bucket metadata notification configuration	Deletes the metadata notification configuration XML associated with a particular bucket.	<a href="#">DELETE Bucket metadata notification configuration request</a>
GET Bucket metadata notification configuration	Returns the metadata notification configuration XML associated with a particular bucket.	<a href="#">GET Bucket metadata notification configuration request</a>
PUT Bucket metadata notification configuration	Configures the metadata notification service for a bucket.	<a href="#">PUT Bucket metadata notification configuration request</a>
PUT Bucket modifications for compliance	Deprecated and not supported: You can no longer create new buckets with Compliance enabled.	<a href="#">Deprecated: PUT Bucket request modifications for compliance</a>
GET Bucket compliance	Deprecated but supported: Returns the compliance settings currently in effect for an existing legacy Compliant bucket.	<a href="#">Deprecated: GET Bucket compliance request</a>
PUT Bucket compliance	Deprecated but supported: Allows you to modify the compliance settings for an existing legacy Compliant bucket.	<a href="#">Deprecated: PUT Bucket compliance request</a>

### Related information

[S3 operations tracked in the audit logs](#)

### Operations on objects

This section describes how the StorageGRID system implements S3 REST API operations for objects.

- [Using S3 Object Lock](#)
- [Using server-side encryption](#)
- [GET Object](#)
- [HEAD Object](#)
- [POST Object restore](#)
- [PUT Object](#)
- [PUT Object - Copy](#)

The following conditions apply to all object operations:

- StorageGRID consistency controls are supported by all operations on objects, with the exception of the following:

- GET Object ACL
- OPTIONS /
- PUT Object legal hold
- PUT Object retention
- 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 cannot be accessed through S3.

The following table describes how StorageGRID implements S3 REST API object operations.

Operation	Implementation
DELETE Object	<p>Multi-Factor Authentication (MFA) and the response header <code>x-amz-mfa</code> are not supported.</p> <p>When processing a DELETE Object 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 cannot 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.</p> <p><b>Versioning</b></p> <p>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>.</p> <ul style="list-style-type: none"> <li>• If an object is deleted without the <code>versionId</code> subresource on a version enabled bucket, 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>.</li> <li>• If an object is deleted without the <code>versionId</code> subresource on a version suspended bucket, 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>.</li> </ul> <p><b>Note:</b> In certain cases, multiple delete markers might exist for an object.</p>
DELETE Multiple Objects	<p>Multi-Factor Authentication (MFA) and the response header <code>x-amz-mfa</code> are not supported.</p> <p>Multiple objects can be deleted in the same request message.</p>

Operation	Implementation
DELETE Object tagging	<p>Uses the <code>tagging</code> subresource to remove all tags from an object. Implemented with all Amazon S3 REST API behavior.</p> <p><b>Versioning</b></p> <p>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>.</p>
GET Object	<a href="#">GET Object</a>
GET Object ACL	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.
GET Object legal hold	<a href="#">Using S3 Object Lock</a>
GET Object retention	<a href="#">Using S3 Object Lock</a>
GET Object tagging	<p>Uses the <code>tagging</code> subresource to return all tags for an object. Implemented with all Amazon S3 REST API behavior</p> <p><b>Versioning</b></p> <p>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>.</p>
HEAD Object	<a href="#">HEAD Object</a>
POST Object restore	<a href="#">POST Object restore</a>
PUT Object	<a href="#">PUT Object</a>
PUT Object - Copy	<a href="#">PUT Object - Copy</a>

Operation	Implementation
PUT Object legal hold	<a href="#">Using S3 Object Lock</a>
PUT Object retention	<a href="#">Using S3 Object Lock</a>
PUT Object tagging	<p>Uses the <code>tagging</code> subresource to add a set of tags to an existing object. Implemented with all Amazon S3 REST API behavior</p> <p><b>Tag updates and ingest behavior</b></p> <p>When you use PUT Object tagging 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.</p> <p>This means that if the ILM rule uses the Strict option for ingest behavior, no action is taken if the required object placements cannot be made (for example, because a newly required location is unavailable). The updated object retains its current placement until the required placement is possible.</p> <p><b>Resolving conflicts</b></p> <p>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.</p> <p><b>Versioning</b></p> <p>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>.</p>

**Related information**

[Consistency controls](#)

[S3 operations tracked in the audit logs](#)

## Using 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 and then specify retain-until-date and legal hold settings for each object version you add to that bucket.

S3 Object Lock allows you to specify object-level settings to prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely.

The StorageGRID S3 Object Lock feature provides a single retention mode that is equivalent to the Amazon S3 compliance mode. By default, a protected object version cannot be overwritten or deleted by any user. The StorageGRID S3 Object Lock feature does not support a governance mode, and it does not allow users with special permissions to bypass retention settings or to delete protected objects.

### Enabling 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. You can use either of these methods:

- Create the bucket using the Tenant Manager.

[Use a tenant account](#)

- Create the bucket using a PUT Bucket request with the `x-amz-bucket-object-lock_enabled` request header.

[Operations on buckets](#)

You cannot add or disable S3 Object Lock after the bucket is created. S3 Object Lock requires bucket versioning, which is enabled automatically when you create the bucket.

A bucket with S3 Object Lock enabled can contain a combination of objects with and without S3 Object Lock settings. StorageGRID does not support default retention for the objects in S3 Object Lock buckets, so the PUT Object Lock Configuration bucket operation is not supported.

### Determining if S3 Object Lock is enabled for a bucket

To determine if S3 Object Lock is enabled, use the GET Object Lock Configuration request.

[Operations on buckets](#)

### Creating an object with S3 Object Lock settings

To specify S3 Object Lock settings when adding an object version to a bucket that has S3 Object Lock enabled, issue a PUT Object, PUT Object - Copy, or Initiate Multipart Upload request. Use the following request headers.



You must enable S3 Object Lock when you create a bucket. You cannot add or disable S3 Object Lock after a bucket is created.

- `x-amz-object-lock-mode`, which must be COMPLIANCE (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 are not 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 PUT Object request. `Content-MD5` is not required for PUT Object - Copy or Initiate Multipart Upload.
- 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 PUT Object 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 HEAD Object 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.
- A subsequent DELETE Object version or DELETE Objects versions request will fail if it is before the `retain-until-date` or if a legal hold is on.

## Updating S3 Object Lock settings

If you need to update the legal hold or retention settings for an existing object version, you can perform the following object subresource operations:

- `PUT Object legal-hold`

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.

- `PUT Object retention`
  - The `mode` value must be `COMPLIANCE` (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 are not allowed.
  - If an object version has an existing `retain-until-date`, you can only increase it. The new value must be in the future.

## Related information

[Manage objects with ILM](#)

[Use a tenant account](#)

[PUT Object](#)

[PUT Object - Copy](#)

[Initiate Multipart Upload](#)

[Object versioning](#)

[Amazon Simple Storage Service User Guide: Using S3 Object Lock](#)

### Using 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.

### Using 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:

- PUT Object
- PUT Object - Copy
- Initiate Multipart Upload

## Using 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:

- GET Object
- HEAD Object
- PUT Object
- PUT Object - Copy
- Initiate Multipart Upload
- Upload Part
- Upload Part - Copy

### 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 CloudMirror replication is configured for the bucket, you cannot ingest SSE-C objects. The ingest operation will fail.

### Related information

[GET Object](#)

[HEAD Object](#)

[PUT Object](#)

[PUT Object - Copy](#)

[Initiate Multipart Upload](#)

[Upload Part](#)

[Upload Part - Copy](#)

[Amazon S3 Developer Guide: Protecting Data Using Server-Side Encryption with Customer-Provided Encryption Keys \(SSE-C\)](#)

### GET Object

You can use the S3 GET Object request to retrieve an object from an S3 bucket.

### partNumber request parameter is not supported

The `partNumber` request parameter is not supported for GET Object requests. You cannot perform a GET request to retrieve a specific part of a multipart object. A 501 Not Implemented error is returned with the following message:

```
GET Object by partNumber is not implemented
```

### 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 “Using server-side encryption.”

## 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 do not return the `x-amz-missing-meta` header if the key name or value includes unprintable characters.

## 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`.

## Behavior of GET Object for Cloud Storage Pool objects

If an object has been stored in a Cloud Storage Pool (see the instructions for managing objects with information lifecycle management), the behavior of a GET Object request depends on the state of the object. See “HEAD Object” 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, GET Object requests will attempt to retrieve data from the grid, before retrieving it from the Cloud Storage Pool.

State of object	Behavior of GET Object
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 POST Object restore 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 POST Object restore 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 GET Object 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 GET Object request might return some data but stop midway through the transfer.
- A subsequent GET Object request might return `403 Forbidden`.

### Related information

[Using server-side encryption](#)

[Manage objects with ILM](#)

[POST Object restore](#)

[S3 operations tracked in the audit logs](#)

### HEAD Object

You can use the S3 HEAD Object 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 HEAD Object to determine the object's transition state.

### 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 "Using server-side encryption."

### 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 do not return the `x-amz-missing-meta` header if the key name or value includes unprintable characters.

### Unsupported request header

The following request header is not supported and returns `XNotImplemented`:

- `x-amz-website-redirect-location`

## Response headers for Cloud Storage Pool objects

If the object is stored in a Cloud Storage Pool (see the instructions for managing objects with information lifecycle management), 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 HEAD object
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.
Object has transitioned to non-retrievable state, but at least one copy also exists on the grid	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>  The value for <code>expiry-date</code> is set to some distant time in the future.  <b>Note:</b> If the copy on the grid is not available (for example, a Storage Node is down), you must issue a POST Object restore request to restore the copy from the Cloud Storage Pool before you can successfully retrieve the object.

State of object	Response to HEAD 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 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 HEAD Object 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.

### 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`.

### Related information

[Using server-side encryption](#)

[Manage objects with ILM](#)

[POST Object restore](#)

[S3 operations tracked in the audit logs](#)

### POST Object restore

You can use the S3 POST Object restore request to restore an object that is stored in a Cloud Storage Pool.

## Supported request type

StorageGRID only supports POST Object restore 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 do not specify `versionId`, the most recent version of the object is restored

## Behavior of POST Object restore on Cloud Storage Pool objects

If an object has been stored in a Cloud Storage Pool (see the instructions for managing objects with information lifecycle management), a POST Object restore request has the following behavior, based on the state of the object. See “HEAD Object” 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 POST Object restore request. Instead, the local copy can be retrieved directly, using a GET Object request.

State of object	Behavior of POST Object restore
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.  <b>Note:</b> Before an object has been transitioned to a non-retrievable state, you cannot change its <code>expiry-date</code> .
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 (Expedited, Standard, or Bulk). If you do not specify <code>Tier</code> , the Standard tier is used.  <b>Attention:</b> If an object has been transitioned to S3 Glacier Deep Archive or the Cloud Storage Pool uses Azure Blob Storage, you cannot restore it using the Expedited tier. The following error is returned 403 Forbidden, InvalidTier: Retrieval option is not supported by this storage class.

State of object	Behavior of POST Object restore
Object in process of being restored from a non-retrievable state	409 Conflict, RestoreAlreadyInProgress
Object fully restored to the Cloud Storage Pool	200 OK  <b>Note:</b> If an object has been restored to a retrievable state, you can change its <code>expiry-date</code> by reissuing the POST Object restore request with a new value for <code>Days</code> . The restoration date is updated relative to the time of the request.

### Related information

[Manage objects with ILM](#)

[HEAD Object](#)

[S3 operations tracked in the audit logs](#)

### PUT Object

You can use the S3 PUT Object request to add an object to a bucket.

### 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.

### Object size

StorageGRID supports objects up to 5 TB in size.

### 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:

- PUT, PUT Object-Copy, GET, and HEAD 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-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 cannot use both a **User Defined Creation Time** for the Reference Time and the Balanced or Strict options for Ingest Behavior. 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`

### Using 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`

### S3 REST API supported operations and limitations

#### Unsupported request headers

The following request headers are not supported:

- The `x-amz-acl` request header is not supported.
- The `x-amz-website-redirect-location` request header is not supported and 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 option for Ingest Behavior, 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 do not, 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 cannot 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 cannot 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.

**Attention:** 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 policy, and does not result in data being stored at lower levels of redundancy in the StorageGRID system.

**Note:** 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`
- **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.

**Attention:** 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 “Using server-side encryption.”

**Note:** 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.

### Related information

[Manage objects with ILM](#)

[Operations on buckets](#)

[S3 operations tracked in the audit logs](#)

[Using server-side encryption](#)

[How client connections can be configured](#)

### PUT Object - Copy

You can use the S3 PUT Object - Copy request to create a copy of an object that is already stored in S3. A PUT Object - Copy operation is the same as performing a GET and then a PUT.

### 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.

### Object size

StorageGRID supports objects up to 5 TB in size.

### 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`

## Using 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`

### Unsupported request headers

The following request headers are not supported:

- Cache-Control
- Content-Disposition
- Content-Encoding
- Content-Language
- Expires
- 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 specifies an Ingest Behavior of Dual commit or Balanced.

- 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 PUT Object - Copy

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 cannot use PUT Object - Copy 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 cannot 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 PUT Object - Copy 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.

**Attention:** 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 “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 PUT Object - Copy request:
  - `x-amz-server-side-encryption`

**Note:** The `server-side-encryption` value of the object cannot 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.

## Related information

[Manage objects with ILM](#)

[Using server-side encryption](#)

[S3 operations tracked in the audit logs](#)

## Operations for multipart uploads

This section describes how StorageGRID supports operations for multipart uploads.

- [List multipart uploads](#)
- [Initiate Multipart Upload](#)
- [Upload Part](#)
- [Upload Part - Copy](#)
- [Complete Multipart Upload](#)

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 List Multipart Uploads 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. Since 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 Strict or Balanced ingest behavior. You should be aware of how this affects object and part placement:
  - If ILM changes while an S3 multipart upload is in progress, when the multipart upload completes some parts of the object might not meet current ILM requirements. Any part that is not placed correctly is queued for ILM re-evaluation, and is 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 do not 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 as a whole, all parts of the object are moved to DC1.
- All of the multipart upload operations support StorageGRID consistency controls.
- 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 Initiate Multipart Upload request only. To use SSE-C (server-side encryption with customer-provided keys), you specify the same three encryption key request headers in the Initiate Multipart Upload request and in each subsequent Upload Part request.

Operation	Implementation
List Multipart Uploads	See <a href="#">List Multipart Uploads</a>

Operation	Implementation
Initiate Multipart Upload	See <a href="#">Initiate Multipart Upload</a>
Upload Part	See <a href="#">Upload Part</a>
Upload Part - Copy	See <a href="#">Upload Part - Copy</a>
Complete Multipart Upload	See <a href="#">Complete Multipart Upload</a>
Abort Multipart Upload	Implemented with all Amazon S3 REST API behavior
List Parts	Implemented with all Amazon S3 REST API behavior

### Related information

[Consistency controls](#)

[Using server-side encryption](#)

### List Multipart Uploads

The List Multipart Uploads operation lists in-progress multipart uploads for a bucket.

The following request parameters are supported:

- `encoding-type`
- `max-uploads`
- `key-marker`
- `prefix`
- `upload-id-marker`

The `delimiter` request parameter is not supported.

### Versioning

Multipart upload consists of separate operations for initiating the upload, listing uploads, uploading parts, assembling the uploaded parts, and completing the upload. When the Complete Multipart Upload operation is performed, that is the point when objects are created (and versioned if applicable).

### Initiate Multipart Upload

The 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 option for Ingest Behavior, 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 do not, 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 cannot 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 cannot 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.

**Attention:** 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 policy, and does not result in data being stored at lower levels of redundancy in the StorageGRID system.

**Note:** 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.

The following request headers are supported:

- `Content-Type`
- `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`

### Using 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`

### S3 REST API supported operations and limitations



For information on how StorageGRID handles UTF-8 characters, see the documentation for PUT Object.

### 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 Initiate Multipart Upload request if you want to encrypt the object with a unique key managed by StorageGRID. Do not specify this header in any of the Upload Part requests.
  - `x-amz-server-side-encryption`
- **SSE-C**: Use all three of these headers in the Initiate Multipart Upload request (and in each subsequent Upload Part 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.

**Attention:** 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 “Using server-side encryption.”

### Unsupported request headers

The following request header is not supported and returns `XNotImplemented`

- `x-amz-website-redirect-location`

### 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 Complete Multipart Upload operation is performed.

### Related information

[Manage objects with ILM](#)

[Using server-side encryption](#)

[PUT Object](#)

### Upload Part

The Upload Part operation uploads a part in a multipart upload for an object.

### Supported request headers

The following request headers are supported:

- `Content-Length`
- `Content-MD5`

### Request headers for server-side encryption

If you specified SSE-C encryption for the Initiate Multipart Upload request, you must also include the following request headers in each Upload Part 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 Initiate Multipart Upload request.
- `x-amz-server-side-encryption-customer-key-MD5`: Specify the same MD5 digest that you provided in the Initiate Multipart Upload 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 “Using 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 Complete Multipart Upload operation is performed.

### Related information

[Using server-side encryption](#)

### Upload Part - Copy

The Upload Part - Copy operation uploads a part of an object by copying data from an existing object as the data source.

The Upload Part - Copy operation is implemented with all Amazon S3 REST API behavior.

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 Initiate Multipart Upload request, you must also include the following request headers in each Upload Part - Copy 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 Initiate Multipart Upload request.
- `x-amz-server-side-encryption-customer-key-MD5`: Specify the same MD5 digest that you provided in the Initiate Multipart Upload request.

If the source object is encrypted using a customer-provided key (SSE-C), you must include the following three headers in the Upload Part - Copy 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 “Using 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 Complete Multipart Upload operation is performed.

### Complete Multipart Upload

The Complete Multipart Upload operation completes a multipart upload of an object by assembling the previously uploaded parts.

### 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.

### Object size

StorageGRID supports objects up to 5 TB in size.

### Request headers

The `x-amz-storage-class` request header is supported, and affects how many object copies StorageGRID creates if the matching ILM rule specifies an Ingest Behavior of Dual commit or Balanced.

- 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.

## Versioning

This operation completes a multipart upload. If versioning is enabled for a bucket, the object version is created upon 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.

If this occurs, an alarm is raised in the Grid Manager on Total Events (SMTT). The Last Event message displays “Failed to publish notifications for bucket-nameobject key” for the last object whose notification failed. (To see this message, select **Nodes** > **Storage Node** > **Events**. View Last Event at the top of the table.) Event messages are also listed in `/var/local/log/bycast-err.log`.

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.

### Related information

[Manage objects with ILM](#)

### 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

<b>Name</b>	<b>HTTP status</b>
IncompleteBody	400 Bad Request
InternalServerError	500 Internal Server Error
InvalidAccessKeyId	403 Forbidden
InvalidArgument	400 Bad Request
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

<b>Name</b>	<b>HTTP status</b>
NoSuchBucket	404 Not Found
NoSuchKey	404 Not Found
NoSuchUpload	404 Not Found
NotImplemented	501 Not Implemented
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**

<b>Name</b>	<b>Description</b>	<b>HTTP status</b>
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

Name	Description	HTTP status
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 S3 REST API operations

There are operations added on to the S3 REST API that are specific to StorageGRID system.

### GET Bucket consistency request

The GET Bucket consistency request allows you to determine the consistency level being applied to a particular bucket.

The default consistency controls are 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: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>
```

#### Response

In the response XML, `<Consistency>` will return one of the following values:

Consistency control	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.

Consistency control	Description
strong-site	Guarantees read-after-write consistency for all client requests within a site.
read-after-new-write	<p>(Default) Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Matches Amazon S3 consistency guarantees.</p> <p><b>Note:</b> If your application uses HEAD requests on objects that do not exist, you might receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable. To prevent these errors, set the consistency control to “available” unless you require consistency guarantees similar to Amazon S3.</p>
available (eventual consistency for HEAD operations)	Behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than “read-after-new-write” if Storage Nodes are unavailable. Differs from Amazon S3 consistency guarantees for HEAD operations only.

### 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 controls](#)

### PUT Bucket consistency request

The PUT Bucket consistency request allows you to specify the consistency level to apply to operations performed on a bucket.

The default consistency controls are set to guarantee read-after-write for newly created objects.

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 control	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	<p>(Default) Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Matches Amazon S3 consistency guarantees.</p> <p><b>Note:</b> If your application uses HEAD requests on objects that do not exist, you might receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable. To prevent these errors, set the consistency control to “available” unless you require consistency guarantees similar to Amazon S3.</p>
available (eventual consistency for HEAD operations)	Behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than “read-after-new-write” if Storage Nodes are unavailable. Differs from Amazon S3 consistency guarantees for HEAD operations only.

**Note:** In general, you should use the “read-after-new-write” consistency control value. If requests are not working correctly, change the application client behavior if possible. Or, configure the client to specify the consistency control for each API request. Set the consistency control at the bucket level only as a last resort.

### Request example

```
PUT /bucket?x-ntap-sg-consistency=strong-global HTTP/1.1
Date: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>
```

## Related information

### [Consistency controls](#)

## GET Bucket last access time request

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: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>
```

### 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 request

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, the **S3 > Buckets > Change Last Access Setting** check box in the Tenant Manager, or the Tenant Management API.

If last access time updates are disabled for a bucket, the following behavior is applied to operations on the bucket:

- GET Object, GET Object ACL, GET Object Tagging, and HEAD Object requests do not update last access time. The object is not added to queues for information lifecycle management (ILM) evaluation.
- PUT Object - Copy and PUT Object Tagging 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, PUT Object - Copy requests do not 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, PUT Object - Copy requests always update last access time. The copy of the object is added to queues for ILM evaluation.
- Complete Multipart Upload 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: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>
```

This example disables last access time for a bucket.

```
PUT /bucket?x-ntap-sg-lastaccesstime=disabled HTTP/1.1
Date: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>
```

#### Related information

[Use a tenant account](#)

#### DELETE Bucket metadata notification configuration request

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: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>
```

### GET Bucket metadata notification configuration request

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: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>
```

### 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	<p>Container tag for rules used to specify the objects and destination for metadata notifications.</p> <p>Contains one or more Rule elements.</p>	Yes
Rule	<p>Container tag for a rule that identifies the objects whose metadata should be added to a specified index.</p> <p>Rules with overlapping prefixes are rejected.</p> <p>Included in the MetadataNotificationConfiguration element.</p>	Yes
ID	<p>Unique identifier for the rule.</p> <p>Included in the Rule element.</p>	No

Name	Description	Required
Status	<p>Status can be 'Enabled' or 'Disabled'. No action is taken for rules that are disabled.</p> <p>Included in the Rule element.</p>	Yes
Prefix	<p>Objects that match the prefix are affected by the rule, and their metadata is sent to the specified destination.</p> <p>To match all objects, specify an empty prefix.</p> <p>Included in the Rule element.</p>	Yes
Destination	<p>Container tag for the destination of a rule.</p> <p>Included in the Rule element.</p>	Yes

Name	Description	Required
Urn	<p>URN of the destination where object metadata is sent. Must be the URN of a StorageGRID endpoint with the following properties:</p> <ul style="list-style-type: none"> <li>• <code>es</code> must be the third element.</li> <li>• The URN must end with the index and type where the metadata is stored, in the form <code>domain-name/myindex/mytype</code>.</li> </ul> <p>Endpoints are configured using the Tenant Manager or Tenant Management API. They take the following form:</p> <ul style="list-style-type: none"> <li>• <code>arn:aws:es:_region:account-ID_:domain/mydomain/myindex/mytype</code></li> <li>• <code>urn:mysite:es:::mydomain/myindex/mytype</code></li> </ul> <p>The endpoint must be configured before the configuration XML is submitted, or configuration will fail with a 404 error.</p> <p>Urn is included in the Destination element.</p>	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-
1:33333333:domain/records/current/2017</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>
```

## Related information

[Use a tenant account](#)

## PUT Bucket metadata notification configuration request

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 are not valid, and are rejected when they are submitted. For example, a configuration that included one rule for 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
MetadataNotificationConfiguration	<p>Container tag for rules used to specify the objects and destination for metadata notifications.</p> <p>Contains one or more Rule elements.</p>	Yes
Rule	<p>Container tag for a rule that identifies the objects whose metadata should be added to a specified index.</p> <p>Rules with overlapping prefixes are rejected.</p> <p>Included in the MetadataNotificationConfiguration element.</p>	Yes
ID	<p>Unique identifier for the rule.</p> <p>Included in the Rule element.</p>	No

Name	Description	Required
Status	<p>Status can be 'Enabled' or 'Disabled'. No action is taken for rules that are disabled.</p> <p>Included in the Rule element.</p>	Yes
Prefix	<p>Objects that match the prefix are affected by the rule, and their metadata is sent to the specified destination.</p> <p>To match all objects, specify an empty prefix.</p> <p>Included in the Rule element.</p>	Yes
Destination	<p>Container tag for the destination of a rule.</p> <p>Included in the Rule element.</p>	Yes

Name	Description	Required
Urn	<p>URN of the destination where object metadata is sent. Must be the URN of a StorageGRID endpoint with the following properties:</p> <ul style="list-style-type: none"> <li>• <code>es</code> must be the third element.</li> <li>• The URN must end with the index and type where the metadata is stored, in the form <code>domain-name/myindex/mytype</code>.</li> </ul> <p>Endpoints are configured using the Tenant Manager or Tenant Management API. They take the following form:</p> <ul style="list-style-type: none"> <li>• <code>arn:aws:es:region:account-ID:domain/mydomain/myindex/mytype</code></li> <li>• <code>urn:mysite:es:::mydomain/myindex/mytype</code></li> </ul> <p>The endpoint must be configured before the configuration XML is submitted, or configuration will fail with a 404 error.</p> <p>Urn is included in the Destination element.</p>	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: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>

<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: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Images-rule</ID>
    <Status>Enabled</Status>
    <Prefix>/images</Prefix>
    <Destination>
      <Urn>arn:aws:es:us-east-1:3333333: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>
```

## Related information

[Use a tenant account](#)

### 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>

Type	Item name	Description
System metadata	size	Object size (in bytes) as visible to an HTTP client
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

**Note:** 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 cannot edit the document's field types in the index.

### 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 GET Service 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 cannot 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: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>
```

#### 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 are not added to the `ObjectCount` total.

## Related information

[Consistency controls](#)

## 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 automatically when you upgrade to StorageGRID 11.5. 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.

[Using S3 Object Lock](#)

[Manage objects with ILM](#)

[NetApp Knowledge Base: How to manage legacy Compliant buckets in StorageGRID 11.5](#)

#### **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.



The StorageGRID Compliance feature that was available in previous StorageGRID versions is deprecated and has been replaced by S3 Object Lock.

[Using S3 Object Lock](#)

[Manage objects with ILM](#)

[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 PUT Bucket 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.
```

#### **Related information**

[Manage objects with ILM](#)

[Use a tenant account](#)

#### **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.

[Using S3 Object Lock](#)

[Manage objects with ILM](#)

[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: <em>date</em>
Authorization: <em>authorization string</em>
Host: <em>host</em>
```

## 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: <em>date</em>
Connection: <em>connection</em>
Server: StorageGRID/11.1.0
x-amz-request-id: <em>request ID</em>
Content-Length: <em>length</em>
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.
LegalHold	<ul style="list-style-type: none"><li>• True: This bucket is currently under a legal hold. Objects in this bucket cannot be deleted until the legal hold is lifted, even if their retention period has expired.</li><li>• False: This bucket is not currently under a legal hold. Objects in this bucket can be deleted when their retention period expires.</li></ul>

Name	Description
AutoDelete	<ul style="list-style-type: none"> <li>• True: The objects in this bucket will be deleted automatically when their retention period expires, unless the bucket is under a legal hold.</li> <li>• 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.</li> </ul>

## 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.

## Related information

[Manage objects with ILM](#)

[Use a tenant account](#)

### 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.

[Using S3 Object Lock](#)

[Manage objects with ILM](#)

[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: <em>date</em>
```

```
Authorization: <em>authorization name</em>
```

```
Host: <em>host</em>
```

```
Content-Length: 152
```

```
<SGCompliance>
```

```
<RetentionPeriodMinutes>1051200</RetentionPeriodMinutes>
```

```
<LegalHold>>false</LegalHold>
```

```
<AutoDelete>>true</AutoDelete>
```

```
</SGCompliance>
```

Name	Description
RetentionPeriodMinutes	<p>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.</p> <p><b>Attention:</b> 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 cannot decrease that value; you can only increase it.</p>
LegalHold	<ul style="list-style-type: none"><li>• True: This bucket is currently under a legal hold. Objects in this bucket cannot be deleted until the legal hold is lifted, even if their retention period has expired.</li><li>• False: This bucket is not currently under a legal hold. Objects in this bucket can be deleted when their retention period expires.</li></ul>
AutoDelete	<ul style="list-style-type: none"><li>• True: The objects in this bucket will be deleted automatically when their retention period expires, unless the bucket is under a legal hold.</li><li>• 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.</li></ul>

### Consistency level 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 level 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 cannot achieve the **strong-global** consistency level 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 level.



Never force the **strong-site** consistency level 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 level 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 level, 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 level, 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](#)

[Use a tenant account](#)

[Manage objects with ILM](#)

## 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 configured using the GET Bucket policy, PUT Bucket policy, and DELETE Bucket policy S3 API operations. 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.

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	<p>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.</p> <p>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.</p>

Element	Description
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.
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 are not supported except to set anonymous access, which grants permission to everyone. For example, you set the wildcard (\*) as the Principal value.

```
"Principal": "*"

```

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:iam:s3::mybucket",
        "arn:aws:iam: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.

### Related information

[Use a tenant account](#)

### Consistency control settings for policies

By default, any updates you make to group policies are eventually consistent. Once 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 also eventually 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 become effective as soon as possible for security reasons.

In this case, you can either set the `Consistency-Control` header in the PUT Bucket policy request, or you can use the PUT Bucket consistency request. When changing the consistency control for this request, you must use the value **all**, which provides the highest guarantee of read-after-write consistency. If you specify any other consistency control value in a header for the PUT Bucket consistency request, the request will be rejected. If you specify any other value for a PUT Bucket policy request, the value will be ignored. Once 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 level to **all** to force a new bucket policy to become effective sooner, be sure to set the bucket-level control back to its original value when you are done. Otherwise, all future bucket requests will use the **all** setting.

## Using the 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 PUT Bucket policy operation must be encoded with charset=UTF-8.

## Specifying 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.

## Related information

### Specifying 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 do not 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.

## Specifying 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 PUT and DELETE Bucket replication 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
<code>s3:CreateBucket</code>	PUT Bucket	
<code>s3&gt;DeleteBucket</code>	DELETE Bucket	
<code>s3&gt;DeleteBucketMetadataNotification</code>	DELETE Bucket metadata notification configuration	Yes
<code>s3&gt;DeleteBucketPolicy</code>	DELETE Bucket policy	
<code>s3&gt;DeleteReplicationConfiguration</code>	DELETE Bucket replication	Yes, separate permissions for PUT and DELETE*
<code>s3:GetBucketAcl</code>	GET Bucket ACL	
<code>s3:GetBucketCompliance</code>	GET Bucket compliance (deprecated)	Yes
<code>s3:GetBucketConsistency</code>	GET Bucket consistency	Yes

Permissions	S3 REST API operations	Custom for StorageGRID
s3:GetBucketCORS	GET Bucket cors	
s3:GetEncryptionConfiguration	GET Bucket encryption	
s3:GetBucketLastAccessTime	GET Bucket last access time	Yes
s3:GetBucketLocation	GET Bucket location	
s3:GetBucketMetadataNotification	GET Bucket metadata notification configuration	Yes
s3:GetBucketNotification	GET Bucket notification	
s3:GetBucketObjectLockConfiguration	GET Object Lock Configuration	
s3:GetBucketPolicy	GET Bucket policy	
s3:GetBucketTagging	GET Bucket tagging	
s3:GetBucketVersioning	GET Bucket versioning	
s3:GetLifecycleConfiguration	GET Bucket lifecycle	
s3:GetReplicationConfiguration	GET Bucket replication	
s3:ListAllMyBuckets	<ul style="list-style-type: none"> <li>• GET Service</li> <li>• GET Storage Usage</li> </ul>	Yes, for GET Storage Usage
s3:ListBucket	<ul style="list-style-type: none"> <li>• GET Bucket (List Objects)</li> <li>• HEAD Bucket</li> <li>• POST Object restore</li> </ul>	
s3:ListBucketMultipartUploads	<ul style="list-style-type: none"> <li>• List Multipart Uploads</li> <li>• POST Object restore</li> </ul>	
s3:ListBucketVersions	GET Bucket versions	
s3:PutBucketCompliance	PUT Bucket compliance (deprecated)	Yes
s3:PutBucketConsistency	PUT Bucket consistency	Yes

Permissions	S3 REST API operations	Custom for StorageGRID
s3:PutBucketCORS	<ul style="list-style-type: none"> <li>• DELETE Bucket cors†</li> <li>• PUT Bucket cors</li> </ul>	
s3:PutEncryptionConfiguration	<ul style="list-style-type: none"> <li>• DELETE Bucket encryption</li> <li>• PUT Bucket encryption</li> </ul>	
s3:PutBucketLastAccessTime	PUT Bucket last access time	Yes
s3:PutBucketMetadataNotification	PUT Bucket metadata notification configuration	Yes
s3:PutBucketNotification	PUT Bucket notification	
s3:PutBucketObjectLockConfiguration	PUT Bucket with the <code>x-amz-bucket-object-lock-enabled: true</code> request header (also requires the s3:CreateBucket permission)	
s3:PutBucketPolicy	PUT Bucket policy	
s3:PutBucketTagging	<ul style="list-style-type: none"> <li>• DELETE Bucket tagging†</li> <li>• PUT Bucket tagging</li> </ul>	
s3:PutBucketVersioning	PUT Bucket versioning	
s3:PutLifecycleConfiguration	<ul style="list-style-type: none"> <li>• DELETE Bucket lifecycle†</li> <li>• PUT Bucket lifecycle</li> </ul>	
s3:PutReplicationConfiguration	PUT Bucket replication	Yes, separate permissions for PUT and DELETE*

#### Permissions that apply to objects

Permissions	S3 REST API operations	Custom for StorageGRID
s3:AbortMultipartUpload	<ul style="list-style-type: none"> <li>• Abort Multipart Upload</li> <li>• POST Object restore</li> </ul>	
s3:DeleteObject	<ul style="list-style-type: none"> <li>• DELETE Object</li> <li>• DELETE Multiple Objects</li> <li>• POST Object restore</li> </ul>	

Permissions	S3 REST API operations	Custom for StorageGRID
s3:DeleteObjectTagging	DELETE Object Tagging	
s3:DeleteObjectVersionTagging	DELETE Object Tagging (a specific version of the object)	
s3:DeleteObjectVersion	DELETE Object (a specific version of the object)	
s3:GetObject	<ul style="list-style-type: none"> <li>• GET Object</li> <li>• HEAD Object</li> <li>• POST Object restore</li> </ul>	
s3:GetObjectAcl	GET Object ACL	
s3:GetObjectLegalHold	GET Object legal hold	
s3:GetObjectRetention	GET Object retention	
s3:GetObjectTagging	GET Object Tagging	
s3:GetObjectVersionTagging	GET Object Tagging (a specific version of the object)	
s3:GetObjectVersion	GET Object (a specific version of the object)	
s3:ListMultipartUploadParts	List Parts, POST Object restore	
s3:PutObject	<ul style="list-style-type: none"> <li>• PUT Object</li> <li>• PUT Object - Copy</li> <li>• POST Object restore</li> <li>• Initiate Multipart Upload</li> <li>• Complete Multipart Upload</li> <li>• Upload Part</li> <li>• Upload Part - Copy</li> </ul>	
s3:PutObjectLegalHold	PUT Object legal hold	
s3:PutObjectRetention	PUT Object retention	
s3:PutObjectTagging	PUT Object Tagging	

Permissions	S3 REST API operations	Custom for StorageGRID
s3:PutObjectVersionTagging	PUT Object Tagging (a specific version of the object)	
s3:PutOverwriteObject	<ul style="list-style-type: none"> <li>• PUT Object</li> <li>• PUT Object - Copy</li> <li>• PUT Object tagging</li> <li>• DELETE Object tagging</li> <li>• Complete Multipart Upload</li> </ul>	Yes
s3:RestoreObject	POST Object restore	

### Using the 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 cannot 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 cannot be overwritten.
    - Any ingest operations in progress are cancelled, and an error is returned.
    - If S3 versioning is enabled, the Deny setting prevents PUT Object tagging or DELETE Object tagging 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 cannot overwrite an object's data, user-defined metadata, or object tagging. In addition, if the **Prevent Client Modification** check box is selected (**Configuration > Grid Options**), that setting overrides the setting of the PutOverwriteObject permission.

### Related information

[S3 group policy examples](#)

### Specifying 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: {
  <em>condition_type</em>: {
    <em>condition_key</em>: <em>condition_values</em>
```

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.

Condition operators	Description
NumericEquals	Compares a key to a numeric value based on exact matching.
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

Category	Applicable condition keys	Description
IP operators	aws:SourceIp	<p>Will compare to the IP address from which the request was sent. Can be used for bucket or object operations.</p> <p><b>Note:</b> 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.</p> <p><b>Note:</b> If a third-party, non-transparent load balancer is used, this will compare to the IP address of that load balancer. Any X-Forwarded-For header will be ignored since its validity cannot be ascertained.</p>
Resource/Identity	aws:username	Will compare to the sender's username from which the request was sent. Can be used for bucket or object operations.
S3:ListBucket and S3:ListBucketVersions permissions	s3:delimiter	Will compare to the delimiter parameter specified in a GET Bucket or GET Bucket Object versions request.
S3:ListBucket and S3:ListBucketVersions permissions	s3:max-keys	Will compare to the max-keys parameter specified in a GET Bucket or GET Bucket Object versions request.
S3:ListBucket and S3:ListBucketVersions permissions	s3:prefix	Will compare to the prefix parameter specified in a GET Bucket or GET Bucket Object versions request.

### Specifying 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 SourceIp key as the provided variable.
<code>\${aws:username}</code>	Uses the username key as the provided variable.
<code>\${s3:prefix}</code>	Uses the service-specific prefix key as the provided variable.
<code>\${s3:max-keys}</code>	Uses the service-specific max-keys key as the provided variable.
<code>\${*}</code>	Special character. Uses the character as a literal * character.
<code>\${?}</code>	Special character. Uses the character as a literal ? character.
<code>\${\$}</code>	Special character. Uses the character as a literal \$ character.

### Creating 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

Policy description	Policy type	Amazon behavior	StorageGRID behavior
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 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 > Grid Options**, and select the **Prevent Client Modification** check box.
- 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 PUT Object 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 do not 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

[Manage objects with ILM](#)

[Creating policies requiring special handling](#)

[How StorageGRID ILM rules manage objects](#)

[S3 group policy examples](#)

#### S3 policy examples

Use the examples in this section to build StorageGRID access policies for buckets and groups.

##### S3 bucket policy examples

Bucket policies specify the access permissions for the bucket that the policy is attached to. Bucket policies are configured using the S3 PutBucketPolicy API.

A bucket policy can be configured using the AWS CLI as per the following command:

```
> aws s3api put-bucket-policy --bucket examplebucket --policy
<em>file://policy.json</em>
```

### 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 Get Object operations on all objects in the bucket. All other operations will be denied. Note that this policy might not be particularly useful since 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 GET Object 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/*"
    }
  ]
}

```

## Related information

[Operations on buckets](#)

## S3 group policy examples

Group policies specify the access permissions for the group that the policy is attached to. There is no `Principal` element in the policy since it is implicit. Group policies are configured using the Tenant Manager or the API.

## Example: Setting the group policy using the Tenant Manager

When using the Tenant Manager to add or edit a group, you can select how you want to create the group policy that defines which S3 access permissions members of this group will have, as follows:

- **No S3 Access:** Default option. Users in this group do not 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 cannot 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 cannot edit this string.
- **Custom:** Users in the group are granted the permissions you specify in the text box.

In this example, members of the group are only permitted to list and access their specific folder (key prefix) in the specified bucket.



No S3 Access

Read Only Access

Full Access

Custom  
(Must be a valid JSON formatted string.)

```
{
  "Statement": [
    {
      "Sid": "AllowListBucketOfASpecificUserPrefix",
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::department-bucket",
      "Condition": {
        "StringLike": {
          "s3:prefix": "${aws:username}/"
        }
      }
    },
    {
      "Sid": "AllowUserSpecificActionsOnlyInTheSpecificFolder",
      "Effect": "Allow",
      "Action": "s3:*Object",
      "Resource": "arn:aws:s3:::department-bucket/${aws:username}/"
    }
  ]
}
```

## 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}/*"
    }
  ]
}

```

### Related information

[Use a tenant account](#)

[Using the PutOverwriteObject permission](#)

[Write-once-read-many \(WORM\) protection](#)

## Configuring security for the REST API

You should review the security measures implemented for the REST API and understand how to secure your system.

### How StorageGRID provides security for the REST API

You should understand how the StorageGRID system implements security, authentication, and authorization for the REST API.

StorageGRID uses the following security measures.

- Client communications with the Load Balancer service use HTTPS if HTTPS is configured for the load balancer endpoint.

When you configure a load balancer endpoint, HTTP can optionally be enabled. For example, you might want to use HTTP for testing or other non-production purposes. See the instructions for administering StorageGRID for more information.

- By default, StorageGRID uses HTTPS for client communications with Storage Nodes and the CLB service on Gateway Nodes.

HTTP can optionally be enabled for these connections. For example, you might want to use HTTP for testing or other non-production purposes. See the instructions for administering StorageGRID for more information.



The CLB service is deprecated.

- 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

Clients can connect to the Load Balancer service on Gateway Nodes or Admin Nodes, directly to Storage Nodes, or to the CLB service on Gateway Nodes.

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 do so using the certificate that was configured for the specific load balancer endpoint used to make the connection. Each endpoint has its own certificate, which is either a custom server certificate uploaded by the grid administrator or a certificate that the grid administrator generated in StorageGRID when configuring the endpoint.
- When client applications connect directly to a Storage Node or to the CLB service on Gateway Nodes, 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.

Clients should be configured to trust the certificate authority that signed whichever certificate they use to establish TLS connections.

See the instructions for administering StorageGRID for information on configuring load balancer endpoints, and for instructions on adding a single custom server certificate for TLS connections directly to Storage Nodes or to the CLB service on Gateway Nodes.

### Summary

The following table shows how security issues are implemented in the S3 and Swift REST APIs:

Security issue	Implementation for REST API
Connection security	TLS
Server authentication	X.509 server certificate signed by system CA or custom server certificate supplied by administrator

Security issue	Implementation for REST API
Client authentication	<ul style="list-style-type: none"> <li>• S3: S3 account (access key ID and secret access key)</li> <li>• Swift: Swift account (user name and password)</li> </ul>
Client authorization	<ul style="list-style-type: none"> <li>• S3: Bucket ownership and all applicable access control policies</li> <li>• Swift: Administrator role access</li> </ul>

### Related information

[Administer StorageGRID](#)

### Supported hashing and encryption algorithms for TLS libraries

The StorageGRID system supports a limited set of cipher suites that client applications can use when establishing a Transport Layer Security (TLS) session.

#### 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.

#### Supported cipher suites

TLS version	IANA name of cipher suite
1.2	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
1.2	TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256
1.2	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
1.3	TLS_AES_256_GCM_SHA384
1.3	TLS_CHACHA20_POLY1305_SHA256
1.3	TLS_AES_128_GCM_SHA256

#### Deprecated cipher suites

The following cipher suites are deprecated. Support for these ciphers will be removed in a future release.

IANA Name
TLS_RSA_WITH_AES_128_GCM_SHA256

<b>IANA Name</b>
TLS_RSA_WITH_AES_256_GCM_SHA384

### Related information

[How client connections can be configured](#)

## Monitoring and auditing operations

You can monitor workloads and efficiencies for client operations by viewing transaction trends for the entire grid, or for specific nodes. You can use audit messages to monitor client operations and transactions.

- [Monitoring object ingest and retrieval rates](#)
- [Accessing and reviewing audit logs](#)

### Monitoring object ingest and retrieval rates

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.

#### Steps

1. Sign in to the Grid Manager using a supported browser.
2. On the Dashboard, locate the Protocol Operations section.

This section summarizes the number of client operations performed by your StorageGRID system. Protocol rates are averaged over the last two minutes.

3. Select **Nodes**.
4. From the Nodes home page (deployment level), click the **Load Balancer** tab.

The charts show trends for all client traffic directed to load balancer endpoints within the grid. You can select a time interval in hours, days, weeks, months, or years, or you can apply a custom interval.

5. From the Nodes home page (deployment level), click the **Objects** tab.

The chart shows ingest and retrieve rates for your entire StorageGRID system in bytes per second and total bytes. You can select a time interval in hours, days, weeks, months, or years, or you can apply a custom interval.

6. To see information for a particular Storage Node, select the node from the list on the left, and click the **Objects** tab.

The chart shows the object ingest and retrieval rates for this Storage Node. The tab also includes metrics for object counts, queries, and verification. You can click the labels to see the definitions of these metrics.

Overview Hardware Network Storage **Objects** ILM Events Tasks

1 hour 1 day 1 week 1 month 1 year Custom

**S3 Ingest and Retrieve**

1.00 Bps  
0.75 Bps  
0.50 Bps  
0.25 Bps  
0 Bps

08:00 08:10 08:20 08:30 08:40 08:50

— Ingest rate — Retrieve rate

**Swift Ingest and Retrieve**

1.00 Bps  
0.75 Bps  
0.50 Bps  
0.25 Bps  
0 Bps

08:00 08:10 08:20 08:30 08:40 08:50

— Ingest rate — Retrieve rate

**Object Counts**

Total Objects	0	<a href="#">↗</a>
Lost Objects	0	<a href="#">↗</a>
S3 Buckets and Swift Containers	0	<a href="#">↗</a>

Help ✕

**S3 Buckets and Swift Containers**

The total number of S3 buckets and Swift containers known by this grid node.

Queries - Successful	12,364	<a href="#">↗</a>
Queries - Failed (timed-out)	0	<a href="#">↗</a>
Queries - Failed (consistency level unmet)	0	<a href="#">↗</a>

**Verification**

Status	No Errors	<a href="#">↗</a>
Rate Setting	Adaptive	<a href="#">↗</a>
Percent Complete	0.00%	<a href="#">↗</a>
Average Stat Time	0.00 microseconds	<a href="#">↗</a>
Objects Verified	0	<a href="#">↗</a>
Object Verification Rate	0.00 objects / second	<a href="#">↗</a>
Data Verified	0 bytes	<a href="#">↗</a>
Data Verification Rate	0.00 bytes / second	<a href="#">↗</a>
Missing Objects	0	<a href="#">↗</a>
Corrupt Objects	0	<a href="#">↗</a>
Quarantined Objects	0	<a href="#">↗</a>

7. If you want even more detail:

- a. Select **Support > Tools > Grid Topology**.
- b. Select **site > Overview > Main**.

The API Operations section displays summary information for the entire grid.

- c. Select **Storage Node > LDR > client application > Overview > Main**

The Operations section displays summary information for the selected Storage Node.

## Accessing and reviewing audit logs

Audit messages are generated by StorageGRID services and stored in text log files. API-specific audit messages in the audit logs provide critical security, operation, and performance monitoring data that can help you evaluate the health of your system.

### What you'll need

- You must have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of an Admin Node.

### About this task

The active audit log file is named `audit.log`, and it is stored on Admin Nodes.

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`.

After a day, the saved file is compressed and renamed, in the format `yyyy-mm-dd.txt.gz`, which preserves the original date.

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
```

### 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.
2. Go to the directory containing the audit log files:

```
cd /var/local/audit/export
```

3. View the current or a saved audit log file, as required.

### S3 operations tracked in the audit logs

Several bucket operations and object operations are tracked in the StorageGRID audit logs.

## Bucket operations tracked in the audit logs

- DELETE Bucket
- DELETE Bucket tagging
- DELETE Multiple Objects
- GET Bucket (List Objects)
- GET Bucket Object versions
- GET Bucket tagging
- HEAD Bucket
- PUT Bucket
- PUT Bucket compliance
- PUT Bucket tagging
- PUT Bucket versioning

## Object operations tracked in the audit logs

- Complete Multipart Upload
- Upload Part (when the ILM rule uses the Strict or Balanced ingest behaviors)
- Upload Part - Copy (when the ILM rule uses the Strict or Balanced ingest behaviors)
- DELETE Object
- GET Object
- HEAD Object
- POST Object restore
- PUT Object
- PUT Object - Copy

### Related information

[Operations on buckets](#)

[Operations on objects](#)

## 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

### Related information

- [Benefits of keeping idle HTTP connections open](#)
- [Benefits of active HTTP connections](#)
- [Benefits of concurrent HTTP connections](#)
- [Separation of HTTP connection pools for read and write operations](#)

### **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 or to the CLB service (deprecated) on Gateway 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 or to the CLB service, limiting active HTTP connections provides the following benefits:

- Enables optimal load balancing across the StorageGRID system.

When using the CLB service, you should prevent long-lived TCP/IP connections to optimize load balancing across the StorageGRID system. You should configure client applications to track the duration of each HTTP connection and close the HTTP connection after a set time so that the HTTP connection can be reestablished and rebalanced.

The CLB service balances load across the StorageGRID system at the time that a client application establishes an HTTP connection. 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.



The CLB service is deprecated.

- 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 may 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.

## Use Swift

Learn how client applications can use the OpenStack Swift API to interface with the StorageGRID system.

- [OpenStack Swift API support in StorageGRID](#)
- [Configuring tenant accounts and connections](#)
- [Swift REST API supported operations](#)
- [StorageGRID Swift REST API operations](#)
- [Configuring security for the REST API](#)
- [Monitoring and auditing operations](#)

## OpenStack Swift API support in StorageGRID

StorageGRID supports the following specific versions of Swift and HTTP.

Item	Version
Swift specification	OpenStack Swift Object Storage API v1 as of November 2015
HTTP	1.1 For more information about HTTP, see HTTP/1.1 (RFCs 7230-35).  <b>Note:</b> StorageGRID does not support HTTP/1.1 pipelining.

### Related information

[OpenStack: Object Storage API](#)

### History of Swift API support in StorageGRID

You should be aware of changes to the StorageGRID system's support for the Swift REST API.

Release	Comments
11.5	Removed Weak consistency control. The Available consistency level will be used instead.
11.4	Added support for TLS 1.3 and updated list of supported TLS cipher suites. CLB is deprecated. Added description of interrelationship between ILM and consistency setting.
11.3	Updated PUT Object operations to describe the impact of ILM rules that use synchronous placement at ingest (the Balanced and Strict options for Ingest Behavior). Added description of client connections that use load balancer endpoints or high availability groups. Updated list of supported TLS cipher suites. TLS 1.1 ciphers are no longer supported.
11.2	Minor editorial changes to document.
11.1	Added support for using HTTP for Swift client connections to grid nodes. Updated the definitions of consistency controls.
11.0	Added support for 1,000 containers for each tenant account.
10.3	Administrative updates and corrections to the document. Removed sections for configuring custom server certificates.
10.2	Initial support of the Swift API by the StorageGRID system. The currently supported version is OpenStack Swift Object Storage API v1.

## How StorageGRID implements the Swift REST API

A client application can use Swift REST API calls to connect to Storage Nodes and Gateway Nodes to create containers and to store and retrieve objects. This enables service-oriented applications developed for OpenStack Swift to connect with on-premise object storage provided by the StorageGRID system.

### Swift object management

After Swift objects have been ingested in the StorageGRID system, they are managed by the information lifecycle management (ILM) rules in the system's active ILM policy. The ILM rules and policy determine how StorageGRID creates and distributes copies of object data and how it manages those copies over time. For example, an ILM rule might apply to objects in specific Swift containers and might specify that multiple object copies be saved to several data centers for a certain number of years.

Contact your StorageGRID administrator if you need to understand how the grid's ILM rules and policies will affect the objects in your Swift tenant account.

### **Conflicting client requests**

Conflicting client requests, such as a 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 Swift clients begin an operation.

### **Consistency guarantees and controls**

By default, StorageGRID provides read-after-write consistency for newly created objects and eventual consistency for object updates and HEAD operations. 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.

StorageGRID also allows you to control consistency on a per container basis. You can change the consistency control to make a trade-off between the availability of the objects and the consistency of those objects across different Storage Nodes and sites, as required by your application.

### **Related information**

[Manage objects with ILM](#)

[GET container consistency request](#)

[PUT container consistency request](#)

### **Recommendations for implementing the Swift REST API**

You should follow these recommendations when implementing the Swift 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 do not expect the object to actually exist, you should use the “Available” consistency control. For example, you should use the “Available” consistency control if your application performs a HEAD operation to a location before performing a PUT operation to that location.

Otherwise, if the HEAD operation does not find the object, you might receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable.

You can set the “Available” consistency control for each container using the PUT container consistency request.

#### **Recommendations for object names**

You should not use random values as the first four characters of object names. Instead, you should use non-random, non-unique prefixes, such as image.

If you do need to use random and unique characters in object name prefixes, you should prefix the object names with a directory name. That is, use this format:

```
mycontainer/mydir/f8e3-image3132.jpg
```

Instead of this format:

```
mycontainer/f8e3-image3132.jpg
```

### Recommendations for “range reads”

If the **Compress Stored Objects** option is selected (**Configuration > System Settings > Grid Options**), Swift client applications should avoid performing GET object operations that specify a range of bytes be returned. These “range read” operations are inefficient because StorageGRID must effectively uncompress the objects to access the requested bytes. GET Object operations that request a small range of bytes from a very large object are especially inefficient; for example, it is very 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.

### Related information

[GET container consistency request](#)

[PUT container consistency request](#)

[Administer StorageGRID](#)

## Configuring tenant accounts and connections

Configuring StorageGRID to accept connections from client applications requires creating one or more tenant accounts and setting up the connections.

### Creating and configuring Swift tenant accounts

A Swift tenant account is required before Swift API clients can store and retrieve objects on StorageGRID. Each tenant account has its own account ID, groups and users, and containers and objects.

Swift tenant accounts are created by a StorageGRID grid administrator using the Grid Manager or the Grid Management API.

When creating a Swift tenant account, the grid administrator specifies the following information:

- Display name for the tenant (the tenant’s account ID is assigned automatically and cannot be changed)
- Optionally, a storage quota for the tenant account—the maximum number of gigabytes, terabytes, or petabytes available for the tenant’s objects. A tenant’s storage quota represents a logical amount (object size), not a physical amount (size on disk).
- If single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid’s identity source, and the initial password for the tenant’s local root user.

- If SSO is enabled, which federated group has Root Access permission to configure the tenant account.

After a Swift tenant account is created, users with the Root Access permission can access the Tenant Manager to perform tasks such as the following:

- Setting up identity federation (unless the identity source is shared with the grid), and creating local groups and users
- Monitoring storage usage



Swift users must have the Root Access permission to access the Tenant Manager. However, the Root Access permission does not allow users to authenticate into the Swift REST API to create containers and ingest objects. Users must have the Swift Administrator permission to authenticate into the Swift REST API.

### Related information

[Administer StorageGRID](#)

[Use a tenant account](#)

[Supported Swift API endpoints](#)

### How client connections can be configured

A grid administrator makes configuration choices that affect how Swift clients connect to StorageGRID to store and retrieve data. The specific information you need to make a connection depends upon the configuration that was chosen.

Client applications can store or retrieve objects by connecting to any of the following:

- The Load Balancer service on Admin Nodes or Gateway Nodes, or optionally, the virtual IP address of a high availability (HA) group of Admin Nodes or Gateway Nodes
- The CLB service on Gateway Nodes, or optionally, the virtual IP address of a high availability group of Gateway Nodes



The CLB service is deprecated. Clients configured before the StorageGRID 11.3 release can continue to use the CLB service on Gateway Nodes. All other client applications that depend on StorageGRID to provide load balancing should connect using the Load Balancer service.

- Storage Nodes, with or without an external load balancer

When configuring StorageGRID, a grid administrator can use the Grid Manager or the Grid Management API to perform the following steps, all of which are optional:

1. Configure endpoints for the Load Balancer service.

You must configure endpoints to use the Load Balancer service. The Load Balancer service on Admin Nodes or Gateway Nodes distributes incoming network connections from client applications to Storage Nodes. When creating a load balancer endpoint, the StorageGRID administrator specifies a port number, whether the endpoint accepts HTTP or HTTPS connections, the type of client (S3 or Swift) that will use the endpoint, and the certificate to be used for HTTPS connections (if applicable).

2. Configure Untrusted Client Networks.

If a StorageGRID administrator configures a node's Client Network to be untrusted, the node only accepts inbound connections on the Client Network on ports that are explicitly configured as load balancer endpoints.

### 3. Configure high availability groups.

If an administrator creates an HA group, the network interfaces of multiple Admin Nodes or Gateway Nodes are placed into an active-backup configuration. Client connections are made using the virtual IP address of the HA group.

For more information about each option, see the instructions for administering StorageGRID.

#### Summary: IP addresses and ports for client connections

Client applications connect to StorageGRID using the IP address of a grid node and the port number of a service on that node. If high availability (HA) groups are configured, client applications can connect using the virtual IP address of the HA group.

#### Information required to make client connections

The 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. Contact your StorageGRID administrator for more information, or see the instructions for administering StorageGRID for a description of how to find this information 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	<ul style="list-style-type: none"> <li>Load balancer endpoint port</li> </ul>
HA group	CLB <b>Note:</b> The CLB service is deprecated.	Virtual IP address of an HA group	Default Swift ports: <ul style="list-style-type: none"> <li>HTTPS: 8083</li> <li>HTTP: 8085</li> </ul>
Admin Node	Load Balancer	IP address of the Admin Node	<ul style="list-style-type: none"> <li>Load balancer endpoint port</li> </ul>
Gateway Node	Load Balancer	IP address of the Gateway Node	<ul style="list-style-type: none"> <li>Load balancer endpoint port</li> </ul>
Gateway Node	CLB <b>Note:</b> The CLB service is deprecated.	IP address of the Gateway Node <b>Note:</b> By default, HTTP ports for CLB and LDR are not enabled.	Default Swift ports: <ul style="list-style-type: none"> <li>HTTPS: 8083</li> <li>HTTP: 8085</li> </ul>

Where connection is made	Service that client connects to	IP address	Port
Storage Node	LDR	IP address of Storage Node	Default Swift ports: <ul style="list-style-type: none"> <li>• HTTPS: 18083</li> <li>• HTTP: 18085</li> </ul>

### Example

To connect a Swift client 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 192.0.2.6 and the port number of a Swift Load Balancer endpoint is 10444, then a Swift client could use the following URL to connect to StorageGRID:

- `https://192.0.2.6:10444`

It is possible to configure a DNS name for the IP address that clients use to connect to StorageGRID. Contact your local network administrator.

### Deciding to use HTTPS or HTTP connections

When client connections are made using a Load Balancer endpoint, connections must be made using the protocol (HTTP or HTTPS) that was specified for that endpoint. To use HTTP for client connections to Storage Nodes or to the CLB service on Gateway Nodes, you must enable its use.

By default, when client applications connect to Storage Nodes or the CLB service on Gateway Nodes, they must use encrypted HTTPS for all connections. Optionally, you can enable less-secure HTTP connections by selecting the **Enable HTTP Connection** grid option in the Grid Manager. For example, a client application might use HTTP when testing the connection to a Storage Node in a non-production environment.



Be careful when enabling HTTP for a production grid since requests will be sent unencrypted.



The CLB service is deprecated.

If the **Enable HTTP Connection** option is selected, clients must use different ports for HTTP than they use for HTTPS. See the instructions for administering StorageGRID.

### Related information

[Administer StorageGRID](#)

### Testing your connection in the Swift API configuration

You can use the Swift CLI to test your connection to the StorageGRID system and to verify that you can read and write objects to the system.

### What you'll need

- You must have downloaded and installed `python-swiftclient`, the Swift command-line client.

- You must have a Swift tenant account in the StorageGRID system.

### About this task

If you have not configured security, you must add the `--insecure` flag to each of these commands.

### Steps

1. Query the info URL for your StorageGRID Swift deployment:

```
swift
-U <Tenant_Account_ID:Account_User_Name>
-K <User_Password>
-A https://<FQDN | IP>:<Port>/info
capabilities
```

This is sufficient to test that your Swift deployment is functional. To further test account configuration by storing an object, continue with the additional steps.

2. Put an object in the container:

```
touch test_object
swift
-U <Tenant_Account_ID:Account_User_Name>
-K <User_Password>
-A https://<FQDN | IP>:<Port>/auth/v1.0
upload test_container test_object
--object-name test_object
```

3. Get the container to verify the object:

```
swift
-U <Tenant_Account_ID:Account_User_Name>
-K <User_Password>
-A https://<FQDN | IP>:<Port>/auth/v1.0
list test_container
```

4. Delete the object:

```
swift
-U <Tenant_Account_ID:Account_User_Name>
-K <User_Password>
-A https://<FQDN | IP>:<Port>/auth/v1.0
delete test_container test_object
```

## 5. Delete the container:

```
swift
-U `<_Tenant_Account_ID:Account_User_Name_>`
-K `<_User_Password_>`
-A `https://<_FQDN_ | _IP_>:<_Port_>/auth/v1.0'
delete test_container
```

### Related information

[Creating and configuring Swift tenant accounts](#)

[Configuring security for the REST API](#)

## Swift REST API supported operations

The StorageGRID system supports most operations in the OpenStack Swift API. Before integrating Swift REST API clients with StorageGRID, review the implementation details for account, container, and object operations.

### Operations supported in StorageGRID

The following Swift API operations are supported:

- [Account operations](#)
- [Container operations](#)
- [Object operations](#)

### Common response headers for all operations

The StorageGRID system implements all common headers for supported operations as defined by the OpenStack Swift Object Storage API v1.

### Related information

[OpenStack: Object Storage API](#)

### Supported Swift API endpoints

StorageGRID supports the following Swift API endpoints: the info URL, the auth URL, and the storage URL.

#### info URL

You can determine the capabilities and limitations of the StorageGRID Swift implementation by issuing a GET request to the Swift base URL with the /info path.

```
https://FQDN | Node IP:Swift Port/info/
```

In the request:

- *FQDN* is the fully qualified domain name.
- *Node IP* is the IP address for the Storage Node or the Gateway Node on the StorageGRID network.
- *Swift Port* is the port number used for Swift API connections on the Storage Node or Gateway Node.

For example, the following info URL would request information from a Storage Node with the IP address of 10.99.106.103 and using port 18083.

```
https://10.99.106.103:18083/info/
```

The response includes the capabilities of the Swift implementation as a JSON dictionary. A client tool can parse the JSON response to determine the capabilities of the implementation and use them as constraints for subsequent storage operations.

The StorageGRID implementation of Swift allows unauthenticated access to the info URL.

### auth URL

A client can use the Swift auth URL to authenticate as a tenant account user.

```
https://FQDN | Node IP:Swift Port/auth/v1.0/
```

You must provide the tenant account ID, user name, and password as parameters in the X-Auth-User and X-Auth-Key request headers, as follows:

```
X-Auth-User: Tenant_Account_ID:Username
```

```
X-Auth-Key: Password
```

In the request headers:

- *Tenant\_Account\_ID* is the account ID assigned by StorageGRID when the Swift tenant was created. This is the same tenant account ID used on the Tenant Manager sign-in page.
- *Username* is the name of a tenant user that has been created in the Tenant Manager. This user must belong to a group that has the Swift Administrator permission. The tenant's root user cannot be configured to use the Swift REST API.

If Identity Federation is enabled for the tenant account, provide the username and password of the federated user from the LDAP server. Alternatively, provide the LDAP user's domain name. For example:

```
X-Auth-User: Tenant_Account_ID:Username@Domain_Name
```

- *Password* is the password for the tenant user. User passwords are created and managed in the Tenant Manager.

The response to a successful authentication request returns a storage URL and an auth token, as follows:

```
X-Storage-Url: https://FQDN | Node_IP:Swift_Port/v1/Tenant_Account_ID
```

```
X-Auth-Token: token
```

```
X-Storage-Token: token
```

By default, the token is valid for 24 hours from generation time.

Tokens are generated for a specific tenant account. A valid token for one account does not authorize a user to access another account.

### storage URL

A client application can issue Swift REST API calls to perform supported account, container, and object operations against a Gateway Node or Storage Node. Storage requests are addressed to the storage URL returned in the authentication response. The request must also include the X-Auth-Token header and value returned from the auth request.

```
https://FQDN | IP:Swift_Port/v1/Tenant_Account_ID
```

```
[/container] [/object]
```

```
X-Auth-Token: token
```

Some storage response headers that contain usage statistics might not reflect accurate numbers for recently modified objects. It might take a few minutes for accurate numbers to appear in these headers.

The following response headers for account and container operations are examples of those that contain usage statistics:

- X-Account-Bytes-Used
- X-Account-Object-Count
- X-Container-Bytes-Used
- X-Container-Object-Count

### Related information

[How client connections can be configured](#)

[Creating and configuring Swift tenant accounts](#)

[Account operations](#)

[Container operations](#)

[Object operations](#)

### Account operations

The following Swift API operations are performed on accounts.

#### GET account

This operation retrieves the container list associated with the account and account usage statistics.

The following request parameter is required:

- Account

The following request header is required:

- X-Auth-Token

The following supported request query parameters are optional:

- Delimiter
- End\_marker
- Format
- Limit
- Marker
- Prefix

A successful execution returns the following headers with an “HTTP/1.1 204 No Content” response if the account is found and has no containers or the container list is empty; or an “HTTP/1.1 200 OK” response if the account is found and the container list is not empty:

- Accept-Ranges
- Content-Length
- Content-Type
- Date
- X-Account-Bytes-Used
- X-Account-Container-Count
- X-Account-Object-Count
- X-Timestamp
- X-Trans-Id

#### **HEAD account**

This operation retrieves account information and statistics from a Swift account.

The following request parameter is required:

- Account

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an “HTTP/1.1 204 No Content” response:

- Accept-Ranges
- Content-Length
- Date
- X-Account-Bytes-Used
- X-Account-Container-Count

- X-Account-Object-Count
- X-Timestamp
- X-Trans-Id

### Related information

[Swift operations tracked in the audit logs](#)

### Container operations

StorageGRID supports a maximum of 1,000 containers per Swift account. The following Swift API operations are performed on containers.

#### DELETE container

This operation removes an empty container from a Swift account in a StorageGRID system.

The following request parameters are required:

- Account
- Container

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an "HTTP/1.1 204 No Content" response:

- Content-Length
- Content-Type
- Date
- X-Trans-Id

#### GET container

This operation retrieves the object list associated with the container along with container statistics and metadata in a StorageGRID system.

The following request parameters are required:

- Account
- Container

The following request header is required:

- X-Auth-Token

The following supported request query parameters are optional:

- Delimiter

- End\_marker
- Format
- Limit
- Marker
- Path
- Prefix

A successful execution returns the following headers with an "HTTP/1.1 200 Success" or a "HTTP/1.1 204 No Content" response:

- Accept-Ranges
- Content-Length
- Content-Type
- Date
- X-Container-Bytes-Used
- X-Container-Object-Count
- X-Timestamp
- X-Trans-Id

#### **HEAD container**

This operation retrieves container statistics and metadata from a StorageGRID system.

The following request parameters are required:

- Account
- Container

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an "HTTP/1.1 204 No Content" response:

- Accept-Ranges
- Content-Length
- Date
- X-Container-Bytes-Used
- X-Container-Object-Count
- X-Timestamp
- X-Trans-Id

## PUT container

This operation creates a container for an account in a StorageGRID system.

The following request parameters are required:

- Account
- Container

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an "HTTP/1.1 201 Created" or "HTTP/1.1 202 Accepted" (if the container already exists under this account) response:

- Content-Length
- Date
- X-Timestamp
- X-Trans-Id

A container name must be unique in the StorageGRID namespace. If the container exists under another account, the following header is returned: "HTTP/1.1 409 Conflict."

## Related information

[Swift operations tracked in the audit logs](#)

## Object operations

The following Swift API operations are performed on objects.

### DELETE object

This operation deletes an object's content and metadata from the StorageGRID system.

The following request parameters are required:

- Account
- Container
- Object

The following request header is required:

- X-Auth-Token

A successful execution returns the following response headers with an HTTP/1.1 204 No Content response:

- Content-Length
- Content-Type

- Date
- X-Trans-Id

When processing a DELETE Object 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 cannot 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.

For more information on how objects are deleted, see the instructions for managing objects with information lifecycle management.

### **GET object**

This operation retrieves the object content and gets the object metadata from a StorageGRID system.

The following request parameters are required:

- Account
- Container
- Object

The following request header is required:

- X-Auth-Token

The following request headers are optional:

- Accept-Encoding
- If-Match
- If-Modified-Since
- If-None-Match
- If-Unmodified-Since
- Range

A successful execution returns the following headers with an HTTP/1.1 200 OK response:

- Accept-Ranges
- Content-Disposition, returned only if Content-Disposition metadata was set
- Content-Encoding, returned only if Content-Encoding metadata was set
- Content-Length
- Content-Type
- Date
- ETag
- Last-Modified

- X-Timestamp
- X-Trans-Id

### HEAD object

This operation retrieves metadata and properties of an ingested object from a StorageGRID system.

The following request parameters are required:

- Account
- Container
- Object

The following request header is required:

- X-Auth-Token

A successful execution returns the following headers with an "HTTP/1.1 200 OK" response:

- Accept-Ranges
- Content-Disposition, returned only if Content-Disposition metadata was set
- Content-Encoding, returned only if Content-Encoding metadata was set
- Content-Length
- Content-Type
- Date
- ETag
- Last-Modified
- X-Timestamp
- X-Trans-Id

### PUT object

This operation creates a new object with data and metadata, or replaces an existing object with data and metadata in a StorageGRID system.

StorageGRID supports objects up to 5 TB in size.



Conflicting client requests, such as a 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 Swift clients begin an operation.

The following request parameters are required:

- Account
- Container

- Object

The following request header is required:

- X-Auth-Token

The following request headers are optional:

- Content-Disposition
- Content-Encoding

Do not use chunked `Content-Encoding` if the ILM rule that applies to an object filters objects based on size and uses synchronous placement on ingest (the `Balanced` or `Strict` options for `Ingest Behavior`).

- Transfer-Encoding

Do not use compressed or chunked `Transfer-Encoding` if the ILM rule that applies to an object filters objects based on size and uses synchronous placement on ingest (the `Balanced` or `Strict` options for `Ingest Behavior`).

- Content-Length

If an ILM rule filters objects by size and uses synchronous placement on ingest, you must specify `Content-Length`.



If you do not follow these guidelines for `Content-Encoding`, `Transfer-Encoding`, and `Content-Length`, StorageGRID must save the object before it can determine object size and apply the ILM rule. In other words, StorageGRID must default to creating interim copies of an object on ingest. That is, StorageGRID must use the `Dual Commit` option for `Ingest Behavior`.

For more information about synchronous placement and ILM rules, see the instructions for managing objects with information lifecycle management.

- Content-Type
- ETag
- X-Object-Meta-`<name\>` (object-related metadata)

If you want to use the **User Defined Creation Time** option as the Reference Time for an ILM rule, you must store the value in a user-defined header named `X-Object-Meta-Creation-Time`. For example:

```
X-Object-Meta-Creation-Time: 1443399726
```

This field is evaluated as seconds since January 1, 1970.

- X-Storage-Class: `reduced_redundancy`

This header affects how many object copies StorageGRID creates if the ILM rule that matches an ingested object specifies an `Ingest Behavior` of `Dual Commit` or `Balanced`.

- **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 cannot immediately make all copies specified in the rule. If StorageGRID can perform synchronous placement, this header has no effect.

The `reduced_redundancy` header 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` header is not recommended in other circumstances because it increases the risk the loss of object data 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.

Note that 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 policy and does not result in data being stored at lower levels of redundancy in the StorageGRID system.

A successful execution returns the following headers with an "HTTP/1.1 201 Created" response:

- Content-Length
- Content-Type
- Date
- ETag
- Last-Modified
- X-Trans-Id

### Related information

[Manage objects with ILM](#)

[Swift operations tracked in the audit logs](#)

### OPTIONS request

The OPTIONS request checks the availability of an individual Swift service. The OPTIONS request is processed by the Storage Node or Gateway Node specified in the URL.

### OPTIONS method

For example, client applications can issue an OPTIONS request to the Swift port on a Storage Node, without providing Swift 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.

When used with the info URL or the storage URL, the OPTIONS method returns a list of supported verbs for the given URL (for example, HEAD, GET, OPTIONS, and PUT). The OPTIONS method cannot be used with the auth URL.

The following request parameter is required:

- Account

The following request parameters are optional:

- Container
- Object

A successful execution returns the following headers with an “HTTP/1.1 204 No Content” response. The OPTIONS request to the storage URL does not require that the target exists.

- Allow (a list of supported verbs for the given URL, for example, HEAD, GET, OPTIONS, and PUT)
- Content-Length
- Content-Type
- Date
- X-Trans-Id

### Related information

[Supported Swift API endpoints](#)

### Error responses to Swift API operations

Understanding the possible error responses can help you troubleshoot operations.

The following HTTP status codes might be returned when errors occur during an operation:

Swift error name	HTTP status
AccountNameTooLong, ContainerNameTooLong, HeaderTooBig, InvalidContainerName, InvalidRequest, InvalidURI, MetadataNameTooLong, MetadataValueTooBig, MissingSecurityHeader, ObjectNameTooLong, TooManyContainers, TooManyMetadataItems, TotalMetadataTooLarge	400 Bad Request
AccessDenied	403 Forbidden
ContainerNotEmpty, ContainerAlreadyExists	409 Conflict
InternalError	500 Internal Server Error
InvalidRange	416 Requested Range Not Satisfiable

Swift error name	HTTP status
MethodNotAllowed	405 Method Not Allowed
MissingContentLength	411 Length Required
NotFound	404 Not Found
NotImplemented	501 Not Implemented
PreconditionFailed	412 Precondition Failed
ResourceNotFound	404 Not Found
Unauthorized	401 Unauthorized
UnprocessableEntity	422 Unprocessable Entity

## StorageGRID Swift REST API operations

There are operations added on to the Swift REST API that are specific to StorageGRID system.

### GET container consistency request

Consistency level makes a trade-off between the availability of the objects and the consistency of those objects across different Storage Nodes and sites. The GET container consistency request allows you to determine the consistency level being applied to a particular container.

#### Request

Request HTTP Header	Description
X-Auth-Token	Specifies the Swift authentication token for the account to use for the request.
x-ntap-sg-consistency	Specifies the type of request, where <code>true</code> = GET container consistency, and <code>false</code> = GET container.
Host	The hostname to which the request is directed.

#### Request example

```
GET /v1/28544923908243208806/Swift container
X-Auth-Token: SGRD_3a877009a2d24cb1801587bfa9050f29
x-ntap-sg-consistency: true
Host: test.com
```

## Response

Response HTTP Header	Description
Date	The date and time of the response.
Connection	Whether the connection to the server is open or closed.
X-Trans-Id	The unique transaction identifier for the request.
Content-Length	The length of the response body.
x-ntap-sg-consistency	<p>The consistency control level being applied to the container. The following values are supported:</p> <ul style="list-style-type: none"><li>• <b>all</b>: All nodes receive the data immediately or the request will fail.</li><li>• <b>strong-global</b>: Guarantees read-after-write consistency for all client requests across all sites.</li><li>• <b>strong-site</b>: Guarantees read-after-write consistency for all client requests within a site.</li><li>• <b>read-after-new-write</b>: Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees.</li></ul> <p><b>Note:</b> If your application uses HEAD requests on objects that do not exist, you might receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable. To prevent these errors, use the “available” level.</p> <ul style="list-style-type: none"><li>• <b>available</b> (eventual consistency for HEAD operations): Behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than “read-after-new-write” if Storage Nodes are unavailable.</li></ul>

## Response example

```
HTTP/1.1 204 No Content
Date: Sat, 29 Nov 2015 01:02:18 GMT
Connection: CLOSE
X-Trans-Id: 1936575373
Content-Length: 0
x-ntap-sg-consistency: strong-site
```

## Related information

[Use a tenant account](#)

## PUT container consistency request

The PUT container consistency request allows you to specify the consistency level to apply to operations performed on a container. By default, new containers are created using the “read-after-new-write” consistency level.

## Request

Request HTTP Header	Description
X-Auth-Token	The Swift authentication token for the account to use for the request.

Request HTTP Header	Description
x-ntap-sg-consistency	<p>The consistency control level to apply to operations on the container. The following values are supported:</p> <ul style="list-style-type: none"> <li>• <b>all</b>: All nodes receive the data immediately or the request will fail.</li> <li>• <b>strong-global</b>: Guarantees read-after-write consistency for all client requests across all sites.</li> <li>• <b>strong-site</b>: Guarantees read-after-write consistency for all client requests within a site.</li> <li>• <b>read-after-new-write</b>: Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees.</li> </ul> <p><b>Note:</b> If your application uses HEAD requests on objects that do not exist, you might receive a high number of 500 Internal Server errors if one or more Storage Nodes are unavailable. To prevent these errors, use the “available” level.</p> <ul style="list-style-type: none"> <li>• <b>available</b> (eventual consistency for HEAD operations): Behaves the same as the “read-after-new-write” consistency level, but only provides eventual consistency for HEAD operations. Offers higher availability for HEAD operations than “read-after-new-write” if Storage Nodes are unavailable.</li> </ul>
Host	The hostname to which the request is directed.

### How consistency controls and ILM rules interact to affect data protection

Both your choice of consistency control and your ILM rule affect how objects are protected. These settings can interact.

For example, the consistency control 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 level and ingest behavior can provide better initial data protection and more predictable system responses.

The following ingest behaviors are available for ILM rules:

- **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.
- **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.



Before selecting the ingest behavior for an ILM rule, read the full description of these settings in the instructions for managing objects with information lifecycle management.

### Example of how the consistency control and ILM rule can interact

Suppose you have a two-site grid with the following ILM rule and the following consistency level setting:

- **ILM rule:** Create two object copies, one at the local site and one at a remote site. The Strict ingest behavior is selected.
- **Consistency level:** “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 level, 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 cannot be retrieved.

The inter-relationship between consistency levels and ILM rules can be complex. Contact NetApp if you require assistance.

### Request example

```
PUT /v1/28544923908243208806/_Swift container_  
X-Auth-Token: SGRD_3a877009a2d24cb1801587bfa9050f29  
x-ntap-sg-consistency: strong-site  
Host: test.com
```

### Response

Response HTTP Header	Description
Date	The date and time of the response.
Connection	Whether the connection to the server is open or closed.
X-Trans-Id	The unique transaction identifier for the request.
Content-Length	The length of the response body.

### Response example

```
HTTP/1.1 204 No Content
Date: Sat, 29 Nov 2015 01:02:18 GMT
Connection: CLOSE
X-Trans-Id: 1936575373
Content-Length: 0
```

### Related information

[Use a tenant account](#)

## Configuring security for the REST API

You should review the security measures implemented for the REST API and understand how to secure your system.

### How StorageGRID provides security for the REST API

You should understand how the StorageGRID system implements security, authentication, and authorization for the REST API.

StorageGRID uses the following security measures.

- Client communications with the Load Balancer service use HTTPS if HTTPS is configured for the load balancer endpoint.

When you configure a load balancer endpoint, HTTP can optionally be enabled. For example, you might want to use HTTP for testing or other non-production purposes. See the instructions for administering StorageGRID for more information.

- By default, StorageGRID uses HTTPS for client communications with Storage Nodes and the CLB service on Gateway Nodes.

HTTP can optionally be enabled for these connections. For example, you might want to use HTTP for testing or other non-production purposes. See the instructions for administering StorageGRID for more information.



The CLB service is deprecated.

- 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

Clients can connect to the Load Balancer service on Gateway Nodes or Admin Nodes, directly to Storage Nodes, or to the CLB service on Gateway Nodes.

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 do so using the certificate that was configured for the specific load balancer endpoint used to make the connection. Each endpoint has its own certificate, which is either a custom server certificate uploaded by the grid administrator or a certificate that the grid administrator generated in StorageGRID when configuring the endpoint.
- When client applications connect directly to a Storage Node or to the CLB service on Gateway Nodes, 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.

Clients should be configured to trust the certificate authority that signed whichever certificate they use to establish TLS connections.

See the instructions for administering StorageGRID for information on configuring load balancer endpoints, and for instructions on adding a single custom server certificate for TLS connections directly to Storage Nodes or to the CLB service on Gateway Nodes.

### Summary

The following table shows how security issues are implemented in the S3 and Swift REST APIs:

Security issue	Implementation for REST API
Connection security	TLS
Server authentication	X.509 server certificate signed by system CA or custom server certificate supplied by administrator
Client authentication	<ul style="list-style-type: none"> <li>• S3: S3 account (access key ID and secret access key)</li> <li>• Swift: Swift account (user name and password)</li> </ul>
Client authorization	<ul style="list-style-type: none"> <li>• S3: Bucket ownership and all applicable access control policies</li> <li>• Swift: Administrator role access</li> </ul>

### Related information

[Administer StorageGRID](#)

### Supported hashing and encryption algorithms for TLS libraries

The StorageGRID system supports a limited set of cipher suites that client applications can use when establishing a Transport Layer Security (TLS) session.

### 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.

## Supported cipher suites

TLS version	IANA name of cipher suite
1.2	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
	TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256
	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
1.3	TLS_AES_256_GCM_SHA384
	TLS_CHACHA20_POLY1305_SHA256
	TLS_AES_128_GCM_SHA256

## Deprecated cipher suites

The following cipher suites are deprecated. Support for these ciphers will be removed in a future release.

IANA Name
TLS_RSA_WITH_AES_128_GCM_SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384

## Related information

[How client connections can be configured](#)

## Monitoring and auditing operations

You can monitor workloads and efficiencies for client operations by viewing transaction trends for the entire grid, or for specific nodes. You can use audit messages to monitor client operations and transactions.

### Monitoring object ingest and retrieval rates

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.

### Steps

1. Sign in to the Grid Manager using a supported browser.
2. On the Dashboard, locate the Protocol Operations section.

This section summarizes the number of client operations performed by your StorageGRID system. Protocol rates are averaged over the last two minutes.

3. Select **Nodes**.

4. From the Nodes home page (deployment level), click the **Load Balancer** tab.


The charts show trends for all client traffic directed to load balancer endpoints within the grid. You can select a time interval in hours, days, weeks, months, or years, or you can apply a custom interval.

5. From the Nodes home page (deployment level), click the **Objects** tab.

The chart shows ingest and retrieve rates for your entire StorageGRID system in bytes per second and total bytes. You can select a time interval in hours, days, weeks, months, or years, or you can apply a custom interval.

6. To see information for a particular Storage Node, select the node from the list on the left, and click the **Objects** tab.

The chart shows the object ingest and retrieval rates for this Storage Node. The tab also includes metrics for object counts, queries, and verification. You can click the labels to see the definitions of these metrics.



7. If you want even more detail:

- a. Select **Support > Tools > Grid Topology**.
- b. Select **site > Overview > Main**.

The API Operations section displays summary information for the entire grid.

- c. Select **Storage Node > LDR > client application > Overview > Main**

The Operations section displays summary information for the selected Storage Node.

## Accessing and reviewing audit logs

Audit messages are generated by StorageGRID services and stored in text log files. API-specific audit messages in the audit logs provide critical security, operation, and performance monitoring data that can help you evaluate the health of your system.

### What you'll need

- You must have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of an Admin Node.

### About this task

The active audit log file is named `audit.log`, and it is stored on Admin Nodes.

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`.

After a day, the saved file is compressed and renamed, in the format `yyyy-mm-dd.txt.gz`, which preserves the original date.

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
```

### 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.
2. Go to the directory containing the audit log files: `cd /var/local/audit/export`
3. View the current or a saved audit log file, as required.

### Related information

[Review audit logs](#)

### Swift operations tracked in the audit logs

All successful storage DELETE, GET, HEAD, POST, and PUT operations are tracked in the StorageGRID audit log. Failures are not logged, nor are `info`, `auth`, or `OPTIONS` requests.

See *Understanding audit messages* for details about the information tracked for the following Swift operations.

## **Account operations**

- GET account
- HEAD account

## **Container operations**

- DELETE container
- GET container
- HEAD container
- PUT container

## **Object operations**

- DELETE object
- GET object
- HEAD object
- PUT object

## **Related information**

[Review audit logs](#)

[Account operations](#)

[Container operations](#)

[Object operations](#)

# Monitor and troubleshoot

## Monitor a StorageGRID system

Learn how to monitor a StorageGRID system and how to assess issues that might occur. Lists all system alerts.

- [Using the Grid Manager for monitoring](#)
- [Information you should monitor regularly](#)
- [Managing alerts and alarms](#)
- [Using SNMP monitoring](#)
- [Collecting additional StorageGRID data](#)
- [Troubleshooting a StorageGRID system](#)
- [Alerts reference](#)
- [Alarms reference \(legacy system\)](#)
- [Log files reference](#)

### Using the Grid Manager for monitoring

The Grid Manager is the most important tool for monitoring your StorageGRID system. This section introduces the Grid Manager Dashboard and provides detailed information about the Nodes pages.

- [Web browser requirements](#)
- [Viewing the Dashboard](#)
- [Viewing the Nodes page](#)

### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84


You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024

<b>Browser width</b>	<b>Pixels</b>
Optimum	1280

## Viewing the Dashboard

When you first sign in to the Grid Manager, you can use the Dashboard to monitor system activities at a glance. The Dashboard includes information about system health, usage metrics, and operational trends and charts.



## Health panel

Description	View additional details	Learn more
<p>Summarizes the system's health. A green checkmark means that there are no current alerts and all grid nodes are connected. Any other icon means that there is at least one current alert or disconnected node.</p>	<p>You might see one or more of the following links:</p> <ul style="list-style-type: none"> <li>• <b>Grid details:</b> Appears if any nodes are disconnected (connection state Unknown or Administratively Down). Click the link, or click the blue or gray icon to determine which node or nodes are affected.</li> <li>• <b>Current alerts:</b> Appears if any alerts are currently active. Click the link, or click <b>Critical, Major,</b> or <b>Minor</b> to see the details on the <b>Alerts &gt; Current</b> page.</li> <li>• <b>Recently resolved alerts:</b> Appears if any alerts triggered in the past week are now resolved. Click the link to see the details on the <b>Alerts &gt; Resolved</b> page.</li> <li>• <b>Legacy alarms:</b> Appears if any alarms (legacy system) are currently active. Click the link to see the details on the <b>Support &gt; Alarms (Legacy) &gt; Current Alarms</b> page.</li> <li>• <b>License:</b> Appears if there is an issue with the software license for this StorageGRID system. Click the link to see the details on the <b>Maintenance &gt; System &gt; License</b> page.</li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring node connection states</a></li> <li>• <a href="#">Viewing current alerts</a></li> <li>• <a href="#">Viewing resolved alerts</a></li> <li>• <a href="#">Viewing legacy alarms</a></li> <li>• <a href="#">Administer StorageGRID</a></li> </ul>

**Available Storage panel**

Description	View additional details	Learn more
<p>Displays the available and used storage capacity in the entire grid, not including archival media.</p> <p>The Overall chart presents grid-wide totals. If this is a multi-site grid, additional charts appear for each data center site.</p> <p>You can use this information to compare the used storage with the available storage. If you have a multi-site grid, you can determine which site is consuming more storage.</p>	<ul style="list-style-type: none"> <li>• To view the capacity, place your cursor over the chart's available and used capacity sections.</li> <li>• To view capacity trends over a date range, click the chart icon  for the overall grid, or for a data center site.</li> <li>• To see details, select <b>Nodes</b>. Then, view the Storage tab for the entire grid, an entire site, or a single Storage Node.</li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">Viewing the Storage tab</a></li> <li>• <a href="#">Monitoring storage capacity</a></li> </ul>

**Information Lifecycle Management (ILM) panel**


Description	View additional details	Learn more
<p>Displays current ILM operations and ILM queues for your system. You can use this information to monitor your system's workload.</p> <ul style="list-style-type: none"> <li>• <b>Awaiting - Client:</b> The total number of objects awaiting ILM evaluation from client operations (for example, ingest).</li> <li>• <b>Awaiting - Evaluation Rate:</b> The current rate at which objects are evaluated against the ILM policy in the grid.</li> <li>• <b>Scan Period - Estimated:</b> The estimated time to complete a full ILM scan of all objects. <b>Note:</b> A full scan does not guarantee that ILM has been applied to all objects.</li> </ul>	<ul style="list-style-type: none"> <li>• To see details, select <b>Nodes</b>. Then, view the ILM tab for the entire grid, an entire site, or a single Storage Node.</li> <li>• To see the existing ILM rules, select <b>ILM &gt; Rules</b>.</li> <li>• To see the existing ILM policies, select <b>ILM &gt; Policies</b>.</li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">Viewing the ILM tab</a></li> <li>• <a href="#">Administer StorageGRID.</a></li> </ul>

**Protocol Operations panel**

Description	View additional details	Learn more
<p>Displays the number of protocol-specific operations (S3 and Swift) performed by your system.</p> <p>You can use this information to monitor your system's workloads and efficiencies. Protocol rates are averaged over the last two minutes.</p>	<ul style="list-style-type: none"> <li>To see details, select <b>Nodes</b>. Then, view the Objects tab for the entire grid, an entire site, or a single Storage Node.</li> <li>To view trends over a date range, click the chart icon  to the right of the S3 or Swift protocol rate.</li> </ul>	<ul style="list-style-type: none"> <li><a href="#">Viewing the Objects tab</a></li> <li><a href="#">Use S3</a></li> <li><a href="#">Use Swift</a></li> </ul>

## Viewing 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.



From the tree view on the left, you can see all the sites and all the nodes in your StorageGRID system. The icon for each node indicates if the node is connected or if there are any active alerts.

### Connection state icons

If a node is disconnected from the grid, the tree view shows a blue or gray connection state icon, not the icon for any underlying alerts.

- Not connected - Unknown** : The node is not connected to the grid for an unknown reason. For example, the network connection between nodes has been lost or the power 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.

- **Not connected - 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.

### Alert icons

If a node is connected to the grid, the tree view shows one of the following icons, depending on if there are any current alerts for the node.

- **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 do not clear on their own to ensure they do not result in a more serious problem.
- **Normal** : No alerts are active, and the node is connected to the grid.

### Viewing details for a system, site, or node

To view the available information, click the appropriate links on the left, as follows:

- Select the grid name to see an aggregate summary of the statistics for your entire StorageGRID system. (The screenshot shows a system named StorageGRID Deployment.)
- 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.

### Viewing 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 grid node.

## DC1-S1 (Storage Node)

Overview

Hardware

Network

Storage

Objects

ILM

Events

Tasks

### Node Information

Name	DC1-S1
Type	Storage Node
ID	5bf57bd4-a68d-467e-b866-bfe09a5c6b96
Connection State	 Connected
Software Version	11.4.0 (build 20200328.0051.269ac98)
IP Addresses	10.96.101.111 <a href="#">Show more</a> 

### Alerts



No active alerts

The overview information for a node includes the following:

- **Name:** The hostname assigned to the node and displayed in the Grid Manager.
- **Type:** The type of node — Admin Node, Storage Node, Gateway Node, or Archive 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.
  - **Not connected - Unknown** : The node is not connected to the grid for an unknown reason. For example, the network connection between nodes has been lost or the power 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.
  - **Not connected - 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.
- **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 Master or the Backup.

## DC1-ADM1 (Admin Node)

Overview Hardware Network Storage Load Balancer Events Tasks

### Node Information

Name	DC1-ADM1
Type	Admin Node
ID	711b7b9b-8d24-4d9f-877a-be3fa3ac27e8
Connection State	 Connected
Software Version	11.4.0 (build 20200515.2346.8edcbbf)
HA Groups	Fabric Pools, Master
IP Addresses	192.168.2.208, 10.224.2.208, 47.47.2.208, 47.47.4.219 <a href="#">Show more</a> 

- **IP Addresses:** The node's IP addresses. Click **Show more** to view the node's IPv4 and IPv6 addresses and interface mappings:
  - eth0: Grid Network
  - eth1: Admin Network
  - eth2: Client Network

## Alerts

The Alerts section of the Overview tab lists any alerts currently affecting this node that have not been silenced. Click the alert name to view additional details and recommended actions.

### Alerts

Name	Severity 	Time triggered	Current values
<b>Low installed node memory</b> The amount of installed memory on a node is low.	 Critical	18 hours ago	Total RAM size: 8.37 GB

## Related information

[Monitoring node connection states](#)

[Viewing current alerts](#)

[Viewing a specific alert](#)

## Viewing the Hardware tab


The Hardware tab displays CPU utilization and memory usage for each node, and additional hardware information about appliances.

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, hover 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.

### Related information

[Viewing information about appliance Storage Nodes](#)

[Viewing information about appliance Admin Nodes and Gateway Nodes](#)

### Viewing 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-S1-226 (Storage Node)

Overview Hardware **Network** Storage Objects ILM Events

1 hour 1 day 1 week 1 month 1 year Custom



## Network Interfaces

Name	Hardware Address	Speed	Duplex	Auto Negotiate	Link Status
eth0	00:50:56:A8:2A:75	10 Gigabit	Full	Off	Up

## Network Communication

### Receive

Interface	Data	Packets	Errors	Dropped	Frame Overruns	Frames
eth0	738.858 GB	904,587,345	0	14,340	0	0

### Transmit

Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	677.555 GB	465,715,998	0	0	0	0

## Related information

[Monitoring network connections and performance](#)

### Viewing 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.



The total values for a site or the grid do not include nodes that not have reported metrics for at least five minutes, such as offline nodes.

DC1-SN1-99-88 (Storage Node)



### Disk Devices, Volumes, and Object Store 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	0.03%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.85%	0 bytes/s	58 KB/s
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	81 bytes/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	14.90 GB	Unknown
/var/local	cvloc	Online	85.86 GB	84.10 GB	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/2	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	96.45 GB	250.90 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

[Monitoring storage capacity for the entire grid](#)

[Monitoring storage capacity for each Storage Node](#)

[Monitoring object metadata capacity for each Storage Node](#)

## Viewing the Events tab

The Events tab displays a count of any system error or fault events for a node, including errors such as network errors.

The Events tab is shown for all nodes.

If you experience issues with a particular node, you can use the Events tab to learn more about the issue. Technical support can also use the information on the Events tab to help with troubleshooting.

**Events** 

Last Event                      No Events

Description	Count	
Abnormal Software Events	0	
Account Service Events	0	
Cassandra Heap Out Of Memory Errors	0	
Cassandra unhandled exceptions	0	
Chunk Service Events	0	
Custom Events	0	
Data-Mover Service Events	0	
File System Errors	0	
Forced Termination Events	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	
Node Errors	0	
Out Of Memory Errors	0	
Replicated State Machine Service Events	0	
SCSI Errors	0	
Stat Service Events	0	
Storage Hardware Events	0	
System Time Events	0	

[Reset event counts](#) 

You can perform these tasks from the Events tab:

- Use the information shown for the **Last Event** field at the top of the table to determine which event occurred most recently.
- Click the chart icon  for a specific event to see when that event occurred over time.

- Reset event counts to zero after resolving any issues.

## Related information

[Monitoring events](#)

[Displaying charts and graphs](#)

[Resetting event counts](#)

## Using the Task tab to reboot a grid node

The Task tab allows you to reboot the selected node. The Task tab is shown for all nodes.

## What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.

## About this task

You can use the Task tab to reboot a node. For appliance nodes, you can also use the Task tab to place the appliance into maintenance mode.

Overview Hardware Network Storage Objects ILM Events **Tasks**

### Reboot

Shuts down and restarts the node.

Reboot

### Maintenance Mode

Places the appliance's compute controller into maintenance mode.

Maintenance Mode

- Rebooting a grid node from the Task tab issues the reboot command on the target node. When you reboot a node, the node shuts down and restarts. All services are restarted automatically.

If you plan to reboot a Storage Node, note the following:

- If an ILM rule specifies an ingest behavior of Dual commit or the rule specifies Balanced and it is not possible to immediately create all required copies, StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site and evaluates ILM later. If you want to reboot two or more Storage Nodes on a given site, you might not be able to access these objects for the duration of the reboot.
- 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.
- You might need to put a StorageGRID appliance into maintenance mode to perform certain procedures, such as changing the link configuration or replacing a storage controller. For instructions, see the hardware

installation and maintenance instructions for the appliance.



Putting an appliance into maintenance mode might make the appliance unavailable for remote access.

## Steps

1. Select **Nodes**.
2. Select the grid node you want to reboot.
3. Select the **Tasks** tab.

### DC3-S3 (Storage Node)

Overview

Hardware

Network

Storage

Objects

ILM

Events

Tasks

## Reboot

Reboot shuts down and restarts the node.

Reboot

4. Click **Reboot**.

A confirmation dialog box appears.

### ⚠ Reboot Node DC3-S3

Reboot shuts down and restarts a node, based on where the node is installed:

- 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 are ready to reboot this node, enter the provisioning passphrase and click OK.

Provisioning Passphrase

Cancel

OK



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 click **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 icon (Administratively Down) appears on the left side of the Nodes page. When all services have started again, the icon changes back to its original color.

**Related information**

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[SG100 & SG1000 services appliances](#)



**Viewing the Objects tab**

The Objects tab provides information about S3 and Swift 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.

Overview Hardware Network Storage **Objects** ILM Events Tasks

1 hour 1 day 1 week 1 month Custom



**Object Counts**

Total Objects	0	
Lost Objects	0	
S3 Buckets and Swift Containers	0	

**Queries**

Average Latency	5.74 milliseconds	
Queries - Successful	12,403	
Queries - Failed (timed-out)	0	
Queries - Failed (consistency level unmet)	0	

**Verification**

Status	No Errors	
Rate Setting	Adaptive	
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	

**Related information**

[Use S3](#)

[Use Swift](#)

## Viewing 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.

### DC1-S1 (Storage Node)

**Evaluation**

<b>Awaiting - All</b>	0 objects	
<b>Awaiting - Client</b>	0 objects	
<b>Evaluation Rate</b>	0.00 objects / second	
<b>Scan Rate</b>	0.00 objects / second	

**Erasure Coding Verification**

<b>Status</b>	Idle	
<b>Next Scheduled</b>	2018-05-23 10:44:47 MDT	
<b>Fragments Verified</b>	0	
<b>Data Verified</b>	0 bytes	
<b>Corrupt Copies</b>	0	
<b>Corrupt Fragments</b>	0	
<b>Missing Fragments</b>	0	

## Related information

[Monitoring information lifecycle management](#)

[Administer StorageGRID](#)

## Viewing 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.”

DC1-SG1000-ADM (Admin Node)



### Load Balancer 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.

### Load Balancer 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

[Monitoring load balancing operations](#)

[Administer StorageGRID](#)

### Viewing 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.

## Data Center 1

Network Storage Objects ILM Platform Services

1 hour 1 day 1 week 1 month 1 year Custom



For more information about S3 platform services, including troubleshooting details, see the instructions for administering StorageGRID.

### Related information

[Administer StorageGRID](#)

### Viewing 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 table on the Overview tab displays the node's ID and name, the node type, the software version installed, and the IP addresses associated with the node. The Interface column contains the name of the interface, 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, which can be bonded or aliased and connected to the StorageGRID Admin Network (eth1).

Node Information 	
Name	SGA-lab11
Type	Storage Node
ID	0b583829-6659-4c6e-b2d0-31461d22ba67
Connection State	 Connected
Software Version	11.4.0 (build 20200527.0043.61839a2)
IP Addresses	192.168.4.138, 10.224.4.138, 169.254.0.1 <a href="#">Show less</a> 
Interface	IP Address
eth0	192.168.4.138
eth0	fd20:331:331:0:2a0:98ff:fea1:831d
eth0	fe80::2a0:98ff:fea1:831d
eth1	10.224.4.138
eth1	fd20:327:327:0:280:e5ff:fe43:a99c
eth1	fd20:8b1e:b255:8154:280:e5ff:fe43:a99c
eth1	fe80::280:e5ff:fe43:a99c
hic2	192.168.4.138
hic4	192.168.4.138
mtc1	10.224.4.138
mtc2	169.254.0.1

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-NetApp-SGA-000-012	
Storage Controller A Management IP	10.224.1.79	
Storage Controller B Management IP	10.224.1.80	
Storage Controller WWID	6d039ea000016fc7000000005fac58f4	
Storage Appliance Chassis Serial Number	721924500062	
Storage Controller Firmware Version	08.70.00.02	
Storage Hardware	Needs Attention	
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	DDP	
Storage Connectivity	Nominal	
Overall Power Supply	Nominal	
Compute Controller BMC IP	10.224.0.13	
Compute Controller Serial Number	721917500067	
Compute Hardware	Nominal	
Compute Controller CPU Temperature	Nominal	
Compute Controller Chassis Temperature	Nominal	

## Storage Shelves

Shelf Chassis Serial Number	Shelf ID	Shelf Status	IOM Status	Power Supply Status	Drawer Status	Fan Status	Drive Slots	Data Drives	Data Drive Size	Cache Drives	Cache Drive Size	Configuration Status
721924500062	99	Nominal 	N/A	Nominal	Nominal	Nominal	60	58	4.00 TB	2	800.17 GB	Configured (in use)

Field in the Appliance table	Description
Appliance Model	The model number for this StorageGRID appliance shown in SANtricity software.
Storage Controller Name	The name for this StorageGRID appliance shown in SANtricity software.
Storage Controller A Management IP	IP address for management port 1 on storage controller A. You use this IP to access SANtricity software 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 software to troubleshoot storage issues.  Some appliance models do not have a storage controller B.
Storage Controller WWID	The worldwide identifier of the storage controller shown in SANtricity software.

<b>Field in the Appliance table</b>	<b>Description</b>
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 Hardware	<p>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.</p> <p>If the status is “needs attention,” first check the storage controller using SANtricity software. Then, ensure that no other alarms exist that apply to the compute controller.</p>
Storage Controller Failed Drive Count	The number of drives that are not optimal.
Storage Controller A	The status of storage controller A.
Storage Controller B	The status of storage controller B. Some appliance models do not 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 disk drive) or SSD (solid state drive).
Storage Data Drive Size	The total capacity including all data drives in the appliance.
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.

Field in the Appliance table	Description
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 do not 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 do not 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. <ul style="list-style-type: none"> <li>• 99: Storage controller shelf</li> <li>• 0: First expansion shelf</li> <li>• 1: Second expansion shelf</li> </ul> <p><b>Note:</b> Expansion shelves apply to the SG6060 only.</p>
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.


Column in the Storage Shelves table	Description
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 Negotiate	Link Status
eth0	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
eth1	D8:C4:97:2A:E4:9E	Gigabit	Full	Off	Up
eth2	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
hic1	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic2	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic3	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic4	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
mtc1	D8:C4:97:2A:E4:9E	Gigabit	Full	On	Up
mtc2	D8:C4:97:2A:E4:9F	Gigabit	Full	On	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 the installation and maintenance instructions for your appliance 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	3.250 TB	5,610,578,144	0	8,327	0	0
eth1	1.205 GB	9,828,095	0	32,049	0	0
eth2	849.829 GB	186,349,407	0	10,269	0	0
hic1	114.864 GB	303,443,393	0	0	0	0
hic2	2.315 TB	5,351,180,956	0	305	0	0
hic3	1.690 TB	1,793,580,230	0	0	0	0
hic4	194.283 GB	331,640,075	0	0	0	0
mtc1	1.205 GB	9,828,096	0	0	0	0
mtc2	1.168 GB	9,564,173	0	32,050	0	0

### Transmit

Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	5.759 TB	5,789,638,626	0	0	0	0
eth1	4.563 MB	41,520	0	0	0	0
eth2	855.404 GB	139,975,194	0	0	0	0
hic1	289.248 GB	326,321,151	5	0	0	5
hic2	1.636 TB	2,640,416,419	18	0	0	18
hic3	3.219 TB	4,571,516,003	33	0	0	33
hic4	1.687 TB	1,658,180,262	22	0	0	22
mtc1	4.563 MB	41,520	0	0	0	0
mtc2	49.678 KB	609	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 software (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	0.03%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.85%	0 bytes/s	58 KB/s
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	81 bytes/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	14.90 GB	Unknown
/var/local	cvloc	Online	85.86 GB	84.10 GB	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB	Enabled
/var/local/rangedb/2	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	96.45 GB	250.90 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

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

### Viewing the SANtricity System Manager tab

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.

The SANtricity System Manager tab is shown for storage appliance nodes.

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 the instructions in [administeringStorageGRID](#).

## Administer StorageGRID

To access SANtricity System Manager through Grid Manager, you must have the Storage Appliance Administrator permission or Root Access permission.



You must have SANtricity firmware 8.70 or higher to access SANtricity System Manager using the Grid Manager.



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 do not apply to monitoring your StorageGRID appliance. To avoid issues, always follow the hardware installation and maintenance instructions for your appliance.

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, hover your cursor over each graph.

For more details on viewing the information accessible from the SANtricity System Manager tab, see the information in the [NetApp E-Series Systems Documentation Center](#)

### Viewing 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 for 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 table on the Overview tab displays the node's ID and name, the node type, the software version installed, and the IP addresses associated with the node. The Interface column contains the name of the interface, 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, which can be bonded or aliased and connected to the StorageGRID Admin Network (eth1).


## Node Information

<b>ID</b>	46702fe0-2bca-4097-8f61-f3fe6b22ed75
<b>Name</b>	GW-SG1000-003-076
<b>Type</b>	Gateway Node
<b>Software Version</b>	11.3.0 (build 20190708.2304.71ba19a)
<b>IP Addresses</b>	169.254.0.1, 172.16.3.76, 10.224.3.76, 47.47.3.76 <a href="#">Show less</a> 

Interface	IP Address
adllb	fe80::c020:17ff:fe59:1cf3
adlli	169.254.0.1
adlli	fd20:327:327:0:408f:84ff:fe80:a9
adlli	fd20:8b1e:b255:8154:408f:84ff:fe80:a9
adlli	fe80::408f:84ff:fe80:a9
eth0	172.16.3.76
eth0	fd20:328:328:0:9a03:9bff:fe98:a272
eth0	fe80::9a03:9bff:fe98:a272
eth1	10.224.3.76
eth1	fd20:327:327:0:b6a9:fcff:fe08:4e49
eth1	fd20:8b1e:b255:8154:b6a9:fcff:fe08:4e49
eth1	fe80::b6a9:fcff:fe08:4e49
eth2	47.47.3.76
eth2	fd20:332:332:0:9a03:9bff:fe98:a272
eth2	fe80::9a03:9bff:fe98:a272
hic1	47.47.3.76
hic2	47.47.3.76
hic3	47.47.3.76
hic4	47.47.3.76
mtc1	10.224.3.76
mtc2	10.224.3.76

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	SG1000	
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	10.224.3.95	
Compute Controller Serial Number	721911500171	
Compute Hardware	Nominal	
Compute Controller CPU Temperature	Nominal	
Compute Controller Chassis Temperature	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 are not optimal.
Storage Data Drive Type	The type of drives in the appliance, such as HDD (hard disk drive) or SSD (solid state drive).
Storage Data Drive Size	The total capacity including all data drives in the appliance.


Field in the Appliance table	Description
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 do not 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 Negotiate	Link Status
adllb	C2:20:17:59:1C:F3	10 Gigabit	Full	Off	Up
adlli	42:8F:84:80:00:A9	10 Gigabit	Full	Off	Up
eth0	98:03:9B:98:A2:72	400 Gigabit	Full	Off	Up
eth1	B4:A9:FC:08:4E:49	10 Gigabit	Full	Off	Up
eth2	98:03:9B:98:A2:72	400 Gigabit	Full	Off	Up
hic1	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
hic2	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
hic3	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
hic4	98:03:9B:98:A2:72	100 Gigabit	Full	On	Up
mtc1	B4:A9:FC:08:4E:49	Gigabit	Full	On	Up
mtc2	B4:A9:FC:08:4E:49	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	3.250 TB	5,610,578,144	0	8,327	0	0
eth1	1.205 GB	9,828,095	0	32,049	0	0
eth2	849.829 GB	186,349,407	0	10,269	0	0
hic1	114.864 GB	303,443,393	0	0	0	0
hic2	2.315 TB	5,351,180,956	0	305	0	0
hic3	1.690 TB	1,793,580,230	0	0	0	0
hic4	194.283 GB	331,640,075	0	0	0	0
mtc1	1.205 GB	9,828,096	0	0	0	0
mtc2	1.168 GB	9,564,173	0	32,050	0	0

### Transmit

Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	5.759 TB	5,789,638,626	0	0	0	0
eth1	4.563 MB	41,520	0	0	0	0
eth2	855.404 GB	139,975,194	0	0	0	0
hic1	289.248 GB	326,321,151	5	0	0	5
hic2	1.636 TB	2,640,416,419	18	0	0	18
hic3	3.219 TB	4,571,516,003	33	0	0	33
hic4	1.687 TB	1,658,180,262	22	0	0	22
mtc1	4.563 MB	41,520	0	0	0	0
mtc2	49.678 KB	609	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](#)[Events](#)[Tasks](#)**Disk Devices**

Name	World Wide Name	I/O Load	Read Rate	Write Rate
croot(253:2,dm-2)	N/A	0.00% 	0 bytes/s 	8 KB/s 
cvloc(253:3,dm-3)	N/A	0.01% 	0 bytes/s 	405 KB/s 

**Volumes**

Mount Point	Device	Status	Size	Available	Write Cache Status
/	croot	Online	21.00 GB	13.09 GB 	Unknown 
/var/local	cvloc	Online	903.78 GB	894.55 GB 	Unknown 

**Related information**[SG100 & SG1000 services appliances](#)**Information you should monitor regularly**

StorageGRID is a fault-tolerant, distributed storage system that is designed to continue operating even when errors occur, or when nodes or sites are unavailable. You must proactively monitor system health, workloads, and usage statistics so that you can take action to address potential issues before they affect the grid's efficiency or availability.

A busy system generates large amounts of information. This section provides guidance about the most important information to monitor on an ongoing basis. This section contains the following sub-sections:

- [Monitoring system health](#)
- [Monitoring storage capacity](#)
- [Monitoring information lifecycle management](#)
- [Monitoring performance, networking, and system resources](#)
- [Monitoring tenant activity](#)
- [Monitoring archival capacity](#)
- [Monitoring load balancing operations](#)
- [Applying hotfixes or upgrading software if necessary](#)

What to monitor	Frequency
The system health data shown on the Grid Manager Dashboard Note if anything has changed from the previous day.	Daily
Rate at which Storage Node object and metadata capacity is being consumed	Weekly
Information lifecycle management operations	Weekly
Performance, networking, and system resources: <ul style="list-style-type: none"> <li>• Query latency</li> <li>• Connectivity and networking</li> <li>• Node-level resources</li> </ul>	Weekly
Tenant activity	Weekly
Capacity of the external archival storage system	Weekly
Load balancing operations	After the initial configuration and after any configuration changes
Availability of software hotfixes and software upgrades	Monthly

## Monitoring system health

You should monitor the overall health of your StorageGRID system on a daily basis.

The StorageGRID system is fault tolerant and can continue to operate even when parts of the grid are unavailable. The first sign of a potential issue with your StorageGRID system is likely to be an alert or an alarm (legacy system) and not necessarily an issue with system operations. Paying attention to system health can help you detect minor issues before they affect operations or grid efficiency.

The Health panel on the Grid Manager Dashboard provides a summary of issues that might be affecting your system. You should investigate any issues that are shown on the Dashboard.



To be notified of alerts as soon as they are triggered, you can set up email notifications for alerts or configure SNMP traps.

1. Sign in to the Grid Manager to view the Dashboard.
2. Review the information in the Health panel.



When issues exist, links appear that allow you to view additional details:

Link	Indicates
<a href="#">Grid details</a>	Appears if any nodes are disconnected (connection state Unknown or Administratively Down). Click the link, or click the blue or gray icon to determine which node or nodes are affected.
<a href="#">Current alerts</a>	Appears if any alerts are currently active. Click the link, or click <b>Critical</b> , <b>Major</b> , or <b>Minor</b> to see the details on the <b>Alerts &gt; Current</b> page.
<a href="#">Recently resolved alerts</a>	Appears if any alerts triggered in the past week are now resolved. Click the link to see the details on the <b>Alerts &gt; Resolved</b> page.
<a href="#">Legacy alarms</a>	Appears if any alarms (legacy system) are currently active. Click the link to see the details on the <b>Support &gt; Alarms (legacy) &gt; Current Alarms</b> page.  <b>Note:</b> While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.
<a href="#">License</a>	Appears if there is an issue with the software license for this StorageGRID system. Click the link to see the details on the <b>Maintenance &gt; System &gt; License</b> page.

### Related information

[Administer StorageGRID](#)

[Setting up email notifications for alerts](#)

[Using SNMP monitoring](#)

## Monitoring node connection states

If one or more nodes are disconnected from the grid, critical StorageGRID operations might be affected. You must monitor node connection states and address any issues promptly.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

### About this task

Nodes can have one of three connection states:

- Not connected - Unknown** : The node is not connected to the grid for an unknown reason. For example, the network connection between nodes has been lost or the power 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.


- Not connected - 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.

### Steps

- If a blue or gray icon appears on the Health panel of the Dashboard, click the icon or click **Grid details**. (The blue or gray icons and the **Grid details** link appear only if at least one node is disconnected from the grid.)

The Overview page for the first blue node in the node tree appears. If there are no blue nodes, the Overview page for the first gray node in the tree appears.

In the example, the Storage Node named DC1-S3 has a blue icon. The **Connection State** on the Node Information panel is **Unknown**, and the **Unable to communicate with node** alert is active. The alert indicates that one or more services are unresponsive, or the node cannot be reached.



The screenshot shows the StorageGRID interface. On the left, a tree view shows the deployment structure: StorageGRID Deployment > Data Center 1 > DC1-ADM1, DC1-ADM2, DC1-S1, DC1-S2, and DC1-S3. DC1-S3 is highlighted with a blue icon. The main panel shows the details for DC1-S3 (Storage Node). The 'Node Information' section includes: Name (DC1-S3), Type (Storage Node), ID (9915f7e1-6c53-45ee-bcde-03753db43aba), Connection State (Unknown), Software Version (11.4.0 (build 20200421.1742.8bf07da)), and IP Addresses (10.96.104.171). The 'Alerts' section shows an active alert: 'Unable to communicate with node' with a severity of 'Major' and a time triggered of '12 minutes ago'. The current values for the alert are: 'Unresponsive acct, adc, chunk, dds, dmv, dynip, idnt, jaegeragent, jmx, ldr, miscd, node, services: rsm, ssm, storagegrid'.

2. If a node has a blue icon, follow these steps:
  - a. Select each alert in the table, and follow the recommended actions.

For example, you might need to restart a service that has stopped or restart the host for the node.

- b. If you are unable to bring the node back online, contact technical support.
3. If a node has a gray icon, follow these steps:

Gray nodes are expected during maintenance procedures and might be associated with one or more alerts. Based on the underlying issue, these “administratively down” nodes often go back online with no intervention.

- a. Review the Alerts section, and determine if any alerts are affecting this node.
- b. If one or more alerts are active, select each alert in the table, and follow the recommended actions.
- c. If you are unable to bring the node back online, contact technical support.

## Related information

[Alerts reference](#)

[Maintain & recover](#)

## Viewing 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. An email notification might also be sent, unless the alert has been silenced.

## What you’ll need

- You must be signed in to the Grid Manager using a supported browser.

## Steps

1. If one or more alerts are active, do either of the following:
  - From the Health panel on the Dashboard, click the alert icon or click **Current alerts**. (An alert icon and the **Current alerts** link appear only if at least one alert is currently active.)
  - Select **Alerts > Current**.

The Current Alerts page appears. It lists all alerts currently affecting your StorageGRID system.

Current Alerts [Learn more](#)

View the current alerts affecting your StorageGRID system.

Name	Severity	Time triggered	Site / Node	Status	Current values
▼ <b>Unable to communicate with node</b> One or more services are unresponsive or cannot be reached by the metrics collection job.	2 Major	9 minutes ago <i>(newest)</i> 19 minutes ago <i>(oldest)</i>		2 Active	
<b>Low root disk capacity</b> The space available on the root disk is low.	Minor	25 minutes ago	Data Center 1 / DC1-S1-99-51	Active	Disk space available: 2.00 GB Total disk space: 21.00 GB
<b>Expiration of server certificate for Storage API Endpoints</b> The server certificate used for the storage API endpoints is about to expire.	Major	31 minutes ago	Data Center 1 / DC1-ADM1-99-49	Active	Days remaining: 14
<b>Expiration of server certificate for Management Interface</b> The server certificate used for the management interface is about to expire.	Minor	31 minutes ago	Data Center 1 / DC1-ADM1-99-49	Active	Days remaining: 30
▼ <b>Low installed node memory</b> The amount of installed memory on a node is low.	8 Critical	a day ago <i>(newest)</i> a day ago <i>(oldest)</i>		8 Active	

By default, 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 are not 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. Review the information in the table.

Column header	Description
Name	The name of the alert and its description.
Severity	<p>The severity of the alert. If multiple alerts are grouped, the title row shows how many instances of that alert are occurring at each severity.</p> <ul style="list-style-type: none"> <li>• <b>Critical</b> : 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.</li> <li>• <b>Major</b> : 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.</li> <li>• <b>Minor</b> : 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 do not clear on their own to ensure they do not result in a more serious problem.</li> </ul>
Time triggered	How long ago the alert was triggered. 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> ).
Site/Node	The name of the site and node where the alert is occurring. If multiple alerts are grouped, the site and node names are not shown in the title row.

Column header	Description
Status	Whether the alert is active or has been silenced. If multiple alerts are grouped and <b>All alerts</b> 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.
Current values	The current 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 <b>Low object data storage</b> alert include the percentage of disk space used, the total amount of disk space, and the amount of disk space used.  <b>Note:</b> If multiple alerts are grouped, current values are not shown in the title row.

3. To expand and collapse groups of alerts:

- To show the individual alerts in a group, click the down caret ▼ in the heading, or click the group's name.
- To hide the individual alerts in a group, click the up caret ▲ in the heading, or click the group's name.

Name	Severity	Time triggered	Site / Node	Status	Current values
▲ <b>Low object data storage</b> The disk space available for storing object data is low.	▲ 5 Minor	a day ago (newest) a day ago (oldest)		5 Active	
<b>Low object data storage</b> The disk space available for storing object data is low.	▲ Minor	a day ago	DC2 231-236 / DC2-S2-233	Active	Disk space remaining: 525.17 GB Disk space used: 243.06 KB Disk space used (%): 0.000%
<b>Low object data storage</b> The disk space available for storing object data is low.	▲ Minor	a day ago	DC1 225-230 / DC1-S1-226	Active	Disk space remaining: 525.17 GB Disk space used: 325.65 KB Disk space used (%): 0.000%
<b>Low object data storage</b> The disk space available for storing object data is low.	▲ Minor	a day ago	DC2 231-236 / DC2-S3-234	Active	Disk space remaining: 525.17 GB Disk space used: 381.55 KB Disk space used (%): 0.000%
<b>Low object data storage</b> The disk space available for storing object data is low.	▲ Minor	a day ago	DC1 225-230 / DC1-S2-227	Active	Disk space remaining: 525.17 GB Disk space used: 282.19 KB Disk space used (%): 0.000%
<b>Low object data storage</b> The disk space available for storing object data is low.	▲ Minor	a day ago	DC2 231-236 / DC2-S1-232	Active	Disk space remaining: 525.17 GB Disk space used: 189.24 KB Disk space used (%): 0.000%

4. To display individual alerts instead of groups of alerts, unselect the **Group alerts** check box at the top of the table.



5. To sort alerts or alert groups, click 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 unselected, 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.

6. To filter the alerts by status, use the drop-down menu at the top of the table.



- Select **All alerts** to view all current alerts (both active and silenced alerts).
- Select **Active** to view only the current alerts that are active.
- Select **Silenced** to view only the current alerts that have been silenced.

7. To view details for a specific alert, select the alert from the table.

A dialog box for the alert appears. See the instructions for viewing a specific alert.

### Related information

[Viewing a specific alert](#)

[Silencing alert notifications](#)

### Viewing resolved alerts

You can search and view a history of alerts that have been resolved.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

### Steps

1. To view resolved alerts, do either of the following:
  - From the Health panel on the Dashboard, click **Recently resolved alerts**.

The **Recently resolved alerts** link appears only if one or more alerts were triggered in the past week and are now resolved.

- Select **Alerts > Resolved**.  
The Resolved Alerts page appears. By default, resolved alerts that were triggered in the last week are shown, with the most recently triggered alerts shown first. The alerts on this page were previously shown on the Current Alerts page or in an email notification.

## Resolved Alerts

Search and view alerts that have been resolved.

When triggered ✕      Severity ✕      Alert rule ✕      Node ✕

Last week      Filter by severity      Filter by rule      Filter by node      Search

Name	IT	Severity <span>ⓘ</span>	IT	Time triggered <span>▼</span>	Time resolved <span>IT</span>	Site / Node	IT	Triggered values
<b>Low installed node memory</b> The amount of installed memory on a node is low.		<span>✖</span> Critical		2 days ago	a day ago	Data Center 1 / DC1-S2		Total RAM size: 8.37 GB
<b>Low installed node memory</b> The amount of installed memory on a node is low.		<span>✖</span> Critical		2 days ago	a day ago	Data Center 1 / DC1-S3		Total RAM size: 8.37 GB
<b>Low installed node memory</b> The amount of installed memory on a node is low.		<span>✖</span> Critical		2 days ago	a day ago	Data Center 1 / DC1-S4		Total RAM size: 8.37 GB
<b>Low installed node memory</b> The amount of installed memory on a node is low.		<span>✖</span> Critical		2 days ago	a day ago	Data Center 1 / DC1-ADM1		Total RAM size: 8.37 GB
<b>Low installed node memory</b> The amount of installed memory on a node is low.		<span>✖</span> Critical		2 days ago	a day ago	Data Center 1 / DC1-ADM2		Total RAM size: 8.37 GB
<b>Low installed node memory</b> The amount of installed memory on a node is low.		<span>✖</span> Critical		2 days ago	a day ago	Data Center 1 / DC1-S1		Total RAM size: 8.37 GB

## 2. Review the information in the table.

Column header	Description
Name	The name of the alert and its description.
Severity	<p>The severity of the alert.</p> <ul style="list-style-type: none"> <li>• <b>Critical</b> <span>✖</span>: 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.</li> <li>• <b>Major</b> <span>!</span>: 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.</li> <li>• <b>Minor</b> <span>!</span>: 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 do not clear on their own to ensure they do not result in a more serious problem.</li> </ul>
Time triggered	How long ago the alert was triggered.
Time resolved	How long ago the alert was resolved.
Site/Node	The name of the site and node where the alert occurred.

Column header	Description
Triggered values	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 <b>Low object data storage</b> alert include the percentage of disk space used, the total amount of disk space, and the amount of disk space used.

3. To sort the entire list of resolved alerts, click the up/down arrows  in each column header.

For example, you might want to sort resolved alerts by **Site/Node** to see the alerts that affected a specific node.

4. Optionally, filter the list of resolved alerts by using the drop-down menus at the top of the table.

- a. Select a time period from the **When triggered** drop-down menu to show resolved alerts based on how long ago they were triggered.

You can search for alerts that were triggered within the following time periods:

- Last hour
- Last day
- Last week (default view)
- Last month
- Any time period
- Custom (allows you to specify the start date and the end date for the time period)

- b. Select one or more severities from the **Severity** drop-down menu to filter on resolved alerts of a specific severity.
- c. 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.
- d. Select one or more nodes from the **Node** drop-down menu to filter on resolved alerts related to a specific node.
- e. Click **Search**.

5. To view details for a specific resolved alert, select the alert from the table.

A dialog box for the alert appears. See the instructions for viewing a specific alert.

## Related information

[Viewing a specific alert](#)

### Viewing a specific alert

You can view detailed information about an alert that is currently affecting your StorageGRID system or an alert that has been resolved. The details include recommended corrective actions, the time the alert was triggered, and the current value of the metrics related to this alert. Optionally, you can silence a current alert or update the

alert rule.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

### Steps

1. Do one of the following, based on whether you want to view a current or resolved alert:

Column header	Description
<b>Current alert</b>	<ul style="list-style-type: none"><li>• From the Health panel on the Dashboard, click the <b>Current alerts</b> link. This link appears only if at least one alert is currently active. This link is hidden if there are no current alerts or if all current alerts have been silenced.</li><li>• Select <b>Alerts &gt; Current</b>.</li><li>• From the <b>Nodes</b> page, select the <b>Overview</b> tab for a node that has an alert icon. Then, in the Alerts section, click the alert name.</li></ul>
<b>Resolved alert</b>	<ul style="list-style-type: none"><li>• From the Health panel on the Dashboard, click the <b>Recently resolved alerts</b> link. (This link appears only if one or more alerts were triggered in the past week and are now resolved. This link is hidden if no alerts were triggered and resolved in the last week.)</li><li>• Select <b>Alerts &gt; Resolved</b>.</li></ul>

2. As required, expand a group of alerts and then select the alert you want to view.



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

^ Low installed node memory The amount of installed memory on a node is low.	8 Critical	a day ago (newest) a day ago (oldest)		8 Active	
<u>Low installed node memory</u> The amount of installed memory on a node is low.	1 Critical	a day ago	Data Center 2 / DC2-S1-99-56	Active	Total RAM size: 8.38 GB

A dialog box appears and provides details for the selected alert.

## Low installed node memory

The amount of installed memory on a node is low.

### Recommended actions

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](#)

### Time triggered

2019-07-15 17:07:41 MDT (2019-07-15 23:07:41 UTC)

Status

Active ([silence this alert](#) )

Site / Node

Data Center 2 / DC2-S1-99-56

Severity

 Critical

Total RAM size

8.38 GB

Condition

[View conditions](#) | [Edit rule](#) 

Close

### 3. Review the alert details.

Information	Description
<i>title</i>	The name of the alert.
<i>first paragraph</i>	The description of the alert.
Recommended actions	The recommended actions for this alert.
Time triggered	The date and time the alert was triggered in your local time and in UTC.
Time resolved	For resolved alerts only, the date and time the alert was resolved in your local time and in UTC.
Status	The status of the alert: Active, Silenced, or Resolved.
Site/Node	The name of the site and node affected by the alert.

Information	Description
Severity	<p>The severity of the alert.</p> <ul style="list-style-type: none"> <li>• <b>Critical</b> : 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.</li> <li>• <b>Major</b> : 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.</li> <li>• <b>Minor</b> : 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 do not clear on their own to ensure they do not result in a more serious problem.</li> </ul>
<i>data values</i>	<p>The current value of the metric for this alert. For some alerts, additional values are shown to help you understand and investigate the alert. For example, the values shown for a <b>Low metadata storage</b> alert include the percent of disk space used, the total amount of disk space, and the amount of disk space used.</p>

4. Optionally, click **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.

5. To view the current conditions for the alert rule:

a. From the alert details, click **View conditions**.

A pop-up appears, listing the Prometheus expression for each defined severity.

Total RAM size  
8.38 GB

Condition  
[View conditions](#) | [Edit rule](#)

**Low installed node memory**

Major `node_memory_MemTotal_bytes < 24000000000`

Critical `node_memory_MemTotal_bytes < 12000000000`

b. To close the pop-up, click anywhere outside of the pop-up.

6. Optionally, click **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.

7. To close the alert details, click **Close**.

## Related information

[Silencing alert notifications](#)

[Editing an alert rule](#)

## Viewing legacy alarms

Alarms (legacy system) are triggered when system attributes reach alarm threshold values. You can view the currently active alarms from the Dashboard or the Current Alarms page.

## What you'll need

- You must be signed in to the Grid Manager using a supported browser.

## About this task

If one or more of the legacy alarms are currently active, the Health panel on the Dashboard includes a **Legacy alarms** link. The number in parentheses indicates how many alarms are currently active.

Health ?

Administratively Down 1

Critical 5

License Status 1

Grid details Current alerts (5) Recently resolved alerts (1) **Legacy alarms (5)** License

The **Legacy alarms** count on the Dashboard is incremented whenever a legacy alarm is triggered. This count is incremented even if you have disabled alarm email notifications. You can typically ignore this number (since

alerts provide a better view of the system), or you can view the alarms that are currently active.



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

### Steps

- To view the legacy alarms that are currently active, do one of the following:
  - From the Health panel on the Dashboard, click **Legacy alarms**. This link appears only if at least one alarm is currently active.
  - Select **Support > Alarms (legacy) > Current Alarms**. The Current Alarms page appears.

The alarm system is the legacy system. The alert system offers significant benefits and is easier to use. See [Managing alerts and alarms](#) in the instructions for monitoring and troubleshooting StorageGRID.

## Current Alarms

Last Refreshed: 2020-05-27 09:41:39 MDT

Show Acknowledged Alarms (1 - 1 of 1)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value
Major	ORSU (Outbound Replication Status)	<a href="#">Data Center 1/DC1-ARC1/ARC</a>	Storage Unavailable	2020-05-26 21:47:18 MDT	Storage Unavailable	Storage Unavailable

Show  Records Per Page  Previous < 1 > Next


The alarm icon indicates the severity of each alarm, as follows:

Icon	Color	Alarm severity	Meaning
	Yellow	Notice	The node is connected to the grid, but an unusual condition exists that does not affect normal operations.
	Light Orange	Minor	The node is connected to the grid, but an abnormal condition exists that could affect operation in the future. You should investigate to prevent escalation.
	Dark Orange	Major	The node is connected to the grid, but an abnormal condition exists that currently affects operation. This requires prompt attention to prevent escalation.

Icon	Color	Alarm severity	Meaning
	Red	Critical	The node is connected to the grid, but an abnormal condition exists that has stopped normal operations. You should address the issue immediately.

- To learn about the attribute that caused the alarm to be triggered, right click the attribute name in the table.
- To view additional details about an alarm, click the service name in the table.

The Alarms tab for the selected service appears (**Support > Tools > Grid Topology > Grid Node > Service > Alarms**).



Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
 Major	ORSU (Outbound Replication Status)	Storage Unavailable	2019-05-23 21:40:08 MDT	Storage Unavailable	Storage Unavailable		<input type="checkbox"/>

- If you want to clear the count of current alarms, you can optionally do the following:
  - Acknowledge the alarm. An acknowledged alarm is no longer included in the count of legacy alarms unless it is triggered at the next severity level or it is resolved and occurs again.
  - Disable a particular Default alarm or Global Custom alarm for the entire system to prevent it from being triggered again.

### Related information

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

[Acknowledging current alarms \(legacy system\)](#)

[Disabling alarms \(legacy system\)](#)

### Monitoring storage capacity

You must monitor the total usable space available on Storage Nodes 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.

## Related information

[Viewing the Storage tab](#)

### Monitoring storage capacity for the entire grid

You must 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.

### What you'll need

You must be signed in to the Grid Manager using a supported browser.

### About this task


The Dashboard in the Grid Manager 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**.
  - b. In the Available Storage panel, note the overall summary of free and used storage capacity.



The summary does not include archival media.




- c. Place your cursor over the chart's Free or Used capacity sections to see exactly how much space is free or used.



- d. For multi-site grids, review the chart for each data center.
- e. Click the chart icon  for the overall chart or for an individual data center to view a graph showing capacity usage over time.

A graph showing Percentage Storage Capacity Used (%) vs. Time appears.

2. Determine how much storage has been used and how much storage remains available for object data and object metadata.
  - a. Select **Nodes**.
  - b. Select **grid > Storage**.



- c. Hover 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 do not include nodes that not have reported metrics for at least five minutes, such as offline nodes.

3. As directed by technical support, view additional details about the storage capacity for your grid.
  - a. Select **Support > Tools > Grid Topology**.
  - b. Select **grid > Overview > Main**.

The screenshot shows the StorageGRID Overview page. On the left, a 'Grid Topology' sidebar displays 'StorageGRID Deployment' with three data centers: Data Center 1, Data Center 2, and Data Center 3. The main content area has tabs for Overview, Alarms, Reports, and Configuration. The 'Overview' tab is active, showing a 'Main' sub-tab. The page title is 'Overview: Summary - StorageGRID Deployment', updated on 2019-03-01 at 11:50:40 MST. Below the title, there are two sections: 'Storage Capacity' and 'ILM Activity'. Each section contains a table of metrics with expandable icons on the right.

Storage Capacity	
Storage Nodes Installed:	9
Storage Nodes Readable:	9
Storage Nodes Writable:	9
Installed Storage Capacity:	2,898 GB
Used Storage Capacity:	100 GB
Used Storage Capacity for Data:	2.31 MB
Used Storage Capacity for Metadata:	5.82 MB
Usable Storage Capacity:	2,797 GB
Percentage Storage Capacity Used:	3.465 %
Percentage Usable Storage Capacity:	96.535 %

ILM Activity	
Awaiting - All:	0
Awaiting - Client:	0
Scan Rate:	0 Objects/s
Scan Period - Estimated:	0 us
Awaiting - Evaluation Rate:	0 Objects/s
Repairs Attempted:	0

- 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 on planning a storage expansion, see the instructions for expanding StorageGRID.

## Related information

[Expand your grid](#)

### Monitoring storage capacity for each Storage Node


You must monitor the total usable space for each Storage Node to ensure that the node has enough space for new object data.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

### 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.



**Total Usable Space = Usable Space 0 + Usable Space 1 + Usable Space 2**

**Steps**


1. Select **Nodes > Storage Node > Storage**.

The graphs and tables for the node appear.

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

The following values are shown:

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



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	0.03%	0 bytes/s	3 KB/s
cvloc(8:2,sda2)	N/A	0.85%	0 bytes/s	58 KB/s
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	81 bytes/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	14.90 GB 	Unknown
/var/local	cvloc	Online	85.86 GB	84.10 GB 	Unknown
/var/local/rangedb/0	sdc	Online	107.32 GB	107.18 GB 	Enabled
/var/local/rangedb/1	sdd	Online	107.32 GB	107.18 GB 	Enabled
/var/local/rangedb/2	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	96.45 GB 	250.90 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

4. Monitor the values over time to estimate the rate at which usable storage space is being consumed.
5. 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 on planning a storage expansion, see the instructions for expanding StorageGRID.

The **Low object data storage** alert and the legacy Storage Status (SSTS) alarm are triggered when insufficient space remains for storing object data on a Storage Node.

#### Related information

[Administer StorageGRID](#)

[Troubleshooting the Low object data storage alert](#)

[Expand your grid](#)

## Monitoring object metadata capacity for each Storage Node

You must 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

### 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

- Select **Nodes** > **Storage Node** > **Storage**.
- Hover your cursor over the Storage Used - Object Metadata graph to see the values for a specific time.



Value	Description	Prometheus metric
Used (%)	The percentage of the allowed metadata space that has been used on this Storage Node.	<code>storagegrid_storage_utilization_metadata_bytes/storagegrid_storage_utilization_metadata_allowed_bytes</code>

Value	Description	Prometheus metric
Used	The bytes of the allowed metadata space that have been used on this Storage Node.	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 instructions for administering StorageGRID.	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 instructions for administering StorageGRID.	storagegrid_storage_utilization_metadata_reserved_bytes



The total values for a site or the grid do not include nodes that have not reported metrics for at least five minutes, such as offline nodes.

- 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.

## Related information

[Troubleshooting the Low metadata storage alert](#)

[Administer StorageGRID](#)

[Expand your grid](#)

## Monitoring 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 required.

## What you'll need

You must be signed in to the Grid Manager using a supported browser.

### About this task


The StorageGRID system manages objects by applying the active ILM policy. The ILM policy 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 cannot be fulfilled in near real time. You can monitor whether StorageGRID is keeping up with client actions by charting the Awaiting - Client attribute.

To chart this attribute:

1. Sign in to the Grid Manager.
2. From the Dashboard, locate the **Awaiting - Client** entry in the Information Lifecycle Management (ILM) panel.
3. Click the chart icon .

The example chart shows a situation where the number of objects awaiting ILM evaluation temporarily increased in an unsustainable manner, then eventually decreased. Such a trend indicates that ILM was temporarily not fulfilled in near real time.



Temporary spikes in the chart of Awaiting - Client are to be expected. But if the value shown on the chart continues to increase and never declines, the grid requires more resources to operate efficiently: either more Storage Nodes, or, if the ILM policy places objects in remote locations, more network bandwidth.

You can further investigate ILM queues using the **Nodes** page.

### Steps

1. Select **Nodes**.
2. Select **grid name > ILM**.
3. Hover 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.

## Monitoring performance, networking, and system resources

You should monitor performance, networking, and system resources to determine whether StorageGRID can handle its current load and to ensure that client performance does not degrade over time.

### Monitoring query latency

Client actions such as storing, retrieving, or deleting objects create queries to the grid's distributed database of object metadata. You should monitor trends in query latency to ensure that grid resources are adequate for the current load.

### What you'll need

You must be signed in to the Grid Manager using a supported browser.

### About this task

Temporary increases in query latency are normal and can be caused by a sudden increase in ingest requests. Failed queries are also normal and can result from transient network issues or nodes that are temporarily unavailable. However, if the average time to perform a query increases, overall grid performance declines.

If you notice that query latency is increasing over time, you should consider adding additional Storage Nodes in an expansion procedure to satisfy future workloads.


The **High latency for metadata queries** alert is triggered if the average time for queries is too long.

## Steps

1. Select **Nodes** > **Storage Node** > **Objects**.
2. Scroll down to the Queries table and view the value for Average Latency.

Queries		
<b>Average Latency</b>	1.22 milliseconds	
<b>Queries - Successful</b>	1,349,103,223	
<b>Queries - Failed (timed-out)</b>	12022	
<b>Queries - Failed (consistency level unmet)</b>	560925	

3. Click the chart icon  to chart the value over time.



The example chart shows spikes in query latency during normal grid operation.

## Related information

[Expand your grid](#)

### Monitoring network connections and performance

Grid nodes must be able to communicate with one another to permit the grid to operate. The integrity of the network between nodes and sites, and the network bandwidth between sites, are critical to efficient operations.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

- You must have specific access permissions.

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 simply to poor ingest performance and ILM backlogs.


You can use the Grid Manager to monitor connectivity and network performance, so you can address any issues promptly.

Additionally, consider creating network traffic classification policies to provide monitoring and limiting for traffic related to specific tenants, buckets, subnets, or load balancer endpoints. See the instructions for administering StorageGRID.

## Steps


1. Select **Nodes**.

The Nodes page appears. The node icons indicate at a glance which nodes are connected (green checkmark icon) and which nodes are disconnected (blue or gray icons).



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 Negotiate	Link Status
eth0	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
eth1	D8:C4:97:2A:E4:9E	Gigabit	Full	Off	Up
eth2	50:6B:4B:42:D7:11	100 Gigabit	Full	Off	Up
hic1	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic2	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic3	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
hic4	50:6B:4B:42:D7:11	25 Gigabit	Full	Off	Up
mtc1	D8:C4:97:2A:E4:9E	Gigabit	Full	On	Up
mtc2	D8:C4:97:2A:E4:9F	Gigabit	Full	On	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	3.250 TB	5,610,578,144	0	8,327	0	0
eth1	1.205 GB	9,828,095	0	32,049	0	0
eth2	849.829 GB	186,349,407	0	10,269	0	0
hic1	114.864 GB	303,443,393	0	0	0	0
hic2	2.315 TB	5,351,180,956	0	305	0	0
hic3	1.690 TB	1,793,580,230	0	0	0	0
hic4	194.283 GB	331,640,075	0	0	0	0
mtc1	1.205 GB	9,828,096	0	0	0	0
mtc2	1.168 GB	9,564,173	0	32,050	0	0

### Transmit

Interface	Data	Packets	Errors	Dropped	Collisions	Carrier
eth0	5.759 TB	5,789,638,626	0	0	0	0
eth1	4.563 MB	41,520	0	0	0	0
eth2	855.404 GB	139,975,194	0	0	0	0
hic1	289.248 GB	326,321,151	5	0	0	5
hic2	1.636 TB	2,640,416,419	18	0	0	18
hic3	3.219 TB	4,571,516,003	33	0	0	33
hic4	1.687 TB	1,658,180,262	22	0	0	22
mtc1	4.563 MB	41,520	0	0	0	0
mtc2	49.678 KB	609	0	0	0	0

3. Use the metrics associated with your traffic classification policies to monitor network traffic.

a. Select **Configuration > Network Settings > 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.

	Name	Description	ID
<input type="radio"/>	ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574
<input checked="" type="radio"/>	Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665bddc894b

Displaying 2 traffic classification policies.

- b. 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**.
- c. 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.

To create, edit, or delete traffic classification policies, see the instructions for administering StorageGRID.

### Related information

[Viewing the Network tab](#)

[Monitoring node connection states](#)

[Administer StorageGRID](#)

### Monitoring node-level resources

You should monitor individual grid nodes to check their resource utilization levels.

### What you'll need


- You must be signed in to the Grid Manager using a supported browser.

### About this task

If nodes are consistently overloaded, more nodes might be required for efficient operations.

### Steps

1. To view information about hardware utilization of a grid node:
  - a. From the **Nodes** page, select the node.
  - b. Select the **Hardware** tab to display graphs of CPU Utilization and Memory Usage.



- c. 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.
- d. 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

[Viewing information about appliance Storage Nodes](#)

[Viewing information about appliance Admin Nodes and Gateway Nodes](#)

### Monitoring tenant activity

All client activity is associated with a tenant account. You can use the Grid Manager to monitor a tenant’s storage usage or network traffic, or you can use the audit log or Grafana dashboards to gather more detailed information about how tenants are using StorageGRID.

### What you’ll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access or Administrator permission.



#### About this task

The Space used values are estimates. These estimates are affected by the timing of ingests, network connectivity, and node status.

### Steps

1. Select **Tenants** to review the amount of storage used by all tenants.

The Space Used, Quota Utilization, Quota, and Object Count are listed for each tenant. If a quota is not set for a tenant, the Quota Utilization field contains a dash (--) and the Quota field indicates “Unlimited.”

## Tenant Accounts

View information for each tenant account.

**Note:** 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 and select **View Details**.

	Display Name  	Space Used  	Quota Utilization  	Quota  	Object Count  	Sign in 
<input checked="" type="radio"/>	Account01	500.00 KB	0.00%	20.00 GB	100	
<input type="radio"/>	Account02	2.50 MB	0.01%	30.00 GB	500	
<input type="radio"/>	Account03	605.00 MB	4.03%	15.00 GB	31,000	
<input type="radio"/>	Account04	1.00 GB	10.00%	10.00 GB	200,000	
<input type="radio"/>	Account05	0 bytes	—	Unlimited	0	

Show  rows per page

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. Use the search box to search for a tenant account by display name or tenant ID.

You can sign in to a tenant account by selecting the link in the **Sign in** column of the table.

2. 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 a .csv file look like the following example:

Tenant ID	Display Name	Space Used (Bytes)	Quota utilization (%)	Quota (Bytes)	Object Count	Protocol
56243391454153665591	Account01	500000	0	20000000000	100	S3
82457136581801590515	Account02	2500000	0.01	30000000000	500	S3
04489086912300179118	Account03	605000000	4.03	15000000000	31000	S3
26417581662098345719	Account04	1000000000	10	10000000000	200000	S3
78472447501213318575	Account05	0			0	S3

You can open the .csv file in a spreadsheet application or use it in automation.

3. To view details for a specific tenant, including usage charts, select the tenant account from the Tenant Accounts page, and then select **View details**.

The Account Details page appears and shows summary information, a chart that represents the amount of quota used and remaining, and a chart that represents the amount of object data in buckets (S3) or containers (Swift).

## Account Details - Account01


Display Name: Account01 [Sign in](#)  
Tenant ID: 6479 6966 4290 3892 3647  
Protocol [?](#): S3  
Allow Platform Services [?](#): Yes  
Uses Own Identity Source [?](#): No

Quota Utilization [?](#): 25.52%  
Logical Space Used [?](#): 127.58 MB  
Quota [?](#): 500.00 MB  
Bucket Count [?](#): 5  
Object Count [?](#): 30


Overview

Bucket Details

Quota [?](#)



Space Used by Buckets [?](#)



Close

### ◦ Quota

If a quota was set for this tenant, the **Quota** chart shows how much of that quota this tenant has used and how much is still available. If no quota was set, the tenant has an unlimited quota, and an informational message is displayed. 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 Used Space segment to see the number of stored objects and the total bytes used. You can place your cursor over the Free Space segment to see how many bytes of storage quota are available.



Quota utilization 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 utilization is recalculated. Quota utilization calculations can take 10 minutes or longer.



A tenant's quota utilization indicates the total amount of object data the tenant has uploaded to StorageGRID (logical size). The quota utilization 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 to determine if tenants are consuming their quotas. If enabled, this alert is triggered when a tenant has used 90% of its quota. For more information, see the alerts reference.

#### ◦ **Space Used**

The **Space Used by Buckets** (S3) or **Space Used by Containers** (Swift) chart shows the largest buckets for the tenant. Space used is the total amount of object data in the bucket. This value does not represent the storage space required for ILM copies and object metadata.

If the tenant has more than nine buckets or containers, they are combined into a segment called Other. Some chart segments might be too small to include a label. You can place your cursor over any of the segments to see the label and obtain more information, including the number of stored objects and total bytes for each bucket or container.

#### **Space Used by Buckets**



4. Select **Bucket Details** (S3) or **Container Details** (Swift) to view a list of the spaced used and number of objects for each of the tenant's buckets or containers.

## Account Details - Account01

Display Name:	Account01 <a href="#">Sign in</a>	Quota Utilization ⓘ :	84.22%
Tenant ID:	6479 6966 4290 3892 3647	Logical Space Used ⓘ :	84.22 MB
Protocol ⓘ :	S3	Quota ⓘ :	100.00 MB
Allow Platform Services ⓘ :	Yes	Bucket Count ⓘ :	3
Uses Own Identity Source ⓘ :	No	Object Count ⓘ :	13

Overview **Bucket Details**

Export to CSV

Bucket Name	Space Used	Number of Objects
bucket-01	88.72 MB	14
bucket-02	21.75 MB	11
bucket-03	15.29 MB	3

Close

- Optionally, select **Export to CSV** to view and export a .csv file containing the usage values for each bucket or container.

You are prompted to open or save the .csv file.

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.

- If traffic classification policies are in place for a tenant, review the network traffic for that tenant.
  - Select **Configuration > Network Settings > 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.

<a href="#">+ Create</a>	<a href="#">✎ Edit</a>	<a href="#">✖ Remove</a>	<a href="#">📊 Metrics</a>
Name	Description	ID	
<input type="radio"/> ERP Traffic Control	Manage ERP traffic into the grid	cd9afbc7-b85e-4208-b6f8-7e8a79e2c574	
<input checked="" type="radio"/> Fabric Pools	Monitor Fabric Pools	223b0cbb-6968-4646-b32d-7665b8dc894b	

Displaying 2 traffic classification policies.

- Review the list of policies to identify the ones that apply to a specific tenant.
- To view metrics associated with a policy, select the radio button to the left of the policy, and then click **Metrics**.
- Analyze the graphs to determine how often the policy is limiting traffic and whether you need to adjust

the policy.

To create, edit, or delete traffic classification policies, see the instructions for administering StorageGRID.

7. Optionally, 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 the instructions for understanding audit messages for more information.

8. 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.

- Select **Help > 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.

### Related information

[Alerts reference](#)

[Review audit logs](#)

[Administer StorageGRID](#)

[Reviewing support metrics](#)

### Monitoring archival capacity

You cannot directly monitor an external archival storage system's capacity through the StorageGRID system. However, you can monitor whether the Archive Node can still send object data to the archival destination, which might indicate that an expansion of archival media is required.

### What you'll need

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

### About this task

You can monitor the Store component to check if the Archive Node can still send object data to the targeted

archival storage system. The Store Failures (ARVF) alarm might also indicate that the targeted archival storage system has reached capacity and can no longer accept object data.

## Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Overview > Main**.
3. Check the Store State and Store Status attributes to confirm that the Store component is Online with No Errors.

Attribute	Value	Icon
ARC State:	Online	
ARC Status:	No Errors	
Tivoli Storage Manager State:	Online	
Tivoli Storage Manager Status:	No Errors	
Store State:	Online	
Store Status:	No Errors	
Retrieve State:	Online	
Retrieve Status:	No Errors	
Inbound Replication Status:	No Errors	
Outbound Replication Status:	No Errors	

An offline Store component or one with errors might indicate that targeted archival storage system can no longer accept object data because it has reached capacity.

## Related information

[Administer StorageGRID](#)

## Monitoring 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.

## What you'll need

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

## About this task

You can use the Load Balancer service on Admin Nodes or Gateway Nodes, an external third-party load balancer, or the CLB service on Gateway Nodes to distribute client requests across multiple Storage Nodes.



The CLB service is deprecated.

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.

See the section on configuring client connections in the instructions for administering StorageGRID.

## Steps

1. If S3 or Swift 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 Master.

Nodes with the role of Master and nodes that are not 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 do not 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 breakdown of traffic being served by the Load Balancer service.

2. If S3 or Swift clients connect using the CLB service (deprecated), perform the following checks:

- a. Select **Nodes**.
- b. Select a Gateway 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 Master.

Nodes with the role of Master and nodes that are not in an HA group should be actively distributing requests to clients.

- d. For each Gateway Node that should be actively distributing client requests, select **Support > Tools > Grid Topology**.
- e. Select **Gateway Node > CLB > HTTP > Overview > Main**.
- f. Review the number of **Incoming Sessions - Established** to verify that the Gateway Node has been actively handling requests.

3. 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.

## Related information

[Administer StorageGRID](#)

[Viewing the Load Balancer tab](#)

## Applying hotfixes or upgrading software if necessary

If a hotfix or a new version of StorageGRID software is available, you should assess whether the update is appropriate for your system, and install it if required.

### About this task

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.

### Steps

1. Go to the NetApp Downloads page for StorageGRID.

[NetApp Downloads: StorageGRID](#)

2. Select the down arrow for the **Type/Select Version** field to see a list of the updates that are available to download:
  - **StorageGRID software versions:** 11.x.y
  - **StorageGRID hotfixes:** 11.x.y.z
3. Review the changes that are included in the update:
  - a. Select the version from the pull-down menu, and click **Go**.
  - b. Sign in using the username and password for your NetApp account.
  - c. Read the End User License Agreement, select the check box, and then select **Accept & Continue**.

The downloads page for the version you selected appears.

4. Learn about the changes included in the software version or hotfix.
  - For a new software version, see the “What’s new” topic in the instructions for upgrading StorageGRID.
  - For a hotfix, download the README file for a summary of the changes included in the hotfix.
5. If you decide a software update is required, locate the instructions before proceeding.
  - For a new software version, carefully follow the instructions for upgrading StorageGRID.
  - For a hotfix, locate the hotfix procedure in the recovery and maintenance instructions

## Related information

[Upgrade software](#)

[Maintain & recover](#)

## Managing alerts and alarms

The StorageGRID alert system is designed to inform you about operational issues that require your attention. As required, you can also use the legacy alarm system to monitor your system. This section contains the following sub-sections:

- [Comparing alerts and alarms](#)
- [Managing alerts](#)
- [Managing alarms \(legacy system\)](#)

StorageGRID includes two systems for informing you about issues.

### Alert system

The alert system is designed to be your primary tool for monitoring any issues that might occur in your StorageGRID system. The alert system provides an easy-to-use interface for detecting, evaluating, and resolving issues.

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 > 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.

### Legacy alarm system

The alarm system is supported, but is considered to be a legacy system. Like alerts, alarms are triggered at specific severity levels when attributes reach defined threshold values. However, unlike alerts, many alarms are triggered for events that you can safely ignore, which might result in an excessive number of email or SNMP notifications.

When an alarm is triggered, the following actions occur:

- The count of legacy alarms on the Dashboard is incremented.
- The alarm appears on the **Support > Alarms (legacy) > Current Alarms** page.
- An email notification is sent, assuming you have configured an SMTP server and configured one or more mailing lists.
- An SNMP notification might be sent, assuming you have configured the StorageGRID SNMP agent. (SNMP notifications are not sent for all alarms or alarm severities.)

### Comparing alerts and alarms

There are a number of similarities between the alert system and the legacy alarm system, but the alert system offers significant benefits and is easier to use.

Refer to the following table to learn how to perform similar operations.

	<b>Alerts</b>	<b>Alarms (legacy system)</b>
How do I see which alerts or alarms are active?	<ul style="list-style-type: none"> <li>Click the <b>Current alerts</b> link on the Dashboard.</li> <li>Click the alert on the <b>Nodes &gt; Overview</b> page.</li> <li>Select <b>Alerts &gt; Current</b>.</li> </ul> <p><a href="#">Viewing current alerts</a></p>	<ul style="list-style-type: none"> <li>Click the <b>Legacy alarms</b> link on the Dashboard.</li> <li>Select <b>Support &gt; Alarms (legacy) &gt; Current Alarms</b>.</li> </ul> <p><a href="#">Viewing legacy alarms</a></p>
What causes an alert or an alert to be triggered?	<p>Alerts are triggered when a Prometheus expression in an alert rule evaluates as true for the specific trigger condition and duration.</p> <p><a href="#">Viewing alert rules</a></p>	<p>Alarms are triggered when a StorageGRID attribute reaches a threshold value.</p> <p><a href="#">Alarm triggering logic (legacy system)</a></p>
If an alert or alarm is triggered, how do I resolve the underlying problem?	<p>The recommended actions for an alert are included in email notifications and are available from the Alerts pages in the Grid Manager.</p> <p>As required, additional information is provided in the StorageGRID documentation.</p> <p><a href="#">Alerts reference</a></p>	<p>You can learn about an alarm by clicking the attribute name, or you can search for an alarm code in the StorageGRID documentation.</p> <p><a href="#">Alarms reference (legacy system)</a></p>
Where can I see a list of alerts or alarms have been resolved?	<ul style="list-style-type: none"> <li>Click the <b>Recently resolved alerts</b> link on the Dashboard.</li> <li>Select <b>Alerts &gt; Resolved</b>.</li> </ul> <p><a href="#">Viewing resolved alerts</a></p>	<p>Select <b>Support &gt; Alarms (legacy) &gt; Historical Alarms</b>.</p> <p><a href="#">Reviewing historical alarms and alarm frequency (legacy system)</a></p>
Where do I manage the settings?	<p>Select <b>Alerts</b>. Then, use the options in the Alerts menu.</p> <p><a href="#">Managing alerts</a></p>	<p>Select <b>Support</b>. Then, use the options in the <b>Alarms (legacy)</b> section of the menu.</p> <p><a href="#">Managing alarms (legacy system)</a></p>

	<b>Alerts</b>	<b>Alarms (legacy system)</b>
What user group permissions do I need?	<ul style="list-style-type: none"> <li>• Anyone who can sign in to the Grid Manager can view current and resolved alerts.</li> <li>• You must have the Manage Alerts permission to manage silences, alert notifications, and alert rules.</li> </ul> <p><a href="#">Administer StorageGRID</a></p>	<ul style="list-style-type: none"> <li>• Anyone who can sign in to the Grid Manager can view legacy alarms.</li> <li>• You must have the Acknowledge Alarms permission to acknowledge alarms.</li> <li>• You must have both the Grid Topology Page Configuration and Other Grid Configuration permissions to manage global alarms and email notifications.</li> </ul> <p><a href="#">Administer StorageGRID</a></p>
How do I manage email notifications?	<p>Select <b>Alerts &gt; Email Setup</b>.</p> <p><b>Note:</b> Because alarms and alerts are independent systems, the email setup used for alarm and AutoSupport notifications is not used for alert notifications. However, you can use the same mail server for all notifications.</p> <p><a href="#">Managing alert notifications</a></p>	<p>Select <b>Support &gt; Alarms (legacy) &gt; Legacy Email Setup</b>.</p> <p><a href="#">Configuring notifications for alarms (legacy system)</a></p>
How do I manage SNMP notifications?	<p>Select <b>Configuration &gt; Monitoring &gt; SNMP Agent</b>. <a href="#">Using SNMP monitoring</a></p>	<p>Select <b>Configuration &gt; Monitoring &gt; SNMP Agent</b>. <a href="#">Using SNMP monitoring</a></p> <p><b>Note:</b> SNMP notifications are not sent for every alarm or alarm severity.</p> <p><a href="#">Alarms that generate SNMP notifications (legacy system)</a></p>

	<b>Alerts</b>	<b>Alarms (legacy system)</b>
How do I control who receives notifications?	<ol style="list-style-type: none"> <li>1. Select <b>Alerts &gt; Email Setup</b>.</li> <li>2. In the <b>Recipients</b> section, enter an email address for each email list or person who should receive an email when an alert occurs.</li> </ol> <p><a href="#">Setting up email notifications for alerts</a></p>	<ol style="list-style-type: none"> <li>1. Select <b>Support &gt; Alarms (legacy) &gt; Legacy Email Setup</b>.</li> <li>2. Creating a mailing list.</li> <li>3. Select <b>Notifications</b>.</li> <li>4. Select the mailing list.</li> </ol> <p><a href="#">Creating mailing lists for alarm notifications (legacy system)</a></p> <p><a href="#">Configuring email notifications for alarms (legacy system)</a></p>
Which Admin Nodes send notifications?	<p>A single Admin Node (the “preferred sender”).</p> <p><a href="#">Administer StorageGRID</a></p>	<p>A single Admin Node (the “preferred sender”).</p> <p><a href="#">Administer StorageGRID</a></p>
How do I suppress some notifications?	<ol style="list-style-type: none"> <li>1. Select <b>Alerts &gt; Silences</b>.</li> <li>2. Select the alert rule you want to silence.</li> <li>3. Specify a duration for the silence.</li> <li>4. Select the severity of alert you want to silence.</li> <li>5. Select to apply the silence to the entire grid, a single site, or a single node.</li> </ol> <p><b>Note:</b> If you have enabled the SNMP agent, silences also suppress SNMP traps and informs.</p> <p><a href="#">Silencing alert notifications</a></p>	<ol style="list-style-type: none"> <li>1. Select <b>Support &gt; Alarms (legacy) &gt; Legacy Email Setup</b>.</li> <li>2. Select <b>Notifications</b>.</li> <li>3. Select a mailing list, and select <b>Suppress</b>.</li> </ol> <p><a href="#">Suppressing alarm notifications for a mailing list (legacy system)</a></p>

	<b>Alerts</b>	<b>Alarms (legacy system)</b>
How do I suppress all notifications?	<p>Select <b>Alerts &gt; Silences</b>. Then, select <b>All rules</b>.</p> <p><b>Note:</b> If you have enabled the SNMP agent, silences also suppress SNMP traps and informs.</p> <p><a href="#">Silencing alert notifications</a></p>	<ol style="list-style-type: none"> <li>1. Select <b>Configuration &gt; System Settings &gt; Display Options</b>.</li> <li>2. Select the <b>Notification Suppress All</b> check box.</li> </ol> <p><b>Note:</b> Suppressing email notifications system wide also suppresses event-triggered AutoSupport emails.</p> <p><a href="#">Suppressing email notifications system wide</a></p>
How do I customize the conditions and triggers?	<ol style="list-style-type: none"> <li>1. Select <b>Alerts &gt; Alert Rules</b>.</li> <li>2. Select a default rule to edit, or select <b>Create custom rule</b>.</li> </ol> <p><a href="#">Editing an alert rule</a></p> <p><a href="#">Creating custom alert rules</a></p>	<ol style="list-style-type: none"> <li>1. Select <b>Support &gt; Alarms (legacy) &gt; Global Alarms</b>.</li> <li>2. Create a Global Custom alarm to override a Default alarm or to monitor an attribute that does not have a Default alarm.</li> </ol> <p><a href="#">Creating Global Custom alarms (legacy system)</a></p>
How do I disable an individual alert or alarm?	<ol style="list-style-type: none"> <li>1. Select <b>Alerts &gt; Alert Rules</b>.</li> <li>2. Select the rule, and click <b>Edit rule</b>.</li> <li>3. Unselect the <b>Enabled</b> check box.</li> </ol> <p><a href="#">Disabling an alert rule</a></p>	<ol style="list-style-type: none"> <li>1. Select <b>Support &gt; Alarms (legacy) &gt; Global Alarms</b>.</li> <li>2. Select the rule, and click the Edit icon.</li> <li>3. Unselect the <b>Enabled</b> check box.</li> </ol> <p><a href="#">Disabling a Default alarm (legacy system)</a></p> <p><a href="#">Disabling Global Custom alarms (legacy system)</a></p>

## Managing alerts

Alerts allow you to monitor various events and conditions within your StorageGRID system. You can manage alerts by creating custom alerts, editing or disabling the default alerts, setting up email notifications for alerts, and silencing alert notifications.

### Related information

[Viewing current alerts](#)

[Viewing resolved alerts](#)

## Viewing a specific alert

### Alerts reference

#### What alerts are

The alert system provides an easy-to-use interface for detecting, evaluating, and resolving the issues that can occur during StorageGRID operation.

- The alert system focuses on actionable problems in the system. Unlike some alarms in the legacy system, alerts are triggered for events that require your immediate attention, not for events that can safely be ignored.
- The Current Alerts page provides a user-friendly interface for viewing current problems. You can sort the listing by individual alerts and alert groups. For example, you might want to sort all alerts by node/site to see which alerts are affecting a specific node. Or, you might want to sort the alerts in a group by time triggered to find the most recent instance of a specific alert.
- The Resolved Alerts page provides similar information as on the Current Alerts page, but it allows you to search and view a history of the alerts that have been resolved, including when the alert was triggered and when it was resolved.
- Multiple alerts of the same type are grouped into one email to reduce the number of notifications. In addition, multiple alerts of the same type are shown as a group on the Alerts page. You can expand and collapse alert groups to show or hide the individual alerts. For example, if several nodes report the **Unable to communicate with node** alert at about the same time, only one email is sent and the alert is shown as a group on the Alerts page.
- Alerts use intuitive names and descriptions to help you quickly understand the problem. Alert notifications include details about the node and site affected, the alert severity, the time when the alert rule was triggered, and the current value of metrics related to the alert.
- Alert emails notifications and the alert listings on the Current Alerts and Resolved Alerts pages provide recommended actions for resolving an alert. These recommended actions often include direct links to the StorageGRID documentation center to make it easier to find and access more detailed troubleshooting procedures.
- If you need to temporarily suppress the notifications for an alert at one or more severity levels, you can easily silence a specific alert rule for a specified duration and for the entire grid, a single site, or a single node. You can also silence all alert rules, for example, during a planned maintenance procedure such as a software upgrade.
- You can edit the default alert rules as required. You can disable an alert rule completely, or change its trigger conditions and duration.
- You can create custom alert rules to target the specific conditions that are relevant to your situation and to provide your own recommended actions. To define the conditions for a custom alert, you create expressions using the Prometheus metrics available from the Metrics section of the Grid Management API.

#### Managing 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.

#### Viewing 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.

## What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Manage Alerts or Root Access permission.

## Steps

1. Select **Alerts > Alert 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			Type	Status
Name	Conditions			
<input type="radio"/> <b>Appliance battery expired</b> The battery in the appliance's storage controller has expired.	storagegrid_appliance_component_failure(type="REC_EXPIRED_BATTERY") Major > 0		Default	Enabled
<input type="radio"/> <b>Appliance battery failed</b> The battery in the appliance's storage controller has failed.	storagegrid_appliance_component_failure(type="REC_FAILED_BATTERY") Major > 0		Default	Enabled
<input type="radio"/> <b>Appliance battery has insufficient learned capacity</b> The battery in the appliance's storage controller has insufficient learned capacity.	storagegrid_appliance_component_failure(type="REC_BATTERY_WARN") Major > 0		Default	Enabled
<input type="radio"/> <b>Appliance battery near expiration</b> The battery in the appliance's storage controller is nearing expiration.	storagegrid_appliance_component_failure(type="REC_BATTERY_NEAR_EXPIRATION") Major > 0		Default	Enabled
<input type="radio"/> <b>Appliance battery removed</b> The battery in the appliance's storage controller is missing.	storagegrid_appliance_component_failure(type="REC_REMOVED_BATTERY") Major > 0		Default	Enabled
<input type="radio"/> <b>Appliance battery too hot</b> The battery in the appliance's storage controller is overheated.	storagegrid_appliance_component_failure(type="REC_BATTERY_OVERTEMP") Major > 0		Default	Enabled
<input type="radio"/> <b>Appliance cache backup device failed</b> A persistent cache backup device has failed.	storagegrid_appliance_component_failure(type="REC_CACHE_BACKUP_DEVICE_FAILED") Major > 0		Default	Enabled
<input type="radio"/> <b>Appliance cache backup device insufficient capacity</b> There is insufficient cache backup device capacity.	storagegrid_appliance_component_failure(type="REC_CACHE_BACKUP_DEVICE_INSUFFICIENT_CAPACITY") Major > 0		Default	Enabled
<input type="radio"/> <b>Appliance cache backup device write-protected</b> A cache backup device is write-protected.	storagegrid_appliance_component_failure(type="REC_CACHE_BACKUP_DEVICE_WRITE_PROTECTED") Major > 0		Default	Enabled
<input type="radio"/> <b>Appliance cache memory size mismatch</b> 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.

Column header	Description
Conditions	<p>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.</p> <ul style="list-style-type: none"> <li>• <b>Critical</b> : 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.</li> <li>• <b>Major</b> : 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.</li> <li>• <b>Minor</b> : 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 do not clear on their own to ensure they do not result in a more serious problem.</li> </ul>
Type	<p>The type of alert rule:</p> <ul style="list-style-type: none"> <li>• <b>Default</b>: 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 cannot remove a default alert rule.</li> <li>• <b>Default*</b>: 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.</li> <li>• <b>Custom</b>: An alert rule that you created. You can disable, edit, and remove custom alert rules.</li> </ul>
Status	<p>Whether this alert rule is currently enabled or disabled. The conditions for disabled alert rules are not evaluated, so no alerts are triggered.</p>

## Related information

[Alerts reference](#)

## Creating custom alert rules

You can create custom alert rules to define your own conditions for triggering alerts.

## What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Manage Alerts or Root Access permission.

## 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 simply 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.

- Do not assume a custom alert is working unless you have validated that the alert is triggered when expected.

## Steps

1. Select **Alerts > Alert 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 unselect the **Enabled** check box to determine if this alert rule is currently enabled.

If an alert rule is disabled, its expressions are not 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.

Field	Description
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.

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.

To see available metrics and to test Prometheus expressions, click the help icon  and follow the link to the Metrics section of the Grid Management API.

To learn about using the Grid Management API, see the instructions for administering StorageGRID. For details on the syntax of Prometheus queries, see the documentation for Prometheus.

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
```

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. Click **Save**.

The dialog box closes, and the new custom alert rule appears in the Alert Rules table.

#### Related information

[Administer StorageGRID](#)

[Commonly used Prometheus metrics](#)

## Editing an alert rule

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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must 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 > Alert 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 cannot be edited.

## Edit Rule - Low installed node memory

Enabled

Unique Name

Description

Recommended Actions (optional) VMware installation- [Red Hat Enterprise Linux or CentOS installation](#)
- [Ubuntu or Debian installation](#)
"/>

### Conditions

Minor

Major

Critical

Enter the amount of time a condition must continuously remain in effect before an alert is triggered.

Duration

Cancel

Save

4. Select or unselect the **Enabled** check box to determine if this alert rule is currently enabled.

If an alert rule is disabled, its expressions are not 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.

5. For custom alert rules, update the following information as required.



You cannot 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.

Field	Description
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, click the three dots to the right of the modified condition.

#### Conditions

Minor	<input type="text"/>
Major	<input type="text" value="node_memory_MemTotal_bytes &lt; 24000000000"/>
Critical	<input type="text" value="node_memory_MemTotal_bytes &lt;= 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.

To see available metrics and to test Prometheus expressions, click the help icon and follow the link to the Metrics section of the Grid Management API.

To learn about using the Grid Management API, see the instructions for administering StorageGRID. For details on the syntax of Prometheus queries, see the documentation for Prometheus.

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. Click **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.

## Related information

[Administer StorageGRID](#)

[Commonly used Prometheus metrics](#)

[Prometheus: Query basics](#)

## Disabling an alert rule

You can change the enabled/disabled state for a default or custom alert rule.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Manage Alerts or Root Access permission.

### About this task

When an alert rule is disabled, its expressions are not 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 > Alert 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 unselect the **Enabled** check box to determine if this alert rule is currently enabled.

If an alert rule is disabled, its expressions are not 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. Click **Save**.

**Disabled** appears in the **Status** column.

## Removing a custom alert rule

You can remove a custom alert rule if you no longer want to use it.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Manage Alerts or Root Access permission.

### Steps

1. Select **Alerts > Alert Rules**.

The Alert Rules page appears.

2. Select the radio button for the custom alert rule you want to remove.

You cannot remove a default alert rule.

3. Click **Remove custom rule**.

A confirmation dialog box appears.

4. Click **OK** to remove the alert rule.

Any active instances of the alert will be resolved within 10 minutes.

## Managing alert notifications

When an alert is triggered, StorageGRID can send email notifications and Simple Network Management Protocol (SNMP) notifications (traps).

## Setting 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.

### About this task

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 the section for using 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 do not 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. Alert notifications are sent by whichever Admin Node is configured to be the preferred sender. By default, the primary Admin Node is selected. For details, see the instructions for administering StorageGRID.



Trap and inform notifications are also sent when certain alarms (legacy system) are triggered at specified severity levels or higher; however, SNMP notifications are not sent for every alarm or every alarm severity.

### Related information

[Using SNMP monitoring](#)

[Silencing alert notifications](#)

[Administer StorageGRID](#)

[Alarms that generate SNMP notifications \(legacy system\)](#)

### Setting 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.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Manage Alerts or Root Access permission.

#### What you'll need

Because alarms and alerts are independent systems, the email setup used for alert notifications is not used for alarm notifications and AutoSupport messages. However, you can use the same email server for all notifications.

If your StorageGRID deployment includes multiple Admin Nodes, you can select which Admin Node should be the preferred sender of alert notifications. The same “preferred sender” is also used for alarm notifications and AutoSupport messages. By default, the primary Admin Node is selected. For details, see the instructions for administering StorageGRID.

### Steps

1. Select **Alerts > Email Setup**.

The Email Setup page appears.

#### Email Setup

You can configure the email server for alert notifications, define filters to limit the number of notifications, and enter email addresses for alert recipients.

Use these settings to define the email server used for alert notifications. These settings are not used for alarm notifications and AutoSupport. See [Managing alerts and alarms](#) in the instructions for monitoring and troubleshooting StorageGRID.

Enable Email Notifications

Save

2. Select the **Enable Email Notifications** check box 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. You must also require TLS and provide a CA certificate.

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.

#### Email (SMTP) Server

Mail Server 

Port 

Username (optional) 

Password (optional) 

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.

Click the plus icon  to add recipients.

## Email Addresses

---

Sender Email Address 	<input type="text" value="storagegrid-alerts@example.com"/>	
Recipient 1 	<input type="text" value="recipient1@example.com"/>	
Recipient 2 	<input type="text" value="recipient2@example.com"/>	 

5. In the Transport Layer Security (TLS) section, select the **Require TLS** check box if Transport Layer Security (TLS) is required for communications with the SMTP server.

- 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 click **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** check box 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 click **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 click **Browse** and select the file.



If you need to edit the email setup, click the pencil icon to update this field.

## Transport Layer Security (TLS)

Require TLS 

CA Certificate 

```
-----BEGIN CERTIFICATE-----  
1234567890abcdefghijklmnopqrstuvwxyz  
ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890  
-----END CERTIFICATE-----
```

Send Client Certificate 

Client Certificate 

```
-----BEGIN CERTIFICATE-----  
1234567890abcdefghijklmnopqrstuvwxyz  
ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890  
-----END CERTIFICATE-----
```

Private Key 

```
-----BEGIN PRIVATE KEY-----  
1234567890abcdefghijklmnopqrstuvwxyz  
ABCDEFGHIJKLMNOPQRSTUVWXYZ1234567890  
-----BEGIN PRIVATE KEY-----
```

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 are not sent for minor alerts.
Critical only	An email notification is sent only when the critical condition for an alert rule is met. Notifications are not sent for minor or major alerts.

## Filters

Severity ⓘ  Minor, major, critical  Major, critical  Critical only

Send Test Email

Save

7. When you are ready to test your email settings, perform these steps:

a. Click **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 alarm notifications and AutoSupport messages, where all Admin Nodes send the test email.

8. Click **Save**.

Sending a test email does not save your settings. You must click **Save**.

The email settings are saved.

### Related information

[Troubleshooting alert email notifications](#)

[Maintain & recover](#)

### 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.

Email notifications include the following information:

## 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

Sent from: DC1-ADM1-225 5

	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.

### Related information

[Silencing alert notifications](#)

### How StorageGRID groups alerts in email notifications

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
<p>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.</p>	<ul style="list-style-type: none"> <li>Alert A is triggered on two nodes at the same time. Only one notification is sent.</li> <li>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.</li> </ul>
<p>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.</p>	<ul style="list-style-type: none"> <li>Alert A is triggered and the minor, major, and critical alert thresholds are reached. One notification is sent for the critical alert.</li> </ul>
<p>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.</p>	<ol style="list-style-type: none"> <li>Alert A is triggered on node 1 at 08:00. No notification is sent.</li> <li>Alert A is triggered on node 2 at 08:01. No notification is sent.</li> <li>At 08:02, a notification is sent to report both instances of the alert.</li> </ol>
<p>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.</p>	<ol style="list-style-type: none"> <li>Alert A is triggered on node 1 at 08:00. A notification is sent at 08:02.</li> <li>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.</li> </ol>
<p>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.</p>	<ol style="list-style-type: none"> <li>Alert A is triggered for node 1. A notification is sent.</li> <li>Alert A is triggered for node 2. A second notification is sent.</li> <li>Alert A is resolved for node 2, but it remains active for node 1.</li> <li>Alert A is triggered again for node 2. No new notification is sent because the alert is still active for node 1.</li> </ol>
<p>StorageGRID continues to send email notifications once every 7 days until all instances of the alert are resolved or the alert rule is silenced.</p>	<ol style="list-style-type: none"> <li>Alert A is triggered for node 1 on March 8. A notification is sent.</li> <li>Alert A is not resolved or silenced. Additional notifications are sent on March 15, March 22, March 29, and so on.</li> </ol>

## Troubleshooting 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must 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 are not 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.

## Related information

[Collecting log files and system data](#)

## Silencing alert notifications

Optionally, you can configure silences to temporarily suppress alert notifications.

## What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must 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.



Because alarms and alerts are independent systems, you cannot use this functionality to suppress alarm notifications.

## 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
<i>No results found.</i>				

### 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.  <b>Note:</b> Select <b>All rules</b> 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.

Field	Description
Duration	<p>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).</p> <p><b>Note:</b> 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 <b>Services appliance link down</b> alerts and the <b>Storage appliance link down</b> alerts.</p>
Severity	<p>Which alert severity or severities should be silenced. If the alert is triggered at one of the selected severities, no notifications are sent.</p>
Nodes	<p>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.</p> <p><b>Note:</b> You cannot 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.</p>

4. Click **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	<ol style="list-style-type: none"> <li>a. Select <b>Alerts &gt; Silences</b>.</li> <li>b. From the table, select the radio button for the silence you want to edit.</li> <li>c. Click <b>Edit</b>.</li> <li>d. Change the description, the amount of time remaining, the selected severities, or the affected node.</li> <li>e. Click <b>Save</b>.</li> </ol>
Remove a silence	<ol style="list-style-type: none"> <li>a. Select <b>Alerts &gt; Silences</b>.</li> <li>b. From the table, select the radio button for the silence you want to remove.</li> <li>c. Click <b>Remove</b>.</li> <li>d. Click <b>OK</b> to confirm you want to remove this silence.</li> </ol> <p><b>Note:</b> 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.</p>

## Related information

## Managing alarms (legacy system)

The StorageGRID alarm system is the legacy system used to identify trouble spots that sometimes occur during normal operation.



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

### Related information

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

[Viewing legacy alarms](#)

[Administer StorageGRID](#)

### Alarm classes (legacy system)

A legacy alarm can belong to one of two mutually exclusive alarm classes.

### Default alarms

Default alarms are provided with each StorageGRID system and cannot be modified. However, you can disable Default alarms or override them by defining Global Custom alarms.

### Global Custom alarms

Global Custom alarms monitor the status of all services of a given type in the StorageGRID system. You can create a Global Custom alarm to override a Default alarm. You can also create a new Global Custom alarm. This can be useful for monitoring any customized conditions of your StorageGRID system.

### Related information

[Viewing Default alarms \(legacy system\)](#)

[Disabling a Default alarm \(legacy system\)](#)

[Creating Global Custom alarms \(legacy system\)](#)

[Disabling Global Custom alarms \(legacy system\)](#)

### Alarm triggering logic (legacy system)

A legacy alarm is triggered when a StorageGRID attribute reaches a threshold value that evaluates to true against a combination of alarm class (Default or Global Custom) and alarm severity level.

Icon	Color	Alarm severity	Meaning
	Yellow	Notice	The node is connected to the grid, but an unusual condition exists that does not affect normal operations.

Icon	Color	Alarm severity	Meaning
	Light Orange	Minor	The node is connected to the grid, but an abnormal condition exists that could affect operation in the future. You should investigate to prevent escalation.
	Dark Orange	Major	The node is connected to the grid, but an abnormal condition exists that currently affects operation. This requires prompt attention to prevent escalation.
	Red	Critical	The node is connected to the grid, but an abnormal condition exists that has stopped normal operations. You should address the issue immediately.

The alarm severity and corresponding threshold value can be set for every numerical attribute. The NMS service on each Admin Node continuously monitors current attribute values against configured thresholds. When an alarm is triggered, a notification is sent to all designated personnel.

Note that a severity level of Normal does not trigger an alarm.

Attribute values are evaluated against the list of enabled alarms defined for that attribute. The list of alarms is checked in the following order to find the first alarm class with a defined and enabled alarm for the attribute:

1. Global Custom alarms with alarm severities from Critical down to Notice.
2. Default alarms with alarm severities from Critical down to Notice.

After an enabled alarm for an attribute is found in the higher alarm class, the NMS service only evaluates within that class. The NMS service will not evaluate against the other lower priority classes. That is, if there is an enabled Global Custom alarm for an attribute, the NMS service only evaluates the attribute value against Global Custom alarms. Default alarms are not evaluated. Thus, an enabled Default alarm for an attribute can meet the criteria needed to trigger an alarm, but it will not be triggered because a Global Custom alarm (that does not meet the specified criteria) for the same attribute is enabled. No alarm is triggered and no notification is sent.

### Alarm triggering example

You can use this example to understand how Global Custom alarms and Default alarms are triggered.

For the following example, an attribute has a Global Custom alarm and a Default alarm defined and enabled as shown in the following table.

	Global Custom alarm threshold (enabled)	Default alarm threshold (enabled)
Notice	>= 1500	>= 1000
Minor	>= 15,000	>= 1000
Major	>=150,000	>= 250,000

If the attribute is evaluated when its value is 1000, no alarm is triggered and no notification is sent.

The Global Custom alarm takes precedence over the Default alarm. A value of 1000 does not reach the threshold value of any severity level for the Global Custom alarm. As a result, the alarm level is evaluated to be Normal.

After the above scenario, if the Global Custom alarm is disabled, nothing changes. The attribute value must be reevaluated before a new alarm level is triggered.

With the Global Custom alarm disabled, when the attribute value is reevaluated, the attribute value is evaluated against the threshold values for the Default alarm. The alarm level triggers a Notice level alarm and an email notification is sent to the designated personnel.

### Alarms of same severity

If two Global Custom alarms for the same attribute have the same severity, the alarms are evaluated with a “top down” priority.

For instance, if UMEM drops to 50MB, the first alarm is triggered (= 50000000), but not the one below it (<=100000000).



## Global Alarms

Updated: 2016-03-17 16:05:31 PDT

### Global Custom Alarms (0 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	Under 50	=	5000		
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	under100	<=	1000		

If the order is reversed, when UMEM drops to 100MB, the first alarm (<=100000000) is triggered, but not the one below it (= 50000000).



Global Custom Alarms (0 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	under10i	<=	1000		
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	Under 50	=	5000		

Default Alarms

Filter by Disabled Defaults

0 Result(s)

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
---------	---------	-----------	----------	---------	----------	-------	---------

Apply Changes

Notifications

A notification reports the occurrence of an alarm or the change of state for a service. Alarm notifications can be sent in email or using SNMP.

To avoid multiple alarms and notifications being sent when an alarm threshold value is reached, the alarm severity is checked against the current alarm severity for the attribute. If there is no change, then no further action is taken. This means that as the NMS service continues to monitor the system, it will only raise an alarm and send notifications the first time it notices an alarm condition for an attribute. If a new value threshold for the attribute is reached and detected, the alarm severity changes and a new notification is sent. Alarms are cleared when conditions return to the Normal level.

The trigger value shown in the notification of an alarm state is rounded to three decimal places. Therefore, an attribute value of 1.9999 triggers an alarm whose threshold is less than (<) 2.0, although the alarm notification shows the trigger value as 2.0.

New services

As new services are added through the addition of new grid nodes or sites, they inherit Default alarms and Global Custom alarms.

Alarms and tables

Alarm attributes displayed in tables can be disabled at the system level. Alarms cannot be disabled for individual rows in a table.

For example, the following table shows two critical Entries Available (VMFI) alarms. (Select **Support > Tools > Grid Topology**. Then, select **Storage Node > SSM > Resources**.)

You can disable the VMFI alarm so that the Critical level VMFI alarm is not triggered (both currently Critical alarms would appear in the table as green); however, you cannot disable a single alarm in a table row so that

one VMFI alarm displays as a Critical level alarm while the other remains green.

## Volumes

Mount Point	Device	Status	Size	Space Available	Total Entries	Entries Available	Write Cache
/	sda1	Online	10.6 GB	7.46 GB	655,360	559,263	Enabled
/var/local	sda3	Online	63.4 GB	59.4 GB	3,932,160	3,931,842	Unknown
/var/local/rangedb/0	sdb	Online	53.4 GB	53.4 GB	52,428,800	52,427,856	Enabled
/var/local/rangedb/1	sdc	Online	53.4 GB	53.4 GB	52,428,800	52,427,848	Enabled
/var/local/rangedb/2	sdd	Online	53.4 GB	53.4 GB	52,428,800	52,427,856	Enabled

## Acknowledging current alarms (legacy system)

Legacy alarms are triggered when system attributes reach alarm threshold values. If you want to reduce or clear the count of legacy alarms on the Dashboard, you can acknowledge the alarms.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Acknowledge Alarms permission.

### About this task

If an alarm from the legacy system is currently active, the Health panel on the Dashboard includes a **Legacy alarms** link. The number in parentheses indicates how many legacy alarms are currently active.

The screenshot shows the 'Health' panel with the following details:

- Administratively Down:** 1 (represented by a grey pentagon icon)
- Critical:** 5 (represented by a red 'X' icon)
- License Status:** 1 (represented by an orange exclamation mark icon)

At the bottom, there are navigation links: [Grid details](#), [Current alerts \(5\)](#), [Recently resolved alerts \(1\)](#), [Legacy alarms \(5\)](#) (highlighted with a yellow box), and [License](#).

Because the legacy alarm system continues to be supported, the number of legacy alarms shown on the Dashboard is incremented whenever a new alarm occurs. This count is incremented even if email notifications are no longer being sent for alarms. You can typically just ignore this number (since alerts provide a better view of the system), or you can acknowledge the alarms.



Optionally, when you have completely transitioned to the alert system, you can disable each legacy alarm to prevent it from being triggered and added to the count of legacy alarms.

When you acknowledge an alarm, it is no longer included in the count of legacy alarms unless the alarm is triggered at the next severity level or it is resolved and occurs again.



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

## Steps

1. To view the alarm, do one of the following:

- From the Health panel on the Dashboard, click **Legacy alarms**. This link appears only if at least one alarm is currently active.
- Select **Support > Alarms (legacy) > Current Alarms**. The Current Alarms page appears.

The alarm system is the legacy system. The alert system offers significant benefits and is easier to use. See [Managing alerts and alarms](#) in the instructions for monitoring and troubleshooting StorageGRID.

## Current Alarms

Last Refreshed: 2020-05-27 09:41:39 MDT

Show Acknowledged Alarms (1 - 1 of 1)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value
 Major	ORSU (Outbound Replication Status)	<a href="#">Data Center 1/DC1-ARC1/ARC</a>	Storage Unavailable	2020-05-26 21:47:18 MDT	Storage Unavailable	Storage Unavailable

Show  Records Per Page  Previous < 1 > Next

2. Click the service name in the table.

The Alarms tab for the selected service appears (**Support > Tools > Grid Topology > Grid Node > Service > Alarms**).

Overview	Alarms	Reports	Configuration
Main	History		

### Alarms: ARC (DC1-ARC1) - Replication

Updated: 2019-05-24 10:46:48 MDT

Severity	Attribute	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time	Acknowledge
 Major	ORSU (Outbound Replication Status)	Storage Unavailable	2019-05-23 21:40:08 MDT	Storage Unavailable	Storage Unavailable		<input type="checkbox"/>

Apply Changes 

3. Select the **Acknowledge** check box for the alarm, and click **Apply Changes**.

The alarm no longer appears on the Dashboard or the Current Alarms page.



When you acknowledge an alarm, the acknowledgment is not copied to other Admin Nodes. For this reason, if you view the Dashboard from another Admin Node, you might continue to see the active alarm.

4. As required, view acknowledged alarms.
- Select **Support > Alarms (legacy) > Current Alarms**.
  - Select **Show Acknowledged Alarms**.

Any acknowledged alarms are shown.

The alarm system is the legacy system. The alert system offers significant benefits and is easier to use. See [Managing alerts and alarms](#) in the instructions for monitoring and troubleshooting StorageGRID.

## Current Alarms

Last Refreshed: 2020-05-27 17:38:58 MDT

Show Acknowledged Alarms (1 - 1 of 1)

Severity	Attribute	Service	Description	Alarm Time	Trigger Value	Current Value	Acknowledge Time
 Major	ORSU (Outbound Replication Status)	<a href="#">Data Center 1/DC1-ARC1/ARC</a>	Storage Unavailable	2020-05-26 21:47:18 MDT	Storage Unavailable	Storage Unavailable	2020-05-27 17:38:14 MDT

Show  Records Per Page  Previous  1  Next

### Related information

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

### Viewing Default alarms (legacy system)

You can view the list of all Default legacy alarms.

### What you'll need

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



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

### Steps

1. Select **Support > Alarms (legacy) > Global Alarms**.
2. For Filter by, select **Attribute Code** or **Attribute Name**.
3. For equals, enter an asterisk: \*
4. Click the arrow  or press **Enter**.

All Default alarms are listed.



Global Custom Alarms (0 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input type="checkbox"/>								

Default Alarms

Filter by Attribute Code ▼ equals

221 Result(s)

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
<input checked="" type="checkbox"/>		IQSZ (Number of Objects)	Major	Greater than 10,000,000	>=	10000000	
<input checked="" type="checkbox"/>		IQSZ (Number of Objects)	Minor	Greater than 1,000,000	>=	1000000	
<input checked="" type="checkbox"/>		IQSZ (Number of Objects)	Notice	Greater than 150,000	>=	150000	
<input checked="" type="checkbox"/>		XCVF (% Completion)	Notice	Foreground Verification Completed	=	100	
<input checked="" type="checkbox"/>	ADC	ADCA (ADC Status)	Minor	Error	>=	10	
<input checked="" type="checkbox"/>	ADC	ADCE (ADC State)	Notice	Standby	=	10	
<input checked="" type="checkbox"/>	ADC	ALIS (Inbound Attribute Sessions)	Notice	Over 100	>=	100	
<input checked="" type="checkbox"/>	ADC	ALOS (Outbound Attribute Sessions)	Notice	Over 200	>=	200	

Reviewing historical alarms and alarm frequency (legacy system)

When troubleshooting an issue, you can review how often a legacy alarm was triggered in the past.

What you'll need

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



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

Steps

1. Follow these steps to get a list of all alarms triggered over a period of time.
  - a. Select **Support > Alarms (legacy) > Historical Alarms**.
  - b. Do one of the following:
    - Click one of the time periods.
    - Enter a custom range, and click **Custom Query**.

2. Follow these steps to find out how often alarms have been triggered for a particular attribute.
  - a. Select **Support > Tools > Grid Topology**.
  - b. Select **grid node > service or component > Alarms > History**.
  - c. Select the attribute from the list.
  - d. Do one of the following:
    - Click one of the time periods.
    - Enter a custom range, and click **Custom Query**.

The alarms are listed in reverse chronological order.

- e. To return to the alarms history request form, click **History**.

## Related information

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

## Creating Global Custom alarms (legacy system)

You might have used Global Custom alarms for the legacy system to address specific monitoring requirements. Global Custom alarms might have alarm levels that override Default alarms, or they might monitor attributes that do not have a Default alarm.

## What you'll need

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



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

Global Custom alarms override Default alarms. You should not change Default alarm values unless absolutely necessary. By changing Default alarms, you run the risk of concealing problems that might otherwise trigger an alarm.



Be very careful if you change alarm settings. For example, if you increase the threshold value for an alarm, you might not detect an underlying problem. Discuss your proposed changes with technical support before changing an alarm setting.

## Steps

1. Select **Support > Alarms (legacy) > Global Alarms**.
2. Add a new row to the Global Custom alarms table:
  - To add a new alarm, click **Edit**  (if this is the first entry) or **Insert** .



Global Custom Alarms (0 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input checked="" type="checkbox"/>	ARC	ARCE (ARC State)	Notice	Standby	=	10		
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)	Minor	At least 6000	>=	6000		
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)	Notice	At least 3000	>=	3000		

Default Alarms

Filter by  equals

9 Result(s)

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
<input checked="" type="checkbox"/>	ARC	ARCE (ARC State)	Notice	Standby	=	10	
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)	Minor	At least 6000	>=	6000	
<input checked="" type="checkbox"/>	ARC	AROQ (Objects Queued)	Notice	At least 3000	>=	3000	
<input checked="" type="checkbox"/>	ARC	ARRF (Request Failures)	Major	At least 1	>=	1	
<input checked="" type="checkbox"/>	ARC	ARRV (Verification Failures)	Major	At least 1	>=	1	
<input checked="" type="checkbox"/>	ARC	ARVF (Store Failures)	Major	At least 1	>=	1	
<input checked="" type="checkbox"/>	NMS	ARRC (Remaining Capacity)	Notice	Below 10	<=	10	
<input checked="" type="checkbox"/>	NMS	ARRS (Repository Status)	Major	Disconnected	<=	9	
<input checked="" type="checkbox"/>	NMS	ARRS (Repository Status)	Notice	Standby	<=	19	

Apply Changes

- To modify a Default alarm, search for the Default alarm.
  - i. Under Filter by, select either **Attribute Code** or **Attribute Name**.
  - ii. Type a search string.

Specify four characters or use wildcards (for example, A??? or AB\*). Asterisks (\*) represent multiple characters, and question marks (?) represent a single character.

- iii. Click the arrow , or press **Enter**.
- iv. In the list of results, click **Copy** next to the alarm you want to modify.

The Default alarm is copied to the Global Custom alarms table.

3. Make any necessary changes to the Global Custom alarms settings:

Heading	Description
Enabled	Select or unselect the check box to enable or disable the alarm.

Heading	Description
Attribute	<p>Select the name and code of the attribute being monitored from the list of all attributes applicable to the selected service or component.</p> <p>To display information about the attribute, click <b>Info</b>  next to the attribute's name.</p>
Severity	The icon and text indicating the level of the alarm.
Message	The reason for the alarm (connection lost, storage space below 10%, and so on).
Operator	<p>Operators for testing the current attribute value against the Value threshold:</p> <ul style="list-style-type: none"> <li>• = equals</li> <li>• &gt; greater than</li> <li>• &lt; less than</li> <li>• &gt;= greater than or equal to</li> <li>• &lt;= less than or equal to</li> <li>• ≠ not equal to</li> </ul>
Value	<p>The alarm's threshold value used to test against the attribute's actual value using the operator.</p> <p>The entry can be a single number, a range of numbers specified with a colon (1:3), or a comma-delineated list of numbers and ranges.</p>
Additional Recipients	<p>A supplementary list of email addresses to be notified when the alarm is triggered. This is in addition to the mailing list configured on the <b>Alarms &gt; Email Setup</b> page. Lists are comma delineated.</p> <p><b>Note:</b> Mailing lists require SMTP server setup in order to operate. Before adding mailing lists, confirm that SMTP is configured.</p> <p>Notifications for Custom alarms can override notifications from Global Custom or Default alarms.</p>
Actions	<p>Control buttons to:</p> <ul style="list-style-type: none"> <li> Edit a row</li> <li> Insert a row</li> <li> Delete a row</li> <li> Drag-and-drop a row up or down</li> <li> Copy a row</li> </ul>

4. Click **Apply Changes**.

## Related information

[Configuring email server settings for alarms \(legacy system\)](#)

### Disabling alarms (legacy system)

The alarms in the legacy alarm system are enabled by default, but you can disable alarms that are not required. You can also disable the legacy alarms after you have completely transitioned to the new alert system.



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

### Disabling a Default alarm (legacy system)

You can disable one of the legacy Default alarms for the entire system.

#### What you'll need

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

#### About this task

Disabling an alarm for an attribute that currently has an alarm triggered does not clear the current alarm. The alarm will be disabled the next time the attribute crosses the alarm threshold, or you can clear the triggered alarm.



Do not disable any of the legacy alarms until you have completely transitioned to the new alert system. Otherwise, you might not detect an underlying problem until it has prevented a critical operation from completing.

#### Steps

1. Select **Support > Alarms (legacy) > Global Alarms**.
2. Search for the Default alarm to disable.
  - a. In the Default Alarms section, select **Filter by > Attribute Code** or **Attribute Name**.
  - b. Type a search string.

Specify four characters or use wildcards (for example, A??? or AB\*). Asterisks (\*) represent multiple characters, and question marks (?) represent a single character.

- c. Click the arrow , or press **Enter**.



Selecting **Disabled Defaults** displays a list of all currently disabled Default alarms.

3. From the search results table, click the Edit icon  for the alarm you want to disable.



Global Custom Alarms (0 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input type="checkbox"/>								

Default Alarms

Filter by  equals

3 Result(s)

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Critical	Under 10000000	<=	10000000	
<input checked="" type="checkbox"/>	SSM	UMEM (Available Memory)	Major	Under 50000000	<=	50000000	
<input type="checkbox"/>	SSM	UMEM (Available Memory)	Minor	Under 100000000	<=	100000000	

Apply Changes

The **Enabled** check box for the selected alarm becomes active.

- Unselect the **Enabled** check box.
- Click **Apply Changes**.

The Default alarm is disabled.

Disabling Global Custom alarms (legacy system)

You can disable a legacy Global Custom alarm for the entire system.

What you'll need

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

About this task

Disabling an alarm for an attribute that currently has an alarm triggered does not clear the current alarm. The alarm will be disabled the next time the attribute crosses the alarm threshold, or you can clear the triggered alarm.

Steps

- Select **Support > Alarms (legacy) > Global Alarms**.
- In the Global Custom Alarms table, click **Edit** next to the alarm you want to disable.
- Unselect the **Enabled** check box.



Global Custom Alarms (1 Result(s))

Enabled	Service	Attribute	Severity	Message	Operator	Value	Additional Recipients	Actions
<input type="checkbox"/>	All	RDTE (Tivoli Storage Manager State)	Major	Offline	=	10		

Default Alarms

Filter by Disabled Defaults

0 Result(s)

Enabled	Service	Attribute	Severity	Message	Operator	Value	Actions
---------	---------	-----------	----------	---------	----------	-------	---------

Apply Changes

4. Click **Apply Changes**.

The Global Custom alarm is disabled.

**Clearing triggered alarms (legacy system)**

If a legacy alarm is triggered, you can clear it instead of acknowledging it.

**What you'll need**

- You must have the `Passwords.txt` file.

Disabling an alarm for an attribute that currently has an alarm triggered against it does not clear the alarm. The alarm will be disabled the next time the attribute changes. You can acknowledge the alarm or, if you want to immediately clear the alarm rather than wait for the attribute value to change (resulting in a change to the alarm state), you can clear the triggered alarm. You might find this helpful if you want to clear an alarm immediately against an attribute whose value does not change often (for example, state attributes).

1. Disable the alarm.
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. Restart the NMS service: `service nms restart`
4. Log out of the Admin Node: `exit`

The alarm is cleared.

**Related information**

## [Disabling alarms \(legacy system\)](#)

### **Configuring notifications for alarms (legacy system)**

StorageGRID system can automatically send email and SNMP notifications when an alarm is triggered or a service state changes.

By default, alarm email notifications are not sent. For email notifications, you must configure the email server and specify the email recipients. For SNMP notifications, you must configure the SNMP agent.

### **Related information**

[Using SNMP monitoring](#)

### **Types of alarm notifications (legacy system)**

When a legacy alarm is triggered, the StorageGRID system sends out two types of alarm notifications: severity level and service state.

### **Severity level notifications**

An alarm email notification is sent when a legacy alarm is triggered at a selected severity level:

- Notice
- Minor
- Major
- Critical

A mailing list receives all notifications related to the alarm for the selected severity. A notification is also sent when the alarm leaves the alarm level — either by being resolved or by entering a different alarm severity level.

### **Service state notifications**

A service state notification is sent when a service (for example, the LDR service or NMS service) enters the selected service state and when it leaves the selected service state. Service state notifications are sent when a service enters or leaves one of the following service states:

- Unknown
- Administratively Down

A mailing list receives all notifications related to changes in the selected state.

### **Related information**

[Configuring email notifications for alarms \(legacy system\)](#)

### **Configuring email server settings for alarms (legacy system)**

If you want StorageGRID to send email notifications when a legacy alarm is triggered, you must specify the SMTP mail server settings. The StorageGRID system only sends email; it cannot receive email.

### **What you'll need**

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

## About this task

Use these settings to define the SMTP server used for legacy alarm email notifications and AutoSupport email messages. These settings are not used for alert notifications.



If you use SMTP as the protocol for AutoSupport messages, you might have already configured an SMTP mail server. The same SMTP server is used for alarm email notifications, so you can skip this procedure. See the instructions for administering StorageGRID.

SMTP is the only protocol supported for sending email.

## Steps

1. Select **Support > Alarms (legacy) > Legacy Email Setup**.
2. From the Email menu, select **Server**.

The Email Server page appears. This page is also used to configure the email server for AutoSupport messages.

Use these settings to define the email server used for alarm notifications and for AutoSupport messages. These settings are not used for alert notifications. See [Managing alerts and alarms in the instructions for monitoring and troubleshooting StorageGRID](#).



## Email Server

Updated: 2016-03-17 11:11:59 PDT

### E-mail Server (SMTP) Information

Mail Server	<input type="text"/>
Port	<input type="text"/>
Authentication	<input type="text" value="Off"/>
Authentication Credentials	Username: <input type="text" value="root"/> Password: <input type="password" value="....."/>
From Address	<input type="text"/>
Test E-mail	To: <input type="text"/> <input type="checkbox"/> Send Test E-mail

Apply Changes

3. Add the following SMTP mail server settings:

Item	Description
Mail Server	IP address of the SMTP mail server. You can enter a hostname rather than an IP address if you have previously configured DNS settings on the Admin Node.

Item	Description
Port	Port number to access the SMTP mail server.
Authentication	Allows for the authentication of the SMTP mail server. By default, authentication is Off.
Authentication Credentials	Username and password of the SMTP mail server. If Authentication is set to On, a username and password to access the SMTP mail server must be provided.

4. Under **From Address**, enter a valid email address that the SMTP server will recognize as the sending email address. This is the official email address from which the email message is sent.
5. Optionally, send a test email to confirm that your SMTP mail server settings are correct.
  - a. In the **Test E-mail > To** box, add one or more addresses that you can access.

You can enter a single email address or a comma-delineated list of email addresses. Because the NMS service does not confirm success or failure when a test email is sent, you must be able to check the test recipient's inbox.

- b. Select **Send Test E-mail**.
6. Click **Apply Changes**.

The SMTP mail server settings are saved. If you entered information for a test email, that email is sent. Test emails are sent to the mail server immediately and are not sent through the notifications queue. In a system with multiple Admin Nodes, each Admin Node sends an email. Receipt of the test email confirms that your SMTP mail server settings are correct and that the NMS service is successfully connecting to the mail server. A connection problem between the NMS service and the mail server triggers the legacy MINS (NMS Notification Status) alarm at the Minor severity level.

## Related information

[Administer StorageGRID](#)

## Creating alarm email templates (legacy system)

Email templates let you customize the header, footer, and subject line of a legacy alarm email notification. You can use email templates to send unique notifications that contain the same body text to different mailing lists.

### What you'll need

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

### About this task

Use these settings to define the email templates used for legacy alarm notifications. These settings are not used for alert notifications.

Different mailing lists might require different contact information. Templates do not include the body text of the email message.

## Steps

1. Select **Support > Alarms (legacy) > Legacy Email Setup**.
2. From the Email menu, select **Templates**.
3. Click **Edit**  (or **Insert**  if this is not the first template).



## Email Templates

Updated: 2016-03-17 11:21:54 PDT

Template (0 - 0 of 0)

Template Name	Subject Prefix	Header	Footer	Actions
<input type="text" value="Template One"/>	<input type="text" value="Notifications"/>	<input type="text" value="All Email Lists"/>	<input type="text" value="From SGWS"/>	

Show  Records Per Page





4. In the new row add the following:

Item	Description
Template Name	Unique name used to identify the template. Template names cannot be duplicated.
Subject Prefix	Optional. Prefix that will appear at the beginning of an email's subject line. Prefixes can be used to easily configure email filters and organize notifications.
Header	Optional. Header text that appears at the beginning of the email message body. Header text can be used to preface the content of the email message with information such as company name and address.
Footer	Optional. Footer text that appears at the end of the email message body. Footer text can be used to close the email message with reminder information such as a contact phone number or a link to a web site.

5. Click **Apply Changes**.

A new template for notifications is added.

## Creating mailing lists for alarm notifications (legacy system)

Mailing lists let you notify recipients when a legacy alarm is triggered or when a service state changes. You must create at least one mailing list before any alarm email notifications can be sent. To send a notification to a single recipient, create a mailing list with one email address.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- If you want to specify an email template for the mailing list (custom header, footer, and subject line), you must have already created the template.

### About this task

Use these settings to define the mailing lists used for legacy alarm email notifications. These settings are not used for alert notifications.

### Steps

1. Select **Support > Alarms (legacy) > Legacy Email Setup**.
2. From the Email menu, select **Lists**.
3. Click **Edit**  (or **Insert**  if this is not the first mailing list).



## Email Lists

Updated: 2018-03-17 11:56:24 PDT

Lists (0 - 0 of 0)

Group Name	Recipients	Template	Actions
<input type="text"/>	<input type="text"/>	<input type="text"/>	  

Show  Records Per Page

« »



4. In the new row, add the following:

Item	Description
Group Name	Unique name used to identify the mailing list. Mailing list names cannot be duplicated.  <b>Note:</b> If you change the name of a mailing list, the change is not propagated to the other locations that use the mailing list name. You must manually update all configured notifications to use the new mailing list name.

Item	Description
Recipients	<p>Single email address, a previously configured mailing list, or a comma-delineated list of email addresses and mailing lists to which notifications will be sent.</p> <p><b>Note:</b> If an email address belongs to multiple mailing lists, only one email notification is sent when a notification triggering event occurs.</p>
Template	<p>Optionally, select an email template to add a unique header, footer, and subject line to notifications sent to all recipients of this mailing list.</p>

5. Click **Apply Changes**.

A new mailing list is created.

### Related information

[Creating alarm email templates \(legacy system\)](#)

### Configuring email notifications for alarms (legacy system)

In order to receive email notifications for the legacy alarm system, recipients must be a member of a mailing list and that list must be added to the Notifications page. Notifications are configured to send email to recipients only when an alarm with a specified severity level is triggered or when a service state changes. Thus, recipients only receive the notifications they need to receive.

#### What you'll need

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

#### About this task

Use these settings to configure notifications for legacy alarms. These settings are not used for alert notifications.

If an email address (or list) belongs to multiple mailing lists, only one email notification is sent when a notification triggering event occurs. For example, one group of administrators within your organization can be configured to receive notifications for all alarms regardless of severity. Another group might only require notifications for alarms with a severity of critical. You can belong to both lists. If a critical alarm is triggered, you receive only one notification.

#### Steps

1. Select **Support > Alarms (legacy) > Legacy Email Setup**.
2. From the Email menu, select **Notifications**.
3. Click **Edit**  (or **Insert**  if this is not the first notification).
4. Under E-mail List, select the mailing list.

5. Select one or more alarm severity levels and service states.
6. Click **Apply Changes**.

Notifications will be sent to the mailing list when alarms with the selected alarm severity level or service state are triggered or changed.

### Related information

[Creating mailing lists for alarm notifications \(legacy system\)](#)

[Types of alarm notifications \(legacy system\)](#)

### Suppressing alarm notifications for a mailing list (legacy system)

You can suppress alarm notifications for a mailing list when you no longer want the mailing list to receive notifications about alarms. For example, you might want to suppress notifications about legacy alarms after you have transitioned to using alert email notifications.

#### What you'll need

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

Use these settings to suppress email notifications for the legacy alarm system. These settings do not apply to alert email notifications.



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

#### Steps

1. Select **Support > Alarms (legacy) > Legacy Email Setup**.
2. From the Email menu, select **Notifications**.
3. Click **Edit**  next to the mailing list for which you want to suppress notifications.
4. Under Suppress, select the check box next to the mailing list you want to suppress, or select **Suppress** at the top of the column to suppress all mailing lists.
5. Click **Apply Changes**.

Legacy alarm notifications are suppressed for the selected mailing lists.

### Suppressing email notifications system wide

You can block the StorageGRID system's ability to send email notifications for legacy alarms and event-triggered AutoSupport messages.

#### What you'll need

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

#### About this task


Use this option to suppress email notifications for legacy alarms and event-triggered AutoSupport messages.



This option does not suppress alert email notifications. It also does not suppress weekly or user-triggered AutoSupport messages.

### Steps

1. Select **Configuration > System Settings > Display Options**.
2. From the Display Options menu, select **Options**.
3. Select **Notification Suppress All**.



**Display Options**  
Updated: 2017-03-23 18:03:48 MDT


---

Current Sender	ADMIN-DC1-ADM1
Preferred Sender	ADMIN-DC1-ADM1
GUI Inactivity Timeout	900
Notification Suppress All	<input checked="" type="checkbox"/>

**Apply Changes**

4. Click **Apply Changes**.

The Notifications page (**Configuration > Notifications**) displays the following message:



**Notifications**  
Updated: 2016-03-17 14:06:48 PDT

---

**All e-mail notifications are now suppressed.**

**Notifications (0 - 0 of 0)**

E-mail List	Suppress	Severity Levels				Service States		Actions
	<input checked="" type="checkbox"/>	Notice	Minor	Major	Critical	Unknown	Administratively Down	
	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	

Show  Records Per Page  « »

**Apply Changes**

### Related information

[Administer StorageGRID](#)

### Using 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.

- [Configuring the SNMP agent](#)
- [Updating the SNMP agent](#)

## Capabilities

Each StorageGRID node runs an SNMP agent, or daemon, that provides a management information base (MIB). The StorageGRID MIB contains table and notification definitions for alerts and alarms. 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.

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** are notifications sent by the SNMP agent that do not 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.

- **Inform**s 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 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 whichever Admin Node is configured to be the preferred sender.
- Certain alarms (legacy system) are triggered at specified severity levels or higher.



SNMP notifications are not sent for every alarm or every alarm severity.

## SNMP version support

The table provides a high-level summary of what is supported for each SNMP version.

	SNMPv1	SNMPv2c	SNMPv3
Queries	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	Traps only	Traps and informs	Traps and informs

	SNMPv1	SNMPv2c	SNMPv3
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.

## Accessing the MIB

You can access the MIB definition file at the following location on any StorageGRID node:

```
/usr/share/snmp/mibs/NETAPP-STORAGEGRID-MIB.txt
```

## Related information

[Alerts reference](#)

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

[Alarms that generate SNMP notifications \(legacy system\)](#)

[Silencing alert notifications](#)

## Configuring the SNMP agent

You can configure the StorageGRID SNMP agent if you want to use a third-party SNMP management system for read-only MIB access and notifications.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

### About this task

The StorageGRID SNMP agent supports all three versions of the SNMP protocol. You can configure the agent for one or more versions.

### Steps

1. Select **Configuration > Monitoring > SNMP Agent**.

The SNMP Agent page appears.

## SNMP Agent

You can configure SNMP for read-only MIB access and notifications. SNMPv1, SNMPv2c, SNMPv3 are supported. For SNMPv3, only User Security Model (USM) authentication is supported. All nodes in the grid share the same SNMP configuration.

Enable SNMP

Save

- To enable the SNMP agent on all grid nodes, select the **Enable SNMP** check box.

The fields for configuring an SNMP agent appear.

## SNMP Agent

You can configure SNMP for read-only MIB access and notifications. SNMPv1, SNMPv2c, SNMPv3 are supported. For SNMPv3, only User Security Model (USM) authentication is supported. All nodes in the grid share the same SNMP configuration.

Enable SNMP

System Contact

System Location

Enable SNMP Agent Notifications

Enable Authentication Traps

### Community Strings

Default Trap Community

Read-Only Community

String 1  +

### Other Configurations

Agent Addresses (0) USM Users (0) Trap Destinations (0)

+ Create Edit Remove

Internet Protocol	Transport Protocol	StorageGRID Network	Port
-------------------	--------------------	---------------------	------

No results found.

Save

- In the **System Contact** field, enter the value you want StorageGRID to provide in SNMP messages for sysContact.

The System Contact typically is an email address. The value you provide applies to all nodes in the StorageGRID system. **System Contact** can be a maximum of 255 characters.

- In the **System Location** field, enter the value you want StorageGRID to provide in SNMP messages for 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. The value you provide applies to all nodes in the StorageGRID system. **System Location** can be a maximum of 255 characters.

5. Keep the **Enable SNMP Agent Notifications** check box selected if you want the StorageGRID SNMP agent to send trap and inform notifications.

If this check box is unselected, the SNMP agent supports read-only MIB access, but it does not send any SNMP notifications.

6. Select the **Enable Authentication Traps** check box if you want the StorageGRID SNMP agent to send an authentication trap if it receives an improperly authenticated protocol message.
7. If you use SNMPv1 or SNMPv2c, complete the Community Strings section.

The fields in this section are used for community-based authentication in SNMPv1 or SNMPv2c. These fields do not apply to SNMPv3.

- a. In the **Default Trap Community** field, optionally enter the default community string you want to use for 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 cannot contain whitespace characters.

- b. For **Read-Only Community**, enter one or more community strings to allow read-only MIB access on IPv4 and IPv6 agent addresses. Click the plus sign **+** to add multiple strings.

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.

Each community string can be a maximum of 32 characters and cannot contain whitespace characters. Up to five strings are allowed.



To ensure the security of your StorageGRID system, do not use “public” as the community string. If you do not enter a community string, the SNMP agent uses the grid ID of your StorageGRID system as the community string.

8. Optionally, 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. Each agent address includes an internet protocol, a transport protocol, a StorageGRID network, and optionally a port.

If you do not configure an agent address, the default listening address is UDP port 161 on all StorageGRID networks.

- a. Click **Create**.

The Create Agent Address dialog box appears.

## Create Agent Address

Internet Protocol  IPv4  IPv6

Transport Protocol  UDP  TCP

StorageGRID Network

Port

- b. For **Internet Protocol**, select whether this address will use IPv4 or IPv6.  
By default, SNMP uses IPv4.
- c. For **Transport Protocol**, select whether this address will use UDP or TCP.  
By default, SNMP uses UDP.
- d. In the **StorageGRID Network** field, select which StorageGRID network the query will be received on.
  - Grid, Admin, and Client Networks: StorageGRID should listen for SNMP queries on all three networks.
  - Grid Network
  - Admin Network
  - Client Network



To ensure that client communications with StorageGRID remain secure, you should not create an agent address for the Client Network.

- e. In the **Port** field, optionally enter 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.



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.

- f. Click **Create**.

The agent address is created and added to the table.

## Other Configurations

Agent Addresses (2)    USM Users (2)    Trap Destinations (2)

**+ Create**   **✎ Edit**   **✕ Remove**

	Internet Protocol	Transport Protocol	StorageGRID Network	Port
<input type="radio"/>	IPv4	UDP	Grid Network	161
<input checked="" type="radio"/>	IPv4	UDP	Admin Network	161

9. If you are using SNMPv3, 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.



This step does not apply if you are only using SNMPv1 or SNMPv2c.

a. Click **Create**.

The Create USM User dialog box appears.

## Create USM User

Username	<input type="text"/>
Read-Only MIB Access 	<input type="checkbox"/>
Authoritative Engine ID 	<input type="text"/>
Security Level 	<input checked="" type="radio"/> authPriv <input type="radio"/> authNoPriv

---

### Authentication

Protocol 	SHA
Password	<input type="text"/>
Confirm Password	<input type="text"/>

---

### Privacy

Protocol 	AES
Password	<input type="text"/>
Confirm Password	<input type="text"/>

- b. Enter a unique **Username** for this USM user.

Usernames have a maximum of 32 characters and cannot contain whitespace characters. The username cannot be changed after the user is created.

- c. Select the **Read-Only MIB Access** check box if this user should have read-only access to the MIB.

If you select **Read-Only MIB Access**, the **Authoritative Engine ID** field is disabled.



USM users who have read-only MIB access cannot have engine IDs.

- d. If this user will be used in an inform destination, enter the **Authoritative Engine ID** for this user.



SNMPv3 inform destinations must have users with engine IDs. SNMPv3 trap destination cannot have users with engine IDs.

The authoritative engine ID can be from 5 to 32 bytes in hexadecimal.

e. Select a 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.

f. Enter and confirm the password this user will use for authentication.



The only authentication protocol supported is SHA (HMAC-SHA-96).

g. If you selected **authPriv**, enter and confirm the password this user will use for privacy.



The only privacy protocol supported is AES.

h. Click **Create**.

The USM user is created and added to the table.

#### Other Configurations

Agent Addresses (2)

USM Users (3)

Trap Destinations (2)

<input type="button" value="+ Create"/> <input type="button" value="✎ Edit"/> <input type="button" value="✕ Remove"/>				
	Username	Read-Only MIB Access	Security Level	Authoritative Engine ID
<input type="radio"/>	user2	✓	authNoPriv	
<input type="radio"/>	user1		authNoPriv	B3A73C2F3D6
<input checked="" type="radio"/>	user3		authPriv	59D39E801256

10. In the Other Configurations section, select the Trap Destinations tab.

The Trap Destinations tab allows you to define one or more destinations for StorageGRID trap or inform notifications. When you enable the SNMP agent and click **Save**, StorageGRID starts sending notifications to each defined destination. Notifications are sent when alerts and alarms are triggered. Standard notifications are also sent for the supported MIB-II entities (for example, ifDown and coldStart).

a. Click **Create**.

The Create Trap Destination dialog box appears.

## Create Trap Destination

Version  SNMPv1  SNMPv2C  SNMPv3

Type  Trap

Host

Port

Protocol   UDP  TCP

Community String   Use the default trap community: No default found  
(Specify the default on the SNMP Agent page.)

Use a custom community string

Custom Community String

- b. In the **Version** field, select which SNMP version will be used for this notification.
- c. Complete the form, based on which version you selected

Version	Specify this information
SNMPv1	<p><b>Note:</b> For SNMPv1, the SNMP agent can only send traps. Informs are not supported.</p> <ol style="list-style-type: none"> <li>i. In the <b>Host</b> field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap.</li> <li>ii. For <b>Port</b>, use the default (162), unless you must use another value. (162 is the standard port for SNMP traps.)</li> <li>iii. For <b>Protocol</b>, use the default (UDP). TCP is also supported. (UDP is the standard SNMP trap protocol.)</li> <li>iv. Use the default trap community, if one was specified on the SNMP Agent page, or enter a custom community string for this trap destination.</li> </ol> <p>The custom community string can be a maximum of 32 characters and cannot contain whitespace.</p>

Version	Specify this information
SNMPv2c	<ul style="list-style-type: none"> <li>i. Select whether the destination will be used for traps or informs.</li> <li>ii. In the <b>Host</b> field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap.</li> <li>iii. For <b>Port</b>, use the default (162), unless you must use another value. (162 is the standard port for SNMP traps.)</li> <li>iv. For <b>Protocol</b>, use the default (UDP). TCP is also supported. (UDP is the standard SNMP trap protocol.)</li> <li>v. Use the default trap community, if one was specified on the SNMP Agent page, or enter a custom community string for this trap destination.  The custom community string can be a maximum of 32 characters and cannot contain whitespace.</li> </ul>
SNMPv3	<ul style="list-style-type: none"> <li>i. Select whether the destination will be used for traps or informs.</li> <li>ii. In the <b>Host</b> field, enter an IPv4 or IPv6 address (or FQDN) to receive the trap.</li> <li>iii. For <b>Port</b>, use the default (162), unless you must use another value. (162 is the standard port for SNMP traps.)</li> <li>iv. For <b>Protocol</b>, use the default (UDP). TCP is also supported. (UDP is the standard SNMP trap protocol.)</li> <li>v. Select the USM user that will be used for authentication. <ul style="list-style-type: none"> <li>◦ If you selected <b>Trap</b>, only USM users without authoritative engine IDs are shown.</li> <li>◦ If you selected <b>Inform</b>, only USM users with authoritative engine IDs are shown.</li> </ul> </li> </ul>

d. Click **Create**.

The trap destination is created and added to the table.

## Other Configurations

Agent Addresses (1)

USM Users (2)

Trap Destinations (2)

<input type="button" value="+ Create"/>	<input type="button" value="Edit"/>	<input type="button" value="Remove"/>				
Version	Type	Host	Port	Protocol	Community/USM User	
<input type="radio"/> SNMPv3	Trap	local		UDP	User: Read only user	
<input type="radio"/> SNMPv3	Inform	10.10.10.10	162	UDP	User: Inform user	

11. When you have completed the SNMP agent configuration, click **Save**

The new SNMP agent configuration becomes active.

### Related information

[Silencing alert notifications](#)

### Updating the SNMP agent

You might want to disable SNMP notifications, update community strings, or add or remove agent addresses, USM users, and trap destinations.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission.

### About this task

Whenever you update the SNMP agent configuration, be aware that you must click **Save** at the bottom on the SNMP Agent page to commit any changes you have made on each tab.

### Steps

1. Select **Configuration > Monitoring > SNMP Agent**.

The SNMP Agent page appears.

2. If you want to disable the SNMP agent on all grid nodes, unselect the **Enable SNMP** check box, and click **Save**.

The SNMP agent is disabled for all grid nodes. If you later re-enable the agent, any previous SNMP configuration settings are retained.

3. Optionally, update the values you entered for **System Contact** and **System Location**.
4. Optionally, unselect the **Enable SNMP Agent Notifications** check box if you no longer want the StorageGRID SNMP agent to send trap and inform notifications.

When this check box is unselected, the SNMP agent supports read-only MIB access, but it does not send any SNMP notifications.

5. Optionally, unselect the **Enable Authentication Traps** check box if you no longer want the StorageGRID SNMP agent to send an authentication trap when it receives an improperly authenticated protocol

message.

6. If you use SNMPv1 or SNMPv2c, optionally update the Community Strings section.

The fields in this section are used for community-based authentication in SNMPv1 or SNMPv2c. These fields do not apply to SNMPv3.



If you want to remove the default community string, you must first ensure that all trap destinations use a custom community string.

7. If you want to update agent addresses, select the Agent Addresses tab in the Other Configurations section.

### Other Configurations

	Internet Protocol	Transport Protocol	StorageGRID Network	Port
<input type="radio"/>	IPv4	UDP	Grid Network	161
<input checked="" type="radio"/>	IPv4	UDP	Admin Network	161

Use this tab to specify one or more “listening addresses.” These are the StorageGRID addresses on which the SNMP agent can receive queries. Each agent address includes an internet protocol, a transport protocol, a StorageGRID network, and a port.

- a. To add an agent address, click **Create**. Then, refer to the step for agent addresses in the instructions for configuring the SNMP agent.
  - b. To edit an agent address, select the radio button for the address, and click **Edit**. Then, refer to the step for agent addresses in the instructions for configuring the SNMP agent.
  - c. To remove an agent address, select the radio button for the address, and click **Remove**. Then, click **OK** to confirm that you want to remove this address.
  - d. To commit your changes, click **Save** at the bottom of the SNMP Agent page.
8. If you want to update USM users, select the USM Users tab in the Other Configurations section.

## Other Configurations

Agent Addresses (2)

USM Users (3)

Trap Destinations (2)

	Username	Read-Only MIB Access	Security Level	Authoritative Engine ID
<input type="radio"/>	user2	✓	authNoPriv	
<input type="radio"/>	user1		authNoPriv	B3A73C2F3D6
<input checked="" type="radio"/>	user3		authPriv	59D39E801256

Use this tab to define the USM users who are authorized to query the MIB or to receive traps and informs.

- To add a USM user, click **Create**. Then, refer to the step for USM users in the instructions for configuring the SNMP agent.
- To edit a USM user, select the radio button for the user, and click **Edit**. Then, refer to the step for USM users in the instructions for configuring the SNMP agent.

The username for an existing USM user cannot 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, as described in step [SNMP trap destination](#). Otherwise, a validation error occurs when you save the SNMP agent configuration.

- To remove a USM user, select the radio button for the user, and click **Remove**. Then, click **OK** to confirm that you want to remove this user.



If the user you removed is currently selected for a trap destination, you must edit or remove the destination, as described in step [SNMP trap destination](#). Otherwise, a validation error occurs when you save the SNMP agent configuration.

### Error

422: Unprocessable Entity

Validation failed. Please check the values you entered for errors.

Undefined trap destination usmUser 'user1'

OK

- To commit your changes, click **Save** at the bottom of the SNMP Agent page.

9. If you want to update trap destinations, select the Trap Destinations tab in the Other Configurations section.

### Other Configurations

Agent Addresses (1)    USM Users (2)    **Trap Destinations (2)**

+ Create    Edit    x Remove						
	Version	Type	Host	Port	Protocol	Community/USM User
<input type="radio"/>	SNMPv3	Trap	local		UDP	User: Read only user
<input type="radio"/>	SNMPv3	Inform	10.10.10.10	162	UDP	User: Inform user

The Trap Destinations tab allows you to define one or more destinations for StorageGRID trap or inform notifications. When you enable the SNMP agent and click **Save**, StorageGRID starts sending notifications to each defined destination. Notifications are sent when alerts and alarms are triggered. Standard notifications are also sent for the supported MIB-II entities (for example, ifDown and coldStart).

- To add a trap destination, click **Create**. Then, refer to the step for trap destinations in the instructions for configuring the SNMP agent.
  - To edit a trap destination, select the radio button for the user, and click **Edit**. Then, refer to the step for trap destinations in the instructions for configuring the SNMP agent.
  - To remove a trap destination, select the radio button for the destination, and click **Remove**. Then, click **OK** to confirm that you want to remove this destination.
  - To commit your changes, click **Save** at the bottom of the SNMP Agent page.
10. When you have updated the SNMP agent configuration, click **Save**.

### Related information

[Configuring the SNMP agent](#)

## Collecting additional StorageGRID data

There are a number of additional ways to collect and analyze data that can be useful when investigating the state of your StorageGRID system or when working with technical support to resolve issues.

- [Using charts and reports](#)
- [Monitoring PUT and GET performance](#)
- [Monitoring object verification operations](#)
- [Monitoring events](#)
- [Reviewing audit messages](#)
- [Collecting log files and system data](#)
- [Manually triggering an AutoSupport message](#)
- [Viewing the Grid Topology tree](#)

- [Reviewing support metrics](#)
- [Running diagnostics](#)
- [Creating custom monitoring applications](#)


## Using charts and reports

You can use charts and reports to monitor the state of the StorageGRID system and troubleshoot problems. The types of charts and reports available in the Grid Manager include pie charts (on the Dashboard only), graphs, and text reports.

### Types of charts and graphs

Charts and graphs summarize the values of specific StorageGRID metrics and attributes.


The Grid Manager Dashboard includes pie (doughnut) charts to 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



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 > Load Balancer** tab for an Admin Node includes four Grafana charts.




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 (click 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 (click 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:



[Close](#)

- **State graph:** Available from the **Support > Tools > Grid Topology** page (click 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

[Viewing the Nodes page](#)

[Viewing the Grid Topology tree](#)

[Reviewing support metrics](#)

#### Chart legend

The lines and colors used to draw charts have specific meaning.

Sample	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.

Sample	Meaning
	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.

### Displaying charts and graphs

The Nodes page contains the graphs and charts 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.

### What you'll need

You must be signed in to the Grid Manager using a supported 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, hover 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, click the chart icon  or  to the right of the attribute name.



Charts are not available for all metrics and attributes.


**Example 1:** From the Objects tab for a Storage Node, you can click the chart icon  to see the average latency for a metadata query over time.

Queries		
Average Latency	14.43 milliseconds	
Queries - Successful	19,786	
Queries - Failed (timed-out)	0	
Queries - Failed (consistency level unmet)	0	



## Reports (Charts): DDS (DC1-S1) - Data Store

Attribute:	Average Query Latency	Vertical Scaling:	<input checked="" type="checkbox"/>	Start Date:	2020/05/20 14:57:46
Quick Query:	Last Hour	Raw Data:	<input type="checkbox"/>	End Date:	2020/05/20 15:57:46
		<input type="button" value="Update"/>			



Close

**Example 2:** From the Objects tab for a Storage Node, you can click 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



1 hour    1 day    1 week    1 month    Custom

From: 2020-10-01 [calendar icon] 12 : 45 PM PDT

To: 2020-10-01 [calendar icon] 01 : 10 PM PDT [Apply](#)



[Close](#)

5. To display charts for attributes that are not shown on the Node page, select **Support > Tools > Grid Topology**.
6. Select *grid node > component or service > Overview > 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

7. Click 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.

#### Generating charts

Charts display a graphical representation of attribute data values. You can report on a data center site, grid node, component, or service.

#### What you'll need

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

#### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **grid node > component or service > Reports > Charts**.
3. Select the attribute to report on from the **Attribute** drop-down list.
4. To force the Y-axis to start at zero, deselect the **Vertical Scaling** check box.
5. To show values at full precision, select the **Raw Data** check box, or to round values to a maximum of three

decimal places (for example, for attributes reported as percentages), deselect the **Raw Data** check box.

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. Click **Update**.

A chart 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.

9. If you want to print the chart, right-click and select **Print**, and modify any necessary printer settings and click **Print**.

#### Types of 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

### Generating 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.

### What you'll need

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

### About this task

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), unselect the **Raw Data** check box.
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.

- 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.

- 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.

- If you want to print the report, right-click and select **Print**, and modify any necessary printer settings and click **Print**.

### Exporting 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

- Select **Support > Tools > Grid Topology**.
- Create a text report.
- Click \*Export\* .



### Text Results for Attribute Send to Relay Rate

2010-07-19 08:42:09 PDT To 2010-07-20 08:42:09 PDT

1 - 5 of 254 

Time Received	Sample Time	Value
2010-07-20 08:40:46	2010-07-20 08:40:46	0.274981485 Messages/s
2010-07-20 08:38:46	2010-07-20 08:38:46	0.274989 Messages/s
2010-07-20 08:36:46	2010-07-20 08:36:46	0.283317543 Messages/s
2010-07-20 08:34:46	2010-07-20 08:34:46	0.274982493 Messages/s
2010-07-20 08:32:46	2010-07-20 08:32:46	0.291646426 Messages/s

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.

## Monitoring 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 and Swift 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 operations:

- **S3:** DELETE, GET, HEAD, Metadata Updated, POST, PUT
- **Swift:** DELETE, GET, HEAD, 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 may require investigation.

- You can download S3tester from github:<https://github.com/s3tester>

## Related information

[Review audit logs](#)

## Monitoring object verification operations

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

### What you'll need

You must be signed in to the Grid Manager using a supported browser.

### About this task

There are two verification processes that 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 Archive Nodes or on objects in a Cloud Storage Pool.



The **Unidentified corrupt object detected** alert is triggered if the system detects a corrupt object that cannot be corrected automatically.

- **Foreground verification** can be triggered by a user to more quickly verify the existence (although not the correctness) of object data.

Foreground verification allows you to verify the existence of replicated and erasure-coded object data on a specific Storage Node, checking that each object that is expected to be present is there. You can run foreground verification on all or some of a Storage Node's object stores to help determine if there are integrity problems with a storage device. Large numbers of missing objects might indicate that there is an issue with storage.

To review results from background and foreground verifications, such as corrupt or missing objects, you can look at the Nodes page for a Storage Node. You should investigate any instances of corrupt or missing object data immediately, to determine the root cause.

## Steps

1. Select **Nodes**.
2. Select **Storage Node > Objects**.
3. To check the verification results:
  - To check replicated object data verification, look at the attributes in the Verification section.

Verification		
Status	No Errors	
Rate Setting	Adaptive	
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	



Click an attribute's name in the table to display help text.

- To check erasure-coded fragment verification, select **Storage Node > ILM** and look at the attributes in the Erasure Coding Verification table.

Erasure Coding Verification		
Status	Idle	
Next Scheduled	2019-03-01 14:20:29 MST	
Fragments Verified	0	
Data Verified	0 bytes	
Corrupt Copies	0	
Corrupt Fragments	0	
Missing Fragments	0	



Click an attribute's name in the table to display help text.

### Related information

[Verifying object integrity](#)

### Monitoring 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.

The SMTT (Total events) alarm can be repeatedly triggered by issues such as network problems, power outages or upgrades. This section has information on 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.

#### Reviewing events from the Nodes page

The Nodes page lists the system events for each grid node.

1. Select **Nodes**.
2. Select **grid node > Events**.
3. At the top of the page, determine if an event is shown for **Last Event**, which describes the last event detected by the grid node.

The event is relayed verbatim from the grid node and includes any log messages with a severity level of ERROR or CRITICAL.

4. Review the table to see if the Count for any event or error is not zero.
5. After resolving issues, click **Reset event counts** to return the counts to zero.

#### Reviewing events from the Grid Topology page

The Grid Topology page also lists the system events for each grid node.

1. Select **Support > Tools > Grid Topology**.
2. Select **site > grid node > SSM > Events > Overview > Main**.

#### Related information

[Resetting event counts](#)

[Log files reference](#)

#### Reviewing previous events

You can generate a list of previous event messages to help isolate issues that occurred in the past.

1. Select **Support > Tools > Grid Topology**.
2. Select **site > grid node > SSM > Events > Reports**.
3. Select **Text**.

The **Last Event** attribute is not shown in the Charts view.

4. Change **Attribute** to **Last Event**.
5. Optionally, select a time period for **Quick Query**.
6. Click **Update**.

Overview Alarms **Reports** Configuration

Charts **Text**

 **Reports (Text): SSM (170-41) - Events**

Attribute: Last Event Results Per Page: 20 Start Date: 2009/04/15 15:19:53  
 Quick Query: Last 5 Minutes Update Raw Data:  End Date: 2009/04/15 15:24:53

**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 }

## Related information

[Using charts and reports](#)

## Resetting event counts

After resolving system events, you can reset event counts to zero.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Grid Topology Page Configuration permission.

### Steps

1. Select **Nodes > Grid Node > Events**.
2. Make sure that any event with a count greater than 0 has been resolved.
3. Click **Reset event counts**.

## Events

Last Event

No Events

Description	Count	
Abnormal Software Events	0	
Account Service Events	0	
Cassandra Heap Out Of Memory Errors	0	
Cassandra unhandled exceptions	0	
Chunk Service Events	0	
Custom Events	0	
Data-Mover Service Events	0	
File System Errors	0	
Forced Termination Events	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	
Node Errors	0	
Out Of Memory Errors	0	
Replicated State Machine Service Events	0	
SCSI Errors	0	
Stat Service Events	0	
Storage Hardware Events	0	
System Time Events	0	

[Reset event counts !\[\]\(2867456836c3afa41be469a30da6e753\_img.jpg\)](#)

## Creating 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. You can view a cumulative Count value for all custom events on the **Nodes > grid node > Events** page.
- To create a custom event based on keywords in the `/var/log/messages` or `/var/log/syslog` 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/log/messages` or `/var/log/syslog` files will be matched unless they satisfy the requirements stated above.

### Steps

1. Select **Configuration > Monitoring > Events**.
2. Click **Edit**  (or **Insert**  if this is not the first event).
3. Enter a custom event string, for example, shutdown



Events  
Updated: 2016-03-24 15:16:20 PDT

Custom Events (1 - 1 of 1)

Event	Actions
xfs internal error	   
shutdown	   

Show 10 Records Per Page Refresh Previous « 1 » Next

Apply Changes 

4. Click **Apply Changes**.
5. Select **Nodes**. Then, select **grid node > Events**.
6. 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.

Events 

Last Event

No Events

Description	Count	
Abnormal Software Events	0	
Account Service Events	0	
Cassandra Heap Out Of Memory Errors	0	
Cassandra unhandled exceptions	0	
Custom Events	0	
File System Errors	0	
Forced Termination Events	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	
Node Errors	0	
Out Of Memory Errors	0	
Replicated State Machine Service Events	0	
SCSI Errors	0	
Stat Service Events	0	
Storage Hardware Events	0	
System Time Events	0	

[Reset event counts !\[\]\(03baa5f5cfb351d60a3016f2ed404e75\_img.jpg\)](#)

### Resetting 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.

### About this task

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.

1. Select **Support > Tools > Grid Topology**.
2. Select **grid node > SSM > Events > Configuration > Main**.
3. Select the **Reset** check box for Custom Events.

Description	Count	Reset
Abnormal Software Events	0	<input type="checkbox"/>
Account Service Events	0	<input type="checkbox"/>
Cassandra Errors	0	<input type="checkbox"/>
Cassandra Heap Out Of Memory Errors	0	<input type="checkbox"/>
Custom Events	0	<input checked="" type="checkbox"/>
File System Errors	0	<input type="checkbox"/>
Forced Termination Events	0	<input type="checkbox"/>

4. Click **Apply Changes**.

## Reviewing 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 or Swift 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.

For easy access to audit logs, you can configure client access to the audit share for both NFS and CIFS (deprecated). You can also access audit log files directly from the command line of the Admin Node.

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 the instructions for audit messages. To learn how to configure audit

client access, see the instructions for administering StorageGRID.

## Related information

[Review audit logs](#)

[Administer StorageGRID](#)

## Collecting 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.

### What you'll need

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

Because application log files can be very large, the destination directory where you download the archived log files must have at least 1 GB of free space.

### Steps

1. Select **Support > Tools > Logs**.

#### Logs

Collect log files from selected grid nodes for the given time range. Download the archive package after all logs are ready.

The screenshot displays the 'Logs' collection interface. On the left, a tree view shows the hierarchy: StorageGRID Webscale Deployment (expanded), Data Center 1 (expanded), and Data Center 2 (expanded). Under Data Center 1, nodes DC1-ADM1, DC1-ARC1, DC1-G1, DC1-S1, DC1-S2, and DC1-S3 are listed with checkboxes. Under Data Center 2, nodes DC2-ADM1, DC2-S1, DC2-S2, and DC2-S3 are listed with checkboxes. On the right, the 'Log Start Time' is set to 2018-04-18 at 01:38 PM MDT, and the 'Log End Time' is set to 2018-04-18 at 05:38 PM MDT. Below these are a 'Notes' text area and a 'Provisioning Passphrase' input field. A blue 'Collect Logs' button is positioned 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. Optionally type 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.

5. Enter the provisioning passphrase for your StorageGRID system in the **Provisioning Passphrase** text box.
6. Click **Collect Logs**.

When you submit a new request, the previous collection of log files is deleted.

## Logs

Collect log files from selected grid nodes for the given time range. Download the archive package after all logs are ready.

Log collection is in progress.

### Last Collected

Log Start Time 2017-05-17 05:01:00 PDT

Log End Time 2017-05-18 09:01:00 PDT

#### Notes

Issues began approximately 7am on the 17th, then multiple alarms propagated throughout the grid.

23%

Collecting logs: 10 of 13 nodes remaining

Download

Delete

Name	Status
DC1-ADM1	Complete
DC1-G1	Error: No route to host - connect(2) for "10.96.104.212" port 22
DC1-S1	Collecting
DC1-S2	Collecting
DC1-S3	Collecting
DC2-S1	Collecting
DC2-S2	Collecting
DC2-S3	Collecting

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.

7. Click **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 click **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.

## Related information

[Log files reference](#)

## Manually triggering an AutoSupport message

To assist technical support in troubleshooting issues with your StorageGRID system, you can manually trigger an AutoSupport message to be sent.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access or Other Grid Configuration permission.

### Steps

1. Select **Support > Tools > AutoSupport**.

The AutoSupport page appears with the **Settings** tab selected.

2. Select **Send User-Triggered AutoSupport**.

StorageGRID attempts to send an AutoSupport message to technical support. 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 message again.



After sending an User-triggered AutoSupport message, refresh the AutoSupport page in your browser after 1 minute to access the most recent results.


## Related information

[Configuring email server settings for alarms \(legacy system\)](#)

## Viewing 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.

### Reviewing support metrics

When troubleshooting an issue, you can work with technical support to review detailed metrics and charts for your StorageGRID system.

#### What you'll need

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

#### About this task

The 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.

#### Steps

1. As directed by technical support, select **Support > Tools > Metrics**.

The Metrics page appears.

## Metrics

Access charts and metrics to help troubleshoot issues.

**i** 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://storagegrid.grid.com/metrics/graph>

### 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.

<a href="#">ADE</a>	<a href="#">Node</a>
<a href="#">Account Service Overview</a>	<a href="#">Node (Internal Use)</a>
<a href="#">Alertmanager</a>	<a href="#">Platform Services Commits</a>
<a href="#">Audit Overview</a>	<a href="#">Platform Services Overview</a>
<a href="#">Cassandra Cluster Overview</a>	<a href="#">Platform Services Processing</a>
<a href="#">Cassandra Network Overview</a>	<a href="#">Replicated Read Path Overview</a>
<a href="#">Cassandra Node Overview</a>	<a href="#">S3 - Node</a>
<a href="#">Cloud Storage Pool Overview</a>	<a href="#">S3 Overview</a>
<a href="#">EC - ADE</a>	<a href="#">Site</a>
<a href="#">EC - Chunk Service</a>	<a href="#">Support</a>
<a href="#">Grid</a>	<a href="#">Traces</a>
<a href="#">ILM</a>	<a href="#">Traffic Classification Policy</a>
<a href="#">Identity Service Overview</a>	<a href="#">Usage Processing</a>
<a href="#">Ingests</a>	<a href="#">Virtual Memory (vmstat)</a>

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.

Enable query history

Expression (press Shift+Enter for newlines)

Execute

- insert metric at cursor -

Graph

Console

Element	Value
no data	

[Remove Graph](#)


Add Graph



Metrics that include *private* in their names are intended for internal use only and are subject to change between StorageGRID releases without notice.

- 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.



## Related information

[Commonly used Prometheus metrics](#)

## Running diagnostics

When troubleshooting an issue, you can work with technical support to run diagnostics on your StorageGRID system and review the results.

### What you'll need

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

### About this task

The 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. In the 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](#)

✓ <b>Cassandra blocked task queue too large</b>	▼
✓ <b>Cassandra commit log latency</b>	▼
✓ <b>Cassandra commit log queue depth</b>	▼
✓ <b>Cassandra compaction queue too large</b>	▼

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 are not used for all diagnostics.)



You cannot 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 >= 75%  
 ⚠ Caution >= 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

## Related information

[Reviewing support metrics](#)[Commonly used Prometheus metrics](#)

## Creating 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 are not 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 and select **Help > API Documentation > metrics**.

## metrics Operations on 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 is beyond the scope of this guide.

### Related information

[Administer StorageGRID](#)

### Alerts reference

The following table lists all default StorageGRID alerts. As required, you can create custom alert rules to fit your system management approach.

See information about the commonly used Prometheus metrics to learn about the metrics used in some of these alerts.

Alert name	Description and recommended actions
Appliance battery expired	<p>The battery in the appliance's storage controller has expired.</p> <ol style="list-style-type: none"><li>1. Replace the battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions.<ul style="list-style-type: none"><li>◦ <a href="#">SG6000 storage appliances</a></li><li>◦ <a href="#">SG5700 storage appliances</a></li><li>◦ <a href="#">SG5600 storage appliances</a></li></ul></li><li>2. If this alert persists, contact technical support.</li></ol>

Alert name	Description and recommended actions
Appliance battery failed	<p>The battery in the appliance's storage controller has failed.</p> <ol style="list-style-type: none"> <li>1. Replace the battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> <li>2. If this alert persists, contact technical support.</li> </ol>
Appliance battery has insufficient learned capacity	<p>The battery in the appliance's storage controller has insufficient learned capacity.</p> <ol style="list-style-type: none"> <li>1. Replace the battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> <li>2. If this alert persists, contact technical support.</li> </ol>
Appliance battery near expiration	<p>The battery in the appliance's storage controller is nearing expiration.</p> <ol style="list-style-type: none"> <li>1. Replace the battery soon. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> <li>2. If this alert persists, contact technical support.</li> </ol>

Alert name	Description and recommended actions
Appliance battery removed	<p>The battery in the appliance's storage controller is missing.</p> <ol style="list-style-type: none"> <li>1. Install a battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> <li>2. If this alert persists, contact technical support.</li> </ol>
Appliance battery too hot	<p>The battery in the appliance's storage controller is overheated.</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>2. Investigate possible reasons for the temperature increase, such as a fan or HVAC failure.</li> <li>3. If this alert persists, contact technical support.</li> </ol>
Appliance BMC communication error	<p>Communication with the baseboard management controller (BMC) has been lost.</p> <ol style="list-style-type: none"> <li>1. Confirm that the BMC is operating normally. Select <b>Nodes</b>, and then select the <b>Hardware</b> tab for the appliance node. Locate the Compute Controller BMC IP field, and browse to that IP.</li> <li>2. Attempt to restore BMC communications by placing the node into maintenance mode and then powering the appliance off and back on. See the installation and maintenance instructions for your appliance. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> </ul> </li> <li>3. If this alert persists, contact technical support.</li> </ol>
Appliance cache backup device failed	<p>A persistent cache backup device has failed.</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>2. Contact technical support.</li> </ol>

Alert name	Description and recommended actions
Appliance cache backup device insufficient capacity	There is insufficient cache backup device capacity. Contact technical support.
Appliance cache backup device write-protected	A cache backup device is write-protected. Contact technical support.
Appliance cache memory size mismatch	The two controllers in the appliance have different cache sizes. Contact technical support.
Appliance compute controller chassis temperature too high	<p>The temperature of the compute controller in a StorageGRID appliance has exceeded a nominal threshold.</p> <ol style="list-style-type: none"> <li>1. Check the hardware components for overheating conditions, and follow the recommended actions: <ul style="list-style-type: none"> <li>◦ If you have an SG100, SG1000, or SG6000, use the BMC.</li> <li>◦ If you have an SG5600 or SG5700, use SANtricity System Manager.</li> </ul> </li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> </ul> </li> </ol>

Alert name	Description and recommended actions
Appliance compute controller CPU temperature too high	<p>The temperature of the CPU in the compute controller in a StorageGRID appliance has exceeded a nominal threshold.</p> <ol style="list-style-type: none"> <li>1. Check the hardware components for overheating conditions, and follow the recommended actions: <ul style="list-style-type: none"> <li>◦ If you have an SG100, SG1000, or SG6000, use the BMC.</li> <li>◦ If you have an SG5600 or SG5700, use SANtricity System Manager.</li> </ul> </li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> </ul> </li> </ol>
Appliance compute controller needs attention	<p>A hardware fault has been detected in the compute controller of a StorageGRID appliance.</p> <ol style="list-style-type: none"> <li>1. Check the hardware components for errors, and follow the recommended actions: <ul style="list-style-type: none"> <li>◦ If you have an SG100, SG1000, or SG6000, use the BMC.</li> <li>◦ If you have an SG5600 or SG5700, use SANtricity System Manager.</li> </ul> </li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> </ul> </li> </ol>

Alert name	Description and recommended actions
Appliance compute controller power supply A has a problem	<p>Power supply A in the compute controller has a problem. This alert might indicate that the power supply has failed or that it has a problem providing power.</p> <ol style="list-style-type: none"> <li>1. Check the hardware components for errors, and follow the recommended actions: <ul style="list-style-type: none"> <li>◦ If you have an SG100, SG1000, or SG6000, use the BMC.</li> <li>◦ If you have an SG5600 or SG5700, use SANtricity System Manager.</li> </ul> </li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> </ul> </li> </ol>
Appliance compute controller power supply B has a problem	<p>Power supply B in the compute controller has a problem. This alert might indicate that the power supply has failed or that it has a problem providing power.</p> <ol style="list-style-type: none"> <li>1. Check the hardware components for errors, and follow the recommended actions: <ul style="list-style-type: none"> <li>◦ If you have an SG100, SG1000, or SG6000, use the BMC.</li> <li>◦ If you have an SG5600 or SG5700, use SANtricity System Manager.</li> </ul> </li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> </ul> </li> </ol>

Alert name	Description and recommended actions
Appliance compute hardware monitor service stalled	<p>The service that monitors storage hardware status has stopped reporting data.</p> <ol style="list-style-type: none"> <li>1. Check the status of the eos-system-status service in the base-os.</li> <li>2. If the service is in a stopped or error state, restart the service.</li> <li>3. If this alert persists, contact technical support.</li> </ol>
Appliance Fibre Channel fault detected	<p>There is a problem with the Fibre Channel connection between the storage and compute controllers in the appliance.</p> <ol style="list-style-type: none"> <li>1. Check the hardware components for errors (<b>Nodes &gt; <i>appliance node</i> &gt; Hardware</b>). If the status of any of the components is not “Nominal”, take these actions: <ol style="list-style-type: none"> <li>a. Verify that the Fibre Channel cables between controllers are completely connected.</li> <li>b. Ensure that the Fibre Channel cables are free of excessive bends.</li> <li>c. Confirm that the SFP+ modules are properly seated.</li> </ol> <p><b>Note:</b> If this problem persists, the StorageGRID system might take the problematic connection offline automatically.</p> </li> <li>2. If necessary, replace components. See the installation and maintenance instructions for your appliance.</li> </ol>
Appliance Fibre Channel HBA port failure	<p>A Fibre Channel HBA port is failing or has failed. Contact technical support.</p>
Appliance flash cache drives non-optimal	<p>The drives used for the SSD cache are non-optimal.</p> <ol style="list-style-type: none"> <li>1. Replace the SSD cache drives. See the appliance installation and maintenance instructions. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> <li>2. If this alert persists, contact technical support.</li> </ol>

Alert name	Description and recommended actions
Appliance interconnect/battery canister removed	<p>The interconnect/battery canister is missing.</p> <ol style="list-style-type: none"> <li>1. Replace the battery. The steps to remove and replace a battery are included in the procedure for replacing a storage controller in the appliance installation and maintenance instructions. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> <li>2. If this alert persists, contact technical support.</li> </ol>
Appliance LACP port missing	<p>A port on a StorageGRID appliance is not participating in the LACP bond.</p> <ol style="list-style-type: none"> <li>1. Check the configuration for the switch. Ensure the interface is configured in the correct link aggregation group.</li> <li>2. If this alert persists, contact technical support.</li> </ol>
Appliance overall power supply degraded	<p>The power of a StorageGRID appliance has deviated from the recommended operating voltage.</p> <ol style="list-style-type: none"> <li>1. Check the status of power supply A and B to determine which power supply is operating abnormally, and follow the recommended actions: <ul style="list-style-type: none"> <li>◦ If you have an SG100, SG1000, or SG6000, use the BMC.</li> <li>◦ If you have an SG5600 or SG5700, use SANtricity System Manager.</li> </ul> </li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> </ul> </li> </ol>

Alert name	Description and recommended actions
Appliance storage controller A failure	<p>Storage controller A in a StorageGRID appliance has failed.</p> <ol style="list-style-type: none"> <li>1. Use SANtricity System Manager to check hardware components, and follow the recommended actions.</li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> </ol>
Appliance storage controller B failure	<p>Storage controller B in a StorageGRID appliance has failed.</p> <ol style="list-style-type: none"> <li>1. Use SANtricity System Manager to check hardware components, and follow the recommended actions.</li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> </ol>
Appliance storage controller drive failure	<p>One or more drives in a StorageGRID appliance has failed or is not optimal.</p> <ol style="list-style-type: none"> <li>1. Use SANtricity System Manager to check hardware components, and follow the recommended actions.</li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> </ol>

Alert name	Description and recommended actions
Appliance storage controller hardware issue	<p>SANtricity software is reporting "Needs attention" for a component in a StorageGRID appliance.</p> <ol style="list-style-type: none"> <li>1. Use SANtricity System Manager to check hardware components, and follow the recommended actions.</li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> </ol>
Appliance storage controller power supply A failure	<p>Power supply A in a StorageGRID appliance has deviated from the recommended operating voltage.</p> <ol style="list-style-type: none"> <li>1. Use SANtricity System Manager to check hardware components, and follow the recommended actions.</li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> </ol>
Appliance storage controller power supply B failure	<p>Power supply B in a StorageGRID appliance has deviated from the recommended operating voltage.</p> <ol style="list-style-type: none"> <li>1. Use SANtricity System Manager to check hardware components, and follow the recommended actions.</li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> </ol>

Alert name	Description and recommended actions
Appliance storage hardware monitor service stalled	<p>The service that monitors storage hardware status has stopped reporting data.</p> <ol style="list-style-type: none"> <li>1. Check the status of the eos-system-status service in the base-os.</li> <li>2. If the service is in a stopped or error state, restart the service.</li> <li>3. If this alert persists, contact technical support.</li> </ol>
Appliance storage shelves degraded	<p>The status of one of the components in the storage shelf for a storage appliance is degraded.</p> <ol style="list-style-type: none"> <li>1. Use SANtricity System Manager to check hardware components, and follow the recommended actions.</li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware: <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> </ol>
Appliance temperature exceeded	<p>The nominal or maximum temperature for the appliance's storage controller has been exceeded.</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>2. Investigate possible reasons for the temperature increase, such as a fan or HVAC failure.</li> <li>3. If this alert persists, contact technical support.</li> </ol>
Appliance temperature sensor removed	<p>A temperature sensor has been removed. Contact technical support.</p>

Alert name	Description and recommended actions
Cassandra auto-compactor error	<p>The Cassandra auto-compactor has experienced an error. The Cassandra auto-compactor exists on all Storage Nodes and manages the size of the Cassandra database for overwrite and delete heavy workloads. While this condition persists, certain workloads will experience unexpectedly high metadata consumption.</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>2. Contact technical support.</li> </ol>
Cassandra auto-compactor metrics out of date	<p>The metrics that describe the Cassandra auto-compactor are out of date. The Cassandra auto-compactor exists on all Storage Nodes and manages the size of the Cassandra database for overwrite and delete heavy workloads. While this alert persists, certain workloads will experience unexpectedly high metadata consumption.</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>2. Contact technical support.</li> </ol>

Alert name	Description and recommended actions
Cassandra communication error	<p>The nodes that run the Cassandra service are having trouble communicating with each other. This alert indicates that something is interfering with node-to-node communications. There might be a network issue or the Cassandra service might be down on one or more Storage Nodes.</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting one or more Storage Nodes. This alert might be resolved when you resolve the other alert.</li> <li>2. Check for a network issue that might be affecting one or more Storage Nodes.</li> <li>3. Select <b>Support &gt; Tools &gt; Grid Topology</b>.</li> <li>4. For each Storage Node in your system, select <b>SSM &gt; Services</b>. Ensure that the status of the Cassandra service is "Running."</li> <li>5. If Cassandra is not running, follow the steps for starting or restarting a service in the recovery and maintenance instructions.</li> <li>6. If all instances of the Cassandra service are now running and the alert is not resolved, contact technical support.</li> </ol> <p><a href="#">Maintain &amp; recover</a></p>
Cassandra compactions overloaded	<p>The Cassandra compaction process is overloaded. If the compaction process is overloaded, read performance might be degraded and RAM might be used up. The Cassandra service might also become unresponsive or crash.</p> <ol style="list-style-type: none"> <li>1. Restart the Cassandra service by following the steps for restarting a service in the recovery and maintenance instructions.</li> <li>2. If this alert persists, contact technical support.</li> </ol> <p><a href="#">Maintain &amp; recover</a></p>
Cassandra repair metrics out of date	<p>The metrics that describe Cassandra repair jobs are out of date. If this condition persists for more than 48 hours, client queries, such as bucket listings, might show deleted data.</p> <ol style="list-style-type: none"> <li>1. Reboot the node. From the Grid Manager, go to <b>Nodes</b>, select the node, and select the <b>Tasks</b> tab.</li> <li>2. If this alert persists, contact technical support.</li> </ol>

Alert name	Description and recommended actions
Cassandra repair progress slow	<p>The progress of Cassandra database repairs is slow. When database repairs are slow, Cassandra data consistency operations are impeded. If this condition persists for more than 48 hours, client queries, such as bucket listings, might show deleted data.</p> <ol style="list-style-type: none"> <li>1. Confirm that all Storage Nodes are online and there are no networking-related alerts.</li> <li>2. Monitor this alert for up to 2 days to see if the issue resolves on its own.</li> <li>3. If database repairs continue to proceed slowly, contact technical support.</li> </ol>
Cassandra repair service not available	<p>The Cassandra repair service is not available. The Cassandra repair service exists on all Storage Nodes and provides critical repair functions for the Cassandra database. If this condition persists for more than 48 hours, client queries, such as bucket listings, might show deleted data.</p> <ol style="list-style-type: none"> <li>1. Select <b>Support &gt; Tools &gt; Grid Topology</b>.</li> <li>2. For each Storage Node in your system, select <b>SSM &gt; Services</b>. Ensure that the status of the Cassandra Reaper service is "Running."</li> <li>3. If Cassandra Reaper is not running, follow the steps for starting or restarting a service in the recovery and maintenance instructions.</li> <li>4. If all instances of the Cassandra Reaper service are now running and the alert is not resolved, contact technical support.</li> </ol> <p><a href="#">Maintain &amp; recover</a></p>
Cloud Storage Pool connectivity error	<p>The health check for Cloud Storage Pools detected one or more new errors.</p> <ol style="list-style-type: none"> <li>1. Go to the Cloud Storage Pools section of the Storage Pools page.</li> <li>2. Look at the Last Error column to determine which Cloud Storage Pool has an error.</li> <li>3. See the instructions for managing objects with information lifecycle management.</li> </ol> <p><a href="#">Manage objects with ILM</a></p>

Alert name	Description and recommended actions
DHCP lease expired	<p>The DHCP lease on a network interface has expired.If the DHCP lease has expired, follow the recommended actions:</p> <ol style="list-style-type: none"> <li>1. Ensure there is connectivity between this node and the DHCP server on the affected interface.</li> <li>2. Ensure there are IP addresses available to assign in the affected subnet on the DHCP server.</li> <li>3. Ensure there is a permanent reservation for the IP address configured in the DHCP server. Or, use the StorageGRID Change IP tool to assign a static IP address outside of the DHCP address pool. See the recovery and maintenance instructions.</li> </ol> <p><a href="#">Maintain &amp; recover</a></p>
DHCP lease expiring soon	<p>The DHCP lease on a network interface is expiring soon.To prevent the DHCP lease from expiring, follow the recommended actions:</p> <ol style="list-style-type: none"> <li>1. Ensure there is connectivity between this node and the DHCP server on the affected interface.</li> <li>2. Ensure there are IP addresses available to assign in the affected subnet on the DHCP server.</li> <li>3. Ensure there is a permanent reservation for the IP address configured in the DHCP server. Or, use the StorageGRID Change IP tool to assign a static IP address outside of the DHCP address pool. See the recovery and maintenance instructions.</li> </ol> <p><a href="#">Maintain &amp; recover</a></p>

Alert name	Description and recommended actions
DHCP server unavailable	<p>The DHCP server is unavailable. The StorageGRID node is unable to contact your DHCP server. The DHCP lease for the node's IP address cannot be validated.</p> <ol style="list-style-type: none"> <li>1. Ensure there is connectivity between this node and the DHCP server on the affected interface.</li> <li>2. Ensure there are IP addresses available to assign in the affected subnet on the DHCP server.</li> <li>3. Ensure there is a permanent reservation for the IP address configured in the DHCP server. Or, use the StorageGRID Change IP tool to assign a static IP address outside of the DHCP address pool. See the recovery and maintenance instructions.</li> </ol> <p><a href="#">Maintain &amp; recover</a></p>
Disk I/O is very slow	<p>Very slow disk I/O might be impacting StorageGRID performance.</p> <ol style="list-style-type: none"> <li>1. If the issue is related to a storage appliance node, use SANtricity System Manager to check for faulty drives, drives with predicted faults, or in-progress drive repairs. Also check the status of the Fibre Channel or SAS links between the appliance compute and storage controllers to see if any links are down or showing excessive error rates.</li> <li>2. Examine the storage system that hosts this node's volumes to determine, and correct, the root cause of the slow I/O.</li> <li>3. If this alert persists, contact technical support.</li> </ol> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>Affected nodes might disable services and reboot themselves to avoid impacting overall grid performance. When the underlying condition is cleared and these nodes detect normal I/O performance, they will return to full service automatically.</p> </div>

Alert name	Description and recommended actions
Email notification failure	<p>The email notification for an alert could not be sent. This alert is triggered when an alert email notification fails or a test email (sent from the <b>Alerts &gt; Email Setup</b> page) cannot be delivered.</p> <ol style="list-style-type: none"> <li>1. Sign in to Grid Manager from the Admin Node listed in the <b>Site/Node</b> column of the alert.</li> <li>2. Go to the <b>Alerts &gt; Email Setup</b> page, check the settings, and change them if required.</li> <li>3. Click <b>Send Test Email</b>, and check the inbox of a test recipient for the email. A new instance of this alert might be triggered if the test email cannot be sent.</li> <li>4. If the test email could not be sent, confirm your email server is online.</li> <li>5. If the server is working, select <b>Support &gt; Tools &gt; Logs</b>, and collect the log for the Admin Node. Specify a time period that is 15 minutes before and after the time of the alert.</li> <li>6. Extract the downloaded archive, and review the contents of <code>prometheus.log</code> (<code>_/GID&lt;gid&gt;&lt;time_stamp&gt;/&lt;site_node&gt;/&lt;time_stamp&gt;/metrics/prometheus.log</code>).</li> <li>7. If you are unable to resolve the problem, contact technical support.</li> </ol>
Expiration of certificates configured on Client Certificates page	<p>One or more certificates configured on the Client Certificates page are about to expire.</p> <ol style="list-style-type: none"> <li>1. Select <b>Configuration &gt; Access Control &gt; Client Certificates</b>.</li> <li>2. Select a certificate that will expire soon.</li> <li>3. Select <b>Edit</b> to upload or generate a new certificate.</li> <li>4. Repeat these steps for each certificate that will expire soon.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>

Alert name	Description and recommended actions
Expiration of load balancer endpoint certificate	<p>One or more load balancer endpoint certificates are about to expire.</p> <ol style="list-style-type: none"> <li>1. Select <b>Configuration &gt; Network Settings &gt; Load Balancer Endpoints</b>.</li> <li>2. Select an endpoint that has a certificate that will expire soon.</li> <li>3. Select <b>Edit endpoint</b> to upload or generate a new certificate.</li> <li>4. Repeat these steps for each endpoint that has an expired certificate or one that will expire soon.</li> </ol> <p>For more information about managing load balancer endpoints, see the instructions for administering StorageGRID.</p> <p><a href="#">Administer StorageGRID</a></p>
Expiration of server certificate for Management Interface	<p>The server certificate used for the management interface is about to expire.</p> <ol style="list-style-type: none"> <li>1. Select <b>Configuration &gt; Network Settings &gt; Server Certificates</b>.</li> <li>2. In the Management Interface Server Certificate section, upload a new certificate.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>
Expiration of server certificate for Storage API Endpoints	<p>The server certificate used for accessing storage API endpoints is about to expire.</p> <ol style="list-style-type: none"> <li>1. Select <b>Configuration &gt; Network Settings &gt; Server Certificates</b>.</li> <li>2. In the Object Storage API Service Endpoints Server Certificate section, upload a new certificate.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>

Alert name	Description and recommended actions
Grid Network MTU mismatch	<p>The maximum transmission unit (MTU) setting for the Grid Network interface (eth0) differs significantly across nodes in the grid. 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.</p> <p><a href="#">Troubleshooting the Grid Network MTU mismatch alert</a></p>
High Java heap use	<p>A high percentage of Java heap space is being used. If the Java heap becomes full, metadata services can become unavailable and client requests can fail.</p> <ol style="list-style-type: none"> <li>1. Review the ILM activity on the Dashboard. This alert might resolve on its own when the ILM workload decreases.</li> <li>2. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>3. If this alert persists, contact technical support.</li> </ol>
High latency for metadata queries	<p>The average time for Cassandra metadata queries is too long. An increase in query latency can be caused by a hardware change, such as replacing a disk, or a workload change, such as a sudden increase in ingests.</p> <ol style="list-style-type: none"> <li>1. Determine if there were any hardware or workload changes around the time the query latency increased.</li> <li>2. If you are unable to resolve the problem, contact technical support.</li> </ol>

Alert name	Description and recommended actions
<p>Identity federation synchronization failure</p>	<p>Unable to synchronize federated groups and users from the identity source.</p> <ol style="list-style-type: none"> <li>1. Confirm that the configured LDAP server is online and available.</li> <li>2. Review the settings on the Identity Federation page. Confirm that all values are current. See “Configuring a federated identity source” in the instructions for administering StorageGRID.</li> <li>3. Click <b>Test Connection</b> to validate the settings for the LDAP server.</li> <li>4. If you cannot resolve the issue, contact technical support.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>
<p>ILM placement unachievable</p>	<p>A placement instruction in an ILM rule cannot be achieved for certain objects. This alert indicates that a node required by a placement instruction is unavailable or that an ILM rule is misconfigured. For example, a rule might specify more replicated copies than there are Storage Nodes.</p> <ol style="list-style-type: none"> <li>1. Ensure that all nodes are online.</li> <li>2. If all nodes are online, review the placement instructions in all ILM rules that are used the active ILM policy. Confirm that there are valid instructions for all objects. See the instructions for managing objects with information lifecycle management.</li> <li>3. As required, update rule settings and activate a new policy.</li> </ol> <p> It might take up to 1 day for the alert to clear.</p> <ol style="list-style-type: none"> <li>4. If the problem persists, contact technical support.</li> </ol> <p> This alert might appear during an upgrade and could persist for 1 day after the upgrade is completed successfully. When this alert is triggered by an upgrade, it will clear on its own.</p> <p><a href="#">Manage objects with ILM</a></p>

Alert name	Description and recommended actions
ILM scan period too long	<p>The time required to scan, evaluate objects, and apply ILM is too long. If the estimated time to complete a full ILM scan of all objects is too long (see <b>Scan Period - Estimated</b> on the Dashboard), the active ILM policy might not be applied to newly ingested objects. Changes to the ILM policy might not be applied to existing objects.</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>2. Confirm that all Storage Nodes are online.</li> <li>3. Temporarily reduce the amount of client traffic. For example, from the Grid Manager, select <b>Configuration &gt; Network Settings &gt; Traffic Classification</b>, and create a policy that limits bandwidth or the number of requests.</li> <li>4. If disk I/O or CPU are overloaded, try to reduce the load or increase the resource.</li> <li>5. If necessary, update ILM rules to use synchronous placement (default for rules created after StorageGRID 11.3).</li> <li>6. If this alert persists, contact technical support.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>
ILM scan rate low	<p>The ILM scan rate is set to less than 100 objects/second. This alert indicates that someone has changed the ILM scan rate for your system to less than 100 objects/second (default: 400 objects/second). The active ILM policy might not be applied to newly ingested objects. Subsequent changes to the ILM policy will not be applied to existing objects.</p> <ol style="list-style-type: none"> <li>1. Determine if a temporary change was made to the ILM scan rate as part of an ongoing support investigation.</li> <li>2. Contact technical support.</li> </ol> <p> Never change the ILM scan rate without contacting technical support.</p>

Alert name	Description and recommended actions
KMS CA certificate expiration	<p>The certificate authority (CA) certificate used to sign the key management server (KMS) certificate is about to expire.</p> <ol style="list-style-type: none"> <li>1. Using the KMS software, update the CA certificate for the key management server.</li> <li>2. From the Grid Manager, select <b>Configuration &gt; System Settings &gt; Key Management Server</b>.</li> <li>3. Select the KMS that has a certificate status warning.</li> <li>4. Select <b>Edit</b>.</li> <li>5. Select <b>Next</b> to go to Step 2 (Upload Server Certificate).</li> <li>6. Select <b>Browse</b> to upload the new certificate.</li> <li>7. Select <b>Save</b>.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>
KMS client certificate expiration	<p>The client certificate for a key management server is about to expire.</p> <ol style="list-style-type: none"> <li>1. From the Grid Manager, select <b>Configuration &gt; System Settings &gt; Key Management Server</b>.</li> <li>2. Select the KMS that has a certificate status warning.</li> <li>3. Select <b>Edit</b>.</li> <li>4. Select <b>Next</b> to go to Step 3 (Upload Client Certificates).</li> <li>5. Select <b>Browse</b> to upload the new certificate.</li> <li>6. Select <b>Browse</b> to upload the new private key.</li> <li>7. Select <b>Save</b>.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>
KMS configuration failed to load	<p>The configuration for the key management server exists but failed to load.</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>2. If this alert persists, contact technical support.</li> </ol>

Alert name	Description and recommended actions
KMS connectivity error	<p>An appliance node could not connect to the key management server for its site.</p> <ol style="list-style-type: none"> <li>1. From the Grid Manager, select <b>Configuration &gt; System Settings &gt; Key Management Server</b>.</li> <li>2. Confirm that the port and hostname entries are correct.</li> <li>3. Confirm that the server certificate, client certificate, and the client certificate private key are correct and not expired.</li> <li>4. Ensure that firewall settings allow the appliance node to communicate with the specified KMS.</li> <li>5. Correct any networking or DNS issues.</li> <li>6. If you need assistance or this alert persists, contact technical support.</li> </ol>
KMS encryption key name not found	<p>The configured key management server does not have an encryption key that matches the name provided.</p> <ol style="list-style-type: none"> <li>1. Confirm that the KMS assigned to the site is using the correct name for the encryption key and any prior versions.</li> <li>2. If you need assistance or this alert persists, contact technical support.</li> </ol>
KMS encryption key rotation failed	<p>All appliance volumes were decrypted, but one or more volumes could not rotate to the latest key. Contact technical support.</p>
KMS is not configured	<p>No key management server exists for this site.</p> <ol style="list-style-type: none"> <li>1. From the Grid Manager, select <b>Configuration &gt; System Settings &gt; Key Management Server</b>.</li> <li>2. Add a KMS for this site or add a default KMS.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>

Alert name	Description and recommended actions
KMS key failed to decrypt an appliance volume	<p>One or more volumes on an appliance with node encryption enabled could not be decrypted with the current KMS key.</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>2. Ensure that the key management server (KMS) has the configured encryption key and any previous key versions.</li> <li>3. If you need assistance or this alert persists, contact technical support.</li> </ol>
KMS server certificate expiration	<p>The server certificate used by the key management server (KMS) is about to expire.</p> <ol style="list-style-type: none"> <li>1. Using the KMS software, update the server certificate for the key management server.</li> <li>2. If you need assistance or this alert persists, contact technical support.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>
Large audit queue	<p>The disk queue for audit messages is full.</p> <ol style="list-style-type: none"> <li>1. Check the load on the system—if there have been a significant number of transactions, the alert should resolve itself over time, and you can ignore the alert.</li> <li>2. If the alert 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.</li> <li>3. 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 (<b>Configuration &gt; Monitoring &gt; Audit</b>).</li> </ol> <p><a href="#">Review audit logs</a></p>
Low audit log disk capacity	<p>The space available for audit logs is low.</p> <ol style="list-style-type: none"> <li>1. Monitor this alert to see if the issue resolves on its own and the disk space becomes available again.</li> <li>2. Contact technical support if the available space continues to decrease.</li> </ol>

Alert name	Description and recommended actions
Low available node memory	<p>The amount of RAM available on a node is low.Low available RAM could indicate a change in the workload or a memory leak with one or more nodes.</p> <ol style="list-style-type: none"> <li>1. Monitor this alert to see if the issue resolves on its own.</li> <li>2. If the available memory falls below the major alert threshold, contact technical support.</li> </ol>
Low free space for storage pool	<p>The amount of space available to store object data in a storage pool is low.</p> <ol style="list-style-type: none"> <li>1. Select <b>ILM &gt; Storage Pools</b>.</li> <li>2. Select the storage pool listed in the alert, and select <b>View details</b>.</li> <li>3. Determine where additional storage capacity is required. You can either add Storage Nodes to each site in the storage pool or add storage volumes (LUNs) to one or more existing Storage Nodes.</li> <li>4. Perform an expansion procedure to increase storage capacity.</li> </ol> <p><a href="#">Expand your grid</a></p>
Low installed node memory	<p>The amount of installed memory on a node is low.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 installation instructions for your platform:</p> <ul style="list-style-type: none"> <li>• <a href="#">Install Red Hat Enterprise Linux or CentOS</a></li> <li>• <a href="#">Install Ubuntu or Debian</a></li> <li>• <a href="#">Install VMware</a></li> </ul>

Alert name	Description and recommended actions
Low metadata storage	<p>The space available for storing object metadata is low.<b>Critical alert</b></p> <ol style="list-style-type: none"> <li>1. Stop ingesting objects.</li> <li>2. Immediately add Storage Nodes in an expansion procedure.</li> </ol> <p><b>Major alert</b></p> <p>Immediately add Storage Nodes in an expansion procedure.</p> <p><b>Minor alert</b></p> <ol style="list-style-type: none"> <li>1. Monitor the rate at which object metadata space is being used. Select <b>Nodes &gt; Storage Node &gt; Storage</b>, and view the Storage Used - Object Metadata graph.</li> <li>2. Add Storage Nodes in an expansion procedure as soon as possible.</li> </ol> <p>Once new Storage Nodes are added, the system automatically rebalances object metadata across all Storage Nodes, and the alarm clears.</p> <p><a href="#">Troubleshooting the Low metadata storage alert</a></p> <p><a href="#">Expand your grid</a></p>
Low metrics disk capacity	<p>The space available for the metrics database is low.</p> <ol style="list-style-type: none"> <li>1. Monitor this alert to see if the issue resolves on its own and the disk space becomes available again.</li> <li>2. Contact technical support if the available space continues to decrease.</li> </ol>
Low object data storage	<p>The space available for storing object data is low.Perform an expansion procedure. You can add storage volumes (LUNs) to existing Storage Nodes, or you can add new Storage Nodes.</p> <p><a href="#">Troubleshooting the Low object data storage alert</a></p> <p><a href="#">Expand your grid</a></p>

Alert name	Description and recommended actions
Low root disk capacity	<p>The space available for the root disk is low.</p> <ol style="list-style-type: none"> <li>1. Monitor this alert to see if the issue resolves on its own and the disk space becomes available again.</li> <li>2. Contact technical support if the available space continues to decrease.</li> </ol>
Low system data capacity	<p>The space available for StorageGRID system data on the /var/local file system is low.</p> <ol style="list-style-type: none"> <li>1. Monitor this alert to see if the issue resolves on its own and the disk space becomes available again.</li> <li>2. Contact technical support if the available space continues to decrease.</li> </ol>
Node network connectivity error	<p>Errors have occurred while transferring data between nodes. Network connectivity errors might clear without manual intervention. Contact technical support if the errors do not clear.</p> <p><a href="#">Troubleshooting the Network Receive Error (NRER) alarm</a></p>
Node network reception frame error	<p>A high percentage of the network frames received by a node had errors. This alert might indicate a hardware issue, such as a bad cable or a failed transceiver on either end of the Ethernet connection.</p> <ol style="list-style-type: none"> <li>1. If you are using an appliance, try replacing each SFP+ or SFP28 transceiver and cable, one at a time, to see if the alert clears.</li> <li>2. If this alert persists, contact technical support.</li> </ol>
Node not in sync with NTP server	<p>The node's time is not in sync with the network time protocol (NTP) server.</p> <ol style="list-style-type: none"> <li>1. Verify that you have specified at least four external NTP servers, each providing a Stratum 3 or better reference.</li> <li>2. Check that all NTP servers are operating normally.</li> <li>3. Verify the connections to the NTP servers. Make sure they are not blocked by a firewall.</li> </ol>

Alert name	Description and recommended actions
Node not locked with NTP server	<p>The node is not locked to a network time protocol (NTP) server.</p> <ol style="list-style-type: none"> <li>1. Verify that you have specified at least four external NTP servers, each providing a Stratum 3 or better reference.</li> <li>2. Check that all NTP servers are operating normally.</li> <li>3. Verify the connections to the NTP servers. Make sure they are not blocked by a firewall.</li> </ol>
Non appliance node network down	<p>One or more network devices are down or disconnected. This alert indicates that a network interface (eth) for a node installed on a virtual machine or Linux host is not accessible.</p> <p>Contact technical support.</p>
Objects lost	<p>One or more objects have been lost from the grid. This alert might indicate that data has been permanently lost and is not retrievable.</p> <ol style="list-style-type: none"> <li>1. Investigate this alert immediately. You might need to take action to prevent further data loss. You also might be able to restore a lost object if you take prompt action.</li> </ol> <p><a href="#">Troubleshooting lost and missing object data</a></p> <ol style="list-style-type: none"> <li>2. When the underlying problem is resolved, reset the counter: <ol style="list-style-type: none"> <li>a. Select <b>Support &gt; Tools &gt; Grid Topology</b>.</li> <li>b. For the Storage Node that raised the alert, select <b>site &gt; grid node &gt; LDR &gt; Data Store &gt; Configuration &gt; Main</b>.</li> <li>c. Select <b>Reset Lost Objects Count</b> and click <b>Apply Changes</b>.</li> </ol> </li> </ol>
Platform services unavailable	<p>Too few Storage Nodes with the RSM service are running or available at a site. Make sure that the majority of the Storage Nodes that have the RSM service at the affected site are running and in a non-error state.</p> <p>See “Troubleshooting platform services” in the instructions for administering StorageGRID.</p> <p><a href="#">Administer StorageGRID</a></p>

Alert name	Description and recommended actions
Services appliance link down on Admin Network port 1	<p>The Admin Network port 1 on the appliance is down or disconnected.</p> <ol style="list-style-type: none"> <li>1. Check the cable and physical connection to Admin Network port 1.</li> <li>2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware.</li> <li>3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select <b>Alerts &gt; Alert Rules</b>, select the rule, and click <b>Edit rule</b>. Then, uncheck the <b>Enabled</b> check box. <ul style="list-style-type: none"> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> <li>◦ <a href="#">Disabling an alert rule</a></li> </ul> </li> </ol>
Services appliance link down on Admin Network (or Client Network)	<p>The appliance interface to the Admin Network (eth1) or the Client Network (eth2) is down or disconnected.</p> <ol style="list-style-type: none"> <li>1. Check the cables, SFPs, and physical connections to the StorageGRID network.</li> <li>2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware.</li> <li>3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select <b>Alerts &gt; Alert Rules</b>, select the rule, and click <b>Edit rule</b>. Then, uncheck the <b>Enabled</b> check box. <ul style="list-style-type: none"> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> <li>◦ <a href="#">Disabling an alert rule</a></li> </ul> </li> </ol>
Services appliance link down on network port 1, 2, 3, or 4	<p>Network port 1, 2, 3, or 4 on the appliance is down or disconnected.</p> <ol style="list-style-type: none"> <li>1. Check the cables, SFPs, and physical connections to the StorageGRID network.</li> <li>2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware.</li> <li>3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select <b>Alerts &gt; Alert Rules</b>, select the rule, and click <b>Edit rule</b>. Then, uncheck the <b>Enabled</b> check box. <ul style="list-style-type: none"> <li>◦ <a href="#">SG100 &amp; SG1000 services appliances</a></li> <li>◦ <a href="#">Disabling an alert rule</a></li> </ul> </li> </ol>

Alert name	Description and recommended actions
Services appliance storage connectivity degraded	<p>One of the two SSDs in a services appliance has failed or is out of synchronization with the other. Appliance functionality is not impacted, but you should address the issue immediately. If both drives fail, the appliance will no longer function.</p> <ol style="list-style-type: none"> <li>1. From the Grid Manager, select <b>Nodes &gt; services appliance</b>, and then select the <b>Hardware</b> tab.</li> <li>2. Review the message in the <b>Storage RAID Mode</b> field.</li> <li>3. If the message shows the progress of a resynchronization operation, wait for the operation to complete and then confirm that the alert is resolved. A resynchronization message means that SSD was replaced recently or that it is being resynchronized for another reason.</li> <li>4. If the message indicates that one of the SSDs has failed, replace the failed drive as soon as possible.</li> </ol> <p>For instructions on how to replace a drive in a services appliance, see the SG100 and SG1000 appliances installation and maintenance guide.</p> <p><a href="#">SG100 &amp; SG1000 services appliances</a></p>
Storage appliance link down on Admin Network port 1	<p>The Admin Network port 1 on the appliance is down or disconnected.</p> <ol style="list-style-type: none"> <li>1. Check the cable and physical connection to Admin Network port 1.</li> <li>2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware.</li> <li>3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select <b>Alerts &gt; Alert Rules</b>, select the rule, and click <b>Edit rule</b>. Then, uncheck the <b>Enabled</b> check box. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> <li>◦ <a href="#">Disabling an alert rule</a></li> </ul> </li> </ol>

Alert name	Description and recommended actions
Storage appliance link down on Admin Network (or Client Network)	<p>The appliance interface to the Admin Network (eth1) or the Client Network (eth2) is down or disconnected.</p> <ol style="list-style-type: none"> <li>1. Check the cables, SFPs, and physical connections to the StorageGRID network.</li> <li>2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware.</li> <li>3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select <b>Alerts &gt; Alert Rules</b>, select the rule, and click <b>Edit rule</b>. Then, uncheck the <b>Enabled</b> check box. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> <li>◦ <a href="#">Disabling an alert rule</a></li> </ul> </li> </ol>
Storage appliance link down on network port 1, 2, 3, or 4	<p>Network port 1, 2, 3, or 4 on the appliance is down or disconnected.</p> <ol style="list-style-type: none"> <li>1. Check the cables, SFPs, and physical connections to the StorageGRID network.</li> <li>2. Address any connection issues. See the installation and maintenance instructions for your appliance hardware.</li> <li>3. If this port is disconnected on purpose, disable this rule. From the Grid Manager, select <b>Alerts &gt; Alert Rules</b>, select the rule, and click <b>Edit rule</b>. Then, uncheck the <b>Enabled</b> check box. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> <li>◦ <a href="#">Disabling an alert rule</a></li> </ul> </li> </ol>

Alert name	Description and recommended actions
Storage appliance storage connectivity degraded	<p data-bbox="816 157 1429 256">There is a problem with one or more connections between the compute controller and storage controller.</p> <ol data-bbox="829 294 1437 525" style="list-style-type: none"><li data-bbox="829 294 1437 357">1. Go to the appliance to check the port indicator lights.</li><li data-bbox="829 373 1437 472">2. If a port's lights are off, confirm the cable is properly connected. As needed, replace the cable.</li><li data-bbox="829 489 1437 525">3. Wait up to five minutes.</li></ol> <div data-bbox="898 567 1461 745" style="border-left: 1px solid #ccc; padding-left: 10px;"><p data-bbox="1015 577 1453 735">If a second cable needs to be replaced, do not unplug it for at least 5 minutes. Otherwise, the root volume might become read-only, which requires a hardware restart.</p></div> <ol data-bbox="829 787 1461 913" style="list-style-type: none"><li data-bbox="829 787 1461 913">4. From the Grid Manager, select <b>Nodes</b>. Then, select the Hardware tab of the node that had the problem. Verify that the alert condition has resolved.</li></ol>

Alert name	Description and recommended actions
Storage device inaccessible	<p>A storage device cannot be accessed. This alert indicates that a volume cannot be mounted or accessed because of a problem with an underlying storage device.</p> <ol style="list-style-type: none"> <li>1. Check the status of all storage devices used for the node: <ul style="list-style-type: none"> <li>◦ If the node is installed on a virtual machine or Linux host, follow the instructions for your operating system to run hardware diagnostics or perform a filesystem check. <ul style="list-style-type: none"> <li>▪ <a href="#">Install Red Hat Enterprise Linux or CentOS</a></li> <li>▪ <a href="#">Install Ubuntu or Debian</a></li> <li>▪ <a href="#">Install VMware</a></li> </ul> </li> <li>◦ If the node is installed on an SG100, SG1000 or SG6000 appliance, use the BMC.</li> <li>◦ If the node is installed on a SG5600 or SG5700 appliance, use SANtricity System Manager.</li> </ul> </li> <li>2. If necessary, replace the component. See the installation and maintenance instructions for your appliance hardware. <ul style="list-style-type: none"> <li>◦ <a href="#">SG6000 storage appliances</a></li> <li>◦ <a href="#">SG5700 storage appliances</a></li> <li>◦ <a href="#">SG5600 storage appliances</a></li> </ul> </li> </ol>
Tenant quota usage high	<p>A high percentage of tenant quota space is being used. If a tenant exceeds its quota, new ingests are rejected.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin: 10px 0;">  <p style="margin: 0;">This alert rule is disabled by default because it might generate a lot of notifications.</p> </div> <ol style="list-style-type: none"> <li>1. From the Grid Manager, select <b>Tenants</b>.</li> <li>2. Sort the table by <b>Quota Utilization</b>.</li> <li>3. Select a tenant whose quota utilization is close to 100%.</li> <li>4. Do either or both of the following: <ul style="list-style-type: none"> <li>◦ Select <b>Edit</b> to increase the storage quota for the tenant.</li> <li>◦ Notify the tenant that their quota utilization is high.</li> </ul> </li> </ol>

Alert name	Description and recommended actions
Unable to communicate with node	<p>One or more services are unresponsive, or the node cannot be reached. This alert indicates that a node is disconnected for an unknown reason. 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.</p> <p>Monitor this alert to see if the issue resolves on its own. If the issue persists:</p> <ol style="list-style-type: none"> <li>1. Determine if there is another alert affecting this node. This alert might be resolved when you resolve the other alert.</li> <li>2. Confirm that all of the services on this node are running. If a service is stopped, try starting it. See the recovery and maintenance instructions.</li> <li>3. Ensure that the host for the node is powered on. If it is not, start the host.</li> </ol> <div style="display: flex; align-items: center; margin: 10px 0;">  <div style="border-left: 1px solid #ccc; padding-left: 10px;"> <p>If more than one host is powered off, see the recovery and maintenance instructions.</p> </div> </div> <ol style="list-style-type: none"> <li>4. Determine if there is a network connectivity issue between this node and the Admin Node.</li> <li>5. If you cannot resolve the alert, contact technical support.</li> </ol> <p><a href="#">Maintain &amp; recover</a></p>
Unexpected node reboot	<p>A node rebooted unexpectedly within the last 24 hours.</p> <ol style="list-style-type: none"> <li>1. Monitor this alert. The alert will be cleared after 24 hours. However, if the node reboots unexpectedly again, this alert will be triggered again.</li> <li>2. If you cannot resolve the alert, there might be a hardware failure. Contact technical support.</li> </ol>

Alert name	Description and recommended actions
Unidentified corrupt object detected	<p>A file was found in replicated object storage that could not be identified as a replicated object.</p> <ol style="list-style-type: none"> <li>1. Determine if there are any issues with the underlying storage on a Storage Node. For example, run hardware diagnostics or perform a filesystem check.</li> <li>2. After resolving any storage issues, run foreground verification to determine if objects are missing and to replace them if possible.</li> <li>3. Monitor this alert. The alert will clear after 24 hours, but will be triggered again if the issue has not been fixed.</li> <li>4. If you cannot resolve the alert, contact technical support.</li> </ol> <p><a href="#">Running foreground verification</a></p>

## Related information

[Commonly used Prometheus metrics](#)

### Commonly used Prometheus metrics

The Prometheus service on Admin Nodes collects time series metrics from the services on all nodes. While Prometheus collects more than a thousand metrics, a relatively small number are required to monitor the most critical StorageGRID operations.

The following table lists the most commonly used Prometheus metrics and provides a mapping of each metric to the equivalent attribute (used in the alarm system).

You can refer to this list to better understand the conditions in the default alert rules or to construct the conditions for custom alert rules. For a complete list of metrics, select **Help > API Documentation**.



Metrics that include *private* in their names are intended for internal use only and are subject to change between StorageGRID releases without notice.



Prometheus metrics are retained for 31 days.

Prometheus metric	Description
alertmanager_notifications_failed_total	The total number of failed alert notifications.
node_filesystem_avail_bytes	The amount of filesystem space available to non-root users in bytes.
node_memory_MemAvailable_bytes	Memory information field MemAvailable_bytes.

Prometheus metric	Description
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 are not 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 or Swift.
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.  <a href="#">Troubleshooting lost and missing object data</a>
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.

Prometheus metric	Description
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.  <b>Note:</b> 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_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.

<b>Prometheus metric</b>	<b>Description</b>
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_load	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_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.

Prometheus metric	Description
storagegrid_storage_state_current	<p>The current state of the storage services. Attribute values are:</p> <ul style="list-style-type: none"> <li>• 10 = Offline</li> <li>• 15 = Maintenance</li> <li>• 20 = Read-only</li> <li>• 30 = Online</li> </ul>
storagegrid_storage_status	<p>The current status of the storage services. Attribute values are:</p> <ul style="list-style-type: none"> <li>• 0 = No Errors</li> <li>• 10 = In Transition</li> <li>• 20 = Insufficient Free Space</li> <li>• 30 = Volume(s) Unavailable</li> <li>• 40 = Error</li> </ul>
storagegrid_storage_utilization_metadata_bytes	<p>An estimate of the total size of replicated and erasure coded object data on the Storage Node.</p>
storagegrid_storage_utilization_metadata_allowed_bytes	<p>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.</p>
storagegrid_storage_utilization_metadata_bytes	<p>The amount of object metadata on storage volume 0, in bytes.</p>
storagegrid_storage_utilization_metadata_reserved_bytes	<p>The total space on volume 0 of each Storage Node that is actually reserved for object metadata. 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.</p>
storagegrid_storage_utilization_total_space_bytes	<p>The total amount of storage space allocated to all object stores.</p>

Prometheus metric	Description
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.

## Alarms reference (legacy system)

The following table lists all of the legacy Default alarms. If an alarm is triggered, you can look up the alarm code in this table to find the recommended actions.



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

Code	Name	Service	Recommended action
ABRL	Available Attribute Relays	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BSSM, BDDS	<p>Restore connectivity to a service (an ADC service) running an Attribute Relay Service as soon as possible. If there are no connected attribute relays, the grid node cannot report attribute values to the NMS service. Thus, the NMS service can no longer monitor the status of the service, or update attributes for the service.</p> <p>If the problem persists, contact technical support.</p>
ACMS	Available Metadata Services	BARC, BLDR, BCMN	<p>An alarm is triggered when an LDR or ARC service loses connection to a DDS service. If this occurs, ingest or retrieve transactions cannot be processed. If the unavailability of DDS services is only a brief transient issue, transactions can be delayed.</p> <p>Check and restore connections to a DDS service to clear this alarm and return the service to full functionality.</p>

Code	Name	Service	Recommended action
ACTS	Cloud Tiering Service Status	ARC	<p>Only available for Archive Nodes with a Target Type of Cloud Tiering - Simple Storage Service (S3).</p> <p>If the ACTS attribute for the Archive Node is set to Read-Only Enabled or Read-Write Disabled, you must set the attribute to Read-Write Enabled.</p> <p>If a major alarm is triggered due to an authentication failure, verify the credentials associated with destination bucket and update values, if necessary.</p> <p>If a major alarm is triggered due to any other reason, contact technical support.</p>
ADCA	ADC Status	ADC	<p>If an alarm is triggered, select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <i>site &gt; grid node &gt; ADC &gt; Overview &gt; Main</i> and <b>ADC &gt; Alarms &gt; Main</b> to determine the cause of the alarm.</p> <p>If the problem persists, contact technical support.</p>
ADCE	ADC State	ADC	<p>If the value of ADC State is Standby, continue monitoring the service and if the problem persists, contact technical support.</p> <p>If the value of ADC State is Offline, restart the service. If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
AITE	Retrieve State	BARC	<p>Only available for Archive Node's with a Target Type of Tivoli Storage Manager (TSM).</p> <p>If the value of Retrieve State is Waiting for Target, check the TSM middleware server and ensure that it is operating correctly. If the Archive Node has just been added to the StorageGRID system, ensure that the Archive Node's connection to the targeted external archival storage system is configured correctly.</p> <p>If the value of Archive Retrieve State is Offline, attempt to update the state to Online. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <b>site &gt; grid node &gt; ARC &gt; Retrieve &gt; Configuration &gt; Main</b>, select <b>Archive Retrieve State &gt; Online</b>, and click <b>Apply Changes</b>.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
AITU	Retrieve Status	BARC	<p>If the value of Retrieve Status is Target Error, check the targeted external archival storage system for errors.</p> <p>If the value of Archive Retrieve Status is Session Lost, check the targeted external archival storage system to ensure it is online and operating correctly. Check the network connection with the target.</p> <p>If the value of Archive Retrieve Status is Unknown Error, contact technical support.</p>
ALIS	Inbound Attribute Sessions	ADC	<p>If the number of inbound attribute sessions on an attribute relay grows too large, it can be an indication that the StorageGRID system has become unbalanced. Under normal conditions, attribute sessions should be evenly distributed amongst ADC services. An imbalance can lead to performance issues.</p> <p>If the problem persists, contact technical support.</p>
ALOS	Outbound Attribute Sessions	ADC	<p>The ADC service has a high number of attribute sessions, and is becoming overloaded. If this alarm is triggered, contact technical support.</p>

<b>Code</b>	<b>Name</b>	<b>Service</b>	<b>Recommended action</b>
ALUR	Unreachable Attribute Repositories	ADC	<p>Check network connectivity with the NMS service to ensure that the service can contact the attribute repository.</p> <p>If this alarm is triggered and network connectivity is good, contact technical support.</p>

Code	Name	Service	Recommended action
AMQS	Audit Messages Queued	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BDDS	<p>If audit messages cannot be immediately forwarded to an audit relay or repository, the messages are stored in a disk queue. If the disk queue becomes full, outages can occur.</p> <p>To allow you to respond in time to prevent an outage, AMQS alarms are triggered when the number of messages in the disk queue reaches the following thresholds:</p> <ul style="list-style-type: none"> <li>• Notice: More than 100,000 messages</li> <li>• Minor: At least 500,000 messages</li> <li>• Major: At least 2,000,000 messages</li> <li>• Critical: At least 5,000,000 messages</li> </ul> <p>If an AMQS alarm is triggered, check the load on the system—if there have been a significant number of transactions, the alarm should resolve itself over time. In this case, you can ignore the alarm.</p> <p>If the 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 to Error or Off. See “Changing audit message levels” in <i>Understanding audit messages</i>.</p>

Code	Name	Service	Recommended action
AOTE	Store State	BARC	<p>Only available for Archive Node's with a Target Type of Tivoli Storage Manager (TSM).</p> <p>If the value of Store State is Waiting for Target, check the external archival storage system and ensure that it is operating correctly. If the Archive Node has just been added to the StorageGRID system, ensure that the Archive Node's connection to the targeted external archival storage system is configured correctly.</p> <p>If the value of Store State is Offline, check the value of Store Status. Correct any problems before moving the Store State back to Online.</p>
AOTU	Store Status	BARC	<p>If the value of Store Status is Session Lost check that the external archival storage system is connected and online.</p> <p>If the value of Target Error, check the external archival storage system for errors.</p> <p>If the value of Store Status is Unknown Error, contact technical support.</p>

Code	Name	Service	Recommended action
APMS	Storage Multipath Connectivity	SSM	<p>If the multipath state alarm appears as “Degraded” (select <b>Support &gt; Tools &gt; Grid Topology</b>, then select <b>site &gt; grid node &gt; SSM &gt; Events</b>), do the following:</p> <ol style="list-style-type: none"> <li>1. Plug in or replace the cable that does not display any indicator lights.</li> <li>2. Wait one to five minutes.</li> </ol> <p>Do not unplug the other cable until at least five minutes after you plug in the first one. Unplugging too early can cause the root volume to become read-only, which requires that the hardware be restarted.</p> <ol style="list-style-type: none"> <li>3. Return to the <b>SSM &gt; Resources</b> page, and verify that the “Degraded” Multipath status has changed to “Nominal” in the Storage Hardware section.</li> </ol>

Code	Name	Service	Recommended action
ARCE	ARC State	ARC	<p>The ARC service has a state of Standby until all ARC components (Replication, Store, Retrieve, Target) have started. It then transitions to Online.</p> <p>If the value of ARC State does not transition from Standby to Online, check the status of the ARC components.</p> <p>If the value of ARC State is Offline, restart the service. If the problem persists, contact technical support.</p>
AROQ	Objects Queued	ARC	<p>This alarm can be triggered if the removable storage device is running slowly due to problems with the targeted external archival storage system, or if it encounters multiple read errors. Check the external archival storage system for errors, and ensure that it is operating correctly.</p> <p>In some cases, this error can occur as a result of a high rate of data requests. Monitor the number of objects queued as system activity declines.</p>

Code	Name	Service	Recommended action
ARRF	Request Failures	ARC	<p>If a retrieval from the targeted external archival storage system fails, the Archive Node retries the retrieval as the failure can be due to a transient issue. However, if the object data is corrupt or has been marked as being permanently unavailable, the retrieval does not fail. Instead, the Archive Node continuously retries the retrieval and the value for Request Failures continues to increase.</p> <p>This alarm can indicate that the storage media holding the requested data is corrupt. Check the external archival storage system to further diagnose the problem.</p> <p>If you determine that the object data is no longer in the archive, the object will have to be removed from the StorageGRID system. For more information, contact technical support.</p> <p>Once the problem that triggered this alarm is addressed, reset the failures count. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <b>site &gt; grid node &gt; ARC &gt; Retrieve &gt; Configuration &gt; Main</b>, select <b>Reset Request Failure Count</b> and click <b>Apply Changes</b>.</p>

Code	Name	Service	Recommended action
ARRV	Verification Failures	ARC	<p>To diagnose and correct this problem, contact technical support.</p> <p>Once the problem that triggered this alarm is addressed, reset the failures count. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <b>site &gt; grid node &gt; ARC &gt; Retrieve &gt; Configuration &gt; Main</b>, select <b>Reset Verification Failure Count</b> and click <b>Apply Changes</b>.</p>
ARVF	Store Failures	ARC	<p>This alarm can occur as a result of errors with the targeted external archival storage system. Check the external archival storage system for errors, and ensure that it is operating correctly.</p> <p>Once the problem that triggered this alarm is addressed, reset the failures count. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <b>site &gt; grid node &gt; ARC &gt; Retrieve &gt; Configuration &gt; Main</b>, select <b>Reset Store Failure Count</b>, and click <b>Apply Changes</b>.</p>
ASXP	Audit Shares	AMS	<p>An alarm is triggered if the value of Audit Shares is Unknown. This alarm can indicate a problem with the installation or configuration of the Admin Node.</p> <p>If the problem persists, contact technical support.</p>

<b>Code</b>	<b>Name</b>	<b>Service</b>	<b>Recommended action</b>
AUMA	AMS Status	AMS	<p>If the value of AMS Status is DB Connectivity Error, restart the grid node.</p> <p>If the problem persists, contact technical support.</p>
AUME	AMS State	AMS	<p>If the value of AMS State is Standby, continue monitoring the StorageGRID system. If the problem persists, contact technical support.</p> <p>If the value of AMS State is Offline, restart the service. If the problem persists, contact technical support.</p>
AUXS	Audit Export Status	AMS	<p>If an alarm is triggered, correct the underlying problem, and then restart the AMS service.</p> <p>If the problem persists, contact technical support.</p>
BADD	Storage Controller Failed Drive Count	SSM	<p>This alarm is triggered when one or more drives in a StorageGRID appliance has failed or is not optimal. Replace the drives as required.</p>
BASF	Available Object Identifiers	CMN	<p>When a StorageGRID system is provisioned, the CMN service is allocated a fixed number of object identifiers. This alarm is triggered when the StorageGRID system begins to exhaust its supply of object identifiers.</p> <p>To allocate more identifiers, contact technical support.</p>

Code	Name	Service	Recommended action
BASS	Identifier Block Allocation Status	CMN	<p>By default, an alarm is triggered when object identifiers cannot be allocated because ADC quorum cannot be reached.</p> <p>Identifier block allocation on the CMN service requires a quorum (50% + 1) of the ADC services to be online and connected. If quorum is unavailable, the CMN service is unable to allocate new identifier blocks until ADC quorum is re-established. If ADC quorum is lost, there is generally no immediate impact on the StorageGRID system (clients can still ingest and retrieve content), as approximately one month's supply of identifiers are cached elsewhere in the grid; however, if the condition continues, the StorageGRID system will lose the ability to ingest new content.</p> <p>If an alarm is triggered, investigate the reason for the loss of ADC quorum (for example, it can be a network or Storage Node failure) and take corrective action.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
BRDT	Compute Controller Chassis Temperature	SSM	<p>An alarm is triggered if the temperature of the compute controller in a StorageGRID appliance exceeds a nominal threshold.</p> <p>Check hardware components and environmental issues for overheated condition. If necessary, replace the component.</p>
BTOF	Offset	BADC, BLDR, BNMS, BAMS, BCLB, BCMN, BARC	<p>An alarm is triggered if the service time (seconds) differs significantly from the operating system time. Under normal conditions, the service should resynchronize itself. If the service time drifts too far from the operating system time, system operations can be affected. Confirm that the StorageGRID system's time source is correct.</p> <p>If the problem persists, contact technical support.</p>
BTSE	Clock State	BADC, BLDR, BNMS, BAMS, BCLB, BCMN, BARC	<p>An alarm is triggered if the service's time is not synchronized with the time tracked by the operating system. Under normal conditions, the service should resynchronize itself. If the time drifts too far from operating system time, system operations can be affected. Confirm that the StorageGRID system's time source is correct.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
CAHP	Java Heap Usage Percent	DDS	<p>An alarm is triggered if Java is unable to perform garbage collection at a rate that allows enough heap space for the system to properly function. An alarm might indicate a user workload that exceeds the resources available across the system for the DDS metadata store. Check the ILM Activity in the Dashboard, or select <b>Support &gt; Tools &gt; Grid Topology</b>, then select <i>site &gt; grid node &gt; DDS &gt; Resources &gt; Overview &gt; Main</i>.</p> <p>If the problem persists, contact technical support.</p>
CAIH	Number Available Ingest Destinations	CLB	This alarm is deprecated.
CAQH	Number Available Destinations	CLB	<p>This alarm clears when underlying issues of available LDR services are corrected. Ensure that the HTTP component of LDR services are online and running normally.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
CASA	Data Store Status	DDS	<p>An alarm is raised if the Cassandra metadata store becomes unavailable.</p> <p>Check the status of Cassandra:</p> <ol style="list-style-type: none"> <li>1. At the Storage Node, log in as admin and su to root using the password listed in the Passwords.txt file.</li> <li>2. Enter: <code>service cassandra status</code></li> <li>3. If Cassandra is not running, restart it: <code>service cassandra restart</code></li> </ol> <p>This alarm might also indicate that the metadata store (Cassandra database) for a Storage Node requires rebuilding.</p> <p><a href="#">Troubleshooting the Services: Status - Cassandra (SVST) alarm</a></p> <p>If the problem persists, contact technical support.</p>
CASE	Data Store State	DDS	<p>This alarm is triggered during installation or expansion to indicate a new data store is joining the grid.</p>
CCES	Incoming Sessions - Established	CLB	<p>This alarm is triggered if there are 20,000 or more HTTP sessions currently active (open) on the Gateway Node. If a client has too many connections, you might see connection failures. You should reduce the workload.</p>

<b>Code</b>	<b>Name</b>	<b>Service</b>	<b>Recommended action</b>
CCNA	Compute Hardware	SSM	This alarm is triggered if the status of the compute controller hardware in a StorageGRID appliance is Needs Attention.

Code	Name	Service	Recommended action
CDLP	Metadata Used Space (Percent)	DDS	<p>This alarm is triggered when the Metadata Effective Space (CEMS) reaches 70% full (minor alarm), 90% full (major alarm), and 100% full (critical alarm).</p> <p>If this alarm reaches the 90% threshold, a warning appears on the Dashboard in the Grid Manager. You must perform an expansion procedure to add new Storage Nodes as soon as possible. See the instructions for expanding a StorageGRID grid.</p> <p>If this alarm reaches the 100% threshold, you must stop ingesting objects and add Storage Nodes immediately. Cassandra requires a certain amount of space to perform essential operations such as compaction and repair. These operations will be impacted if object metadata uses more than 100% of the allowed space. Undesirable results can occur.</p> <p><b>Note:</b> Contact technical support if you are unable to add Storage Nodes.</p> <p>Once new Storage Nodes are added, the system automatically rebalances object metadata across all Storage Nodes, and the alarm clears.</p> <p><a href="#">Troubleshooting the Low metadata storage alert</a></p> <p><a href="#">Expand your grid</a></p>

Code	Name	Service	Recommended action
CLBA	CLB Status	CLB	<p>If an alarm is triggered, select <b>Support &gt; Tools &gt; Grid Topology</b>, then select <i>site</i> &gt; <i>grid node</i> &gt; <b>CLB &gt; Overview &gt; Main</b> and <b>CLB &gt; Alarms &gt; Main</b> to determine the cause of the alarm and to troubleshoot the problem.</p> <p>If the problem persists, contact technical support.</p>
CLBE	CLB State	CLB	<p>If the value of CLB State is Standby, continue monitoring the situation and if the problem persists, contact technical support.</p> <p>If the state is Offline and there are no known server hardware issues (for example, the server is unplugged) or scheduled downtime, restart the service. If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
CMNA	CMN Status	CMN	<p>If the value of CMN Status is Error, select <b>Support &gt; Tools &gt; Grid Topology</b>, then select <i>site</i> &gt; <b>grid node &gt; CMN &gt; Overview &gt; Main</b> and <b>CMN &gt; Alarms &gt; Main</b> to determine the cause of the error and to troubleshoot the problem.</p> <p>An alarm is triggered and the value of CMN Status is No Online CMN during a hardware refresh of the primary Admin Node when the CMNs are switched (the value of the old CMN State is Standby and the new is Online).</p> <p>If the problem persists, contact technical support.</p>
CPRC	Remaining Capacity	NMS	<p>An alarm is triggered if the remaining capacity (number of available connections that can be opened to the NMS database) falls below the configured alarm severity.</p> <p>If an alarm is triggered, contact technical support.</p>
CPSA	Compute Controller Power Supply A	SSM	<p>An alarm is triggered if there is an issue with power supply A in the compute controller for a StorageGRID appliance.</p> <p>If necessary, replace the component.</p>

<b>Code</b>	<b>Name</b>	<b>Service</b>	<b>Recommended action</b>
CPSB	Compute Controller Power Supply B	SSM	<p>An alarm is triggered if there is an issue with power supply B in the compute controller for a StorageGRID appliance.</p> <p>If necessary, replace the component.</p>
CPUT	Compute Controller CPU Temperature	SSM	<p>An alarm is triggered if the temperature of the CPU in the compute controller in a StorageGRID appliance exceeds a nominal threshold.</p> <p>If the Storage Node is a StorageGRID appliance, the StorageGRID system indicates that the controller needs attention.</p> <p>Check hardware components and environment issues for overheated condition. If necessary, replace the component.</p>
DNST	DNS Status	SSM	<p>After installation completes, a DNST alarm is triggered in the SSM service. After the DNS is configured and the new server information reaches all grid nodes, the alarm is canceled.</p>

Code	Name	Service	Recommended action
ECCD	Corrupt Fragments Detected	LDR	<p>An alarm is triggered when the background verification process detects a corrupt erasure coded fragment. If a corrupt fragment is detected, an attempt is made to rebuild the fragment. Reset the Corrupt Fragments Detected and Copies Lost attributes to zero and monitor them to see if counts go up again. If counts do go up, there may be a problem with the Storage Node's underlying storage. A copy of erasure coded object data is not considered missing until such time that the number of lost or corrupt fragments breaches the erasure code's fault tolerance; therefore, it is possible to have corrupt fragment and to still be able to retrieve the object.</p> <p>If the problem persists, contact technical support.</p>
ECST	Verification Status	LDR	<p>This alarm indicates the current status of the background verification process for erasure coded object data on this Storage Node.</p> <p>A major alarm is triggered if there is an error in the background verification process.</p>
FOPN	Open File Descriptors	BADDC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BSSM, BDDS	<p>FOPN can become large during peak activity. If it does not diminish during periods of slow activity, contact technical support.</p>

Code	Name	Service	Recommended action
HSTE	HTTP State	BLDR	See recommended actions for HSTU.
HSTU	HTTP Status	BLDR	<p>HSTE and HSTU are related to the HTTP protocol for all LDR traffic, including S3, Swift, and other internal StorageGRID traffic. An alarm indicates that one of the following situations has occurred:</p> <ul style="list-style-type: none"> <li>• The HTTP protocol has been taken offline manually.</li> <li>• The Auto-Start HTTP attribute has been disabled.</li> <li>• The LDR service is shutting down.</li> </ul> <p>The Auto-Start HTTP attribute is enabled by default. If this setting is changed, HTTP could remain offline after a restart.</p> <p>If necessary, wait for the LDR service to restart.</p> <p>Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <b>Storage Node &gt; LDR &gt; Configuration</b>. If the HTTP protocol is offline, place it online. Verify that the Auto-Start HTTP attribute is enabled.</p> <p>If the HTTP protocol remains offline, contact technical support.</p>
HTAS	Auto-Start HTTP	LDR	Specifies whether to start HTTP services automatically on start-up. This is a user-specified configuration option.

Code	Name	Service	Recommended action
IRSU	Inbound Replication Status	BLDR, BARC	An alarm indicates that inbound replication has been disabled. Confirm configuration settings: Select <b>Support &gt; Tools &gt; Grid Topology</b> . Then select <b>site &gt; grid node &gt; LDR &gt; Replication &gt; Configuration &gt; Main</b> .
LATA	Average Latency	NMS	<p>Check for connectivity issues.</p> <p>Check system activity to confirm that there is an increase in system activity. An increase in system activity will result in an increase to attribute data activity. This increased activity will result in a delay to the processing of attribute data. This can be normal system activity and will subside.</p> <p>Check for multiple alarms. An increase in average latency times can be indicated by an excessive number of triggered alarms.</p> <p>If the problem persists, contact technical support.</p>
LDRE	LDR State	LDR	<p>If the value of LDR State is Standby, continue monitoring the situation and if the problem persists, contact technical support.</p> <p>If the value of LDR State is Offline, restart the service. If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
LOST	Lost Objects	DDS, LDR	<p>Triggered when the StorageGRID system fails to retrieve a copy of the requested object from anywhere in the system. Before a LOST (Lost Objects) alarm is triggered, the system attempts to retrieve and replace a missing object from elsewhere in the system.</p> <p>Lost objects represent a loss of data. The Lost Objects attribute is incremented whenever the number of locations for an object drops to zero without the DDS service purposely purging the content to satisfy the ILM policy.</p> <p>Investigate LOST (LOST Object) alarms immediately. If the problem persists, contact technical support.</p> <p><a href="#">Troubleshooting lost and missing object data</a></p>
MCEP	Management Interface Certificate Expiry	CMN	<p>Triggered when the certificate used for accessing the management interface is about to expire.</p> <ol style="list-style-type: none"> <li>1. Go to <b>Configuration &gt; Server Certificates</b>.</li> <li>2. In the Management Interface Server Certificate section, upload a new certificate.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>

Code	Name	Service	Recommended action
MINQ	E-mail Notifications Queued	NMS	<p>Check the network connections of the servers hosting the NMS service and the external mail server. Also confirm that the email server configuration is correct.</p> <p><a href="#">Configuring email server settings for alarms (legacy system)</a></p>
MINS	E-mail Notifications Status	BNMS	<p>A minor alarm is triggered if the NMS service is unable to connect to the mail server. Check the network connections of the servers hosting the NMS service and the external mail server. Also confirm that the email server configuration is correct.</p> <p><a href="#">Configuring email server settings for alarms (legacy system)</a></p>
MISS	NMS Interface Engine Status	BNMS	<p>An alarm is triggered if the NMS interface engine on the Admin Node that gathers and generates interface content is disconnected from the system. Check Server Manager to determine if the server individual application is down.</p>
NANG	Network Auto Negotiate Setting	SSM	<p>Check the network adapter configuration. The setting must match preferences of your network routers and switches.</p> <p>An incorrect setting can have a severe impact on system performance.</p>

Code	Name	Service	Recommended action
NDUP	Network Duplex Setting	SSM	<p>Check the network adapter configuration. The setting must match preferences of your network routers and switches.</p> <p>An incorrect setting can have a severe impact on system performance.</p>
NLNK	Network Link Detect	SSM	<p>Check the network cable connections on the port and at the switch.</p> <p>Check the network router, switch, and adapter configurations.</p> <p>Restart the server.</p> <p>If the problem persists, contact technical support.</p>
NRER	Receive Errors	SSM	<p>The following can be causes of NRER alarms:</p> <ul style="list-style-type: none"> <li>• Forward error correction (FEC) mismatch</li> <li>• Switch port and NIC MTU mismatch</li> <li>• High link error rates</li> <li>• NIC ring buffer overrun</li> </ul> <p><a href="#">Troubleshooting the Network Receive Error (NRER) alarm</a></p>

<b>Code</b>	<b>Name</b>	<b>Service</b>	<b>Recommended action</b>
NRLY	Available Audit Relays	BADC, BARC, BCLB, BCMN, BLDR, BNMS, BDDS	<p>If audit relays are not connected to ADC services, audit events cannot be reported. They are queued and unavailable to users until the connection is restored.</p> <p>Restore connectivity to an ADC service as soon as possible.</p> <p>If the problem persists, contact technical support.</p>
NSCA	NMS Status	NMS	<p>If the value of NMS Status is DB Connectivity Error, restart the service. If the problem persists, contact technical support.</p>
NSCE	NMS State	NMS	<p>If the value of NMS State is Standby, continue monitoring and if the problem persists, contact technical support.</p> <p>If the value of NMS State is Offline, restart the service. If the problem persists, contact technical support.</p>
NSPD	Speed	SSM	<p>This can be caused by network connectivity or driver compatibility issues. If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
NTBR	Free Tablespace	NMS	<p>If an alarm is triggered, check how fast database usage has been changing. A sudden drop (as opposed to a gradual change over time) indicates an error condition. If the problem persists, contact technical support.</p> <p>Adjusting the alarm threshold allows you to proactively manage when additional storage needs to be allocated.</p> <p>If the available space reaches a low threshold (see alarm threshold), contact technical support to change the database allocation.</p>
NTER	Transmit Errors	SSM	<p>These errors can clear without being manually reset. If they do not clear, check network hardware. Check that the adapter hardware and driver are correctly installed and configured to work with your network routers and switches.</p> <p>When the underlying problem is resolved, reset the counter. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <b>site &gt; grid node &gt; SSM &gt; Resources &gt; Configuration &gt; Main</b>, select <b>Reset Transmit Error Count</b>, and click <b>Apply Changes</b>.</p>

<b>Code</b>	<b>Name</b>	<b>Service</b>	<b>Recommended action</b>
NTFQ	NTP Frequency Offset	SSM	If the frequency offset exceeds the configured threshold, there is likely a hardware problem with the local clock. If the problem persists, contact technical support to arrange a replacement.
NCLK	NTP Lock	SSM	If the NTP daemon is not locked to an external time source, check network connectivity to the designated external time sources, their availability, and their stability.
NTOF	NTP Time Offset	SSM	If the time offset exceeds the configured threshold, there is likely a hardware problem with the oscillator of the local clock. If the problem persists, contact technical support to arrange a replacement.
NTSJ	Chosen Time Source Jitter	SSM	<p>This value indicates the reliability and stability of the time source that NTP on the local server is using as its reference.</p> <p>If an alarm is triggered, it can be an indication that the time source's oscillator is defective, or that there is a problem with the WAN link to the time source.</p>
NTSU	NTP Status	SSM	If the value of NTP Status is Not Running, contact technical support.

Code	Name	Service	Recommended action
OPST	Overall Power Status	SSM	<p>An alarm is triggered if the power of a StorageGRID appliance deviates from the recommended operating voltage.</p> <p>Check the status of Power Supply A or B to determine which power supply is operating abnormally.</p> <p>If necessary, replace the power supply.</p>
OQRT	Objects Quarantined	LDR	<p>After the objects are automatically restored by the StorageGRID system, the quarantined objects can be removed from the quarantine directory.</p> <ol style="list-style-type: none"> <li>1. Select <b>Support &gt; Tools &gt; Grid Topology</b>.</li> <li>2. Select <b>site &gt; Storage Node &gt; LDR &gt; Verification &gt; Configuration &gt; Main</b>.</li> <li>3. Select <b>Delete Quarantined Objects</b>.</li> <li>4. Click <b>Apply Changes</b>.</li> </ol> <p>The quarantined objects are removed, and the count is reset to zero.</p>

Code	Name	Service	Recommended action
ORSU	Outbound Replication Status	BLDR, BARC	<p>An alarm indicates that outbound replication is not possible: storage is in a state where objects cannot be retrieved. An alarm is triggered if outbound replication is disabled manually. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <i>site &gt; grid node &gt; LDR &gt; Replication &gt; Configuration</i>.</p> <p>An alarm is triggered if the LDR service is unavailable for replication. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <i>site &gt; grid node &gt; LDR &gt; Storage</i>.</p>
OSLF	Shelf Status	SSM	<p>An alarm is triggered if the status of one of the components in the storage shelf for a storage appliance is degraded. Storage shelf components include the IOMs, fans, power supplies, and drive drawers. If this alarm is triggered, see the maintenance instructions for your appliance.</p>

Code	Name	Service	Recommended action
PMEM	Service Memory Usage (Percent)	BADC, BAMS, BARC, BCLB, BCMN, BLDR, BNMS, BSSM, BDDS	<p>Can have a value of Over Y% RAM, where Y represents the percentage of memory being used by the server.</p> <p>Figures under 80% are normal. Over 90% is considered a problem.</p> <p>If memory usage is high for a single service, monitor the situation and investigate.</p> <p>If the problem persists, contact technical support.</p>
PSAS	Power Supply A Status	SSM	<p>An alarm is triggered if power supply A in a StorageGRID appliance deviates from the recommended operating voltage.</p> <p>If necessary, replace power supply A.</p>
PSBS	Power Supply B Status	SSM	<p>An alarm is triggered if power supply B in a StorageGRID appliance deviates from the recommended operating voltage.</p> <p>If necessary, replace the power supply B.</p>

Code	Name	Service	Recommended action
RDTE	Tivoli Storage Manager State	BARC	<p>Only available for Archive Nodes with a Target Type of Tivoli Storage Manager (TSM).</p> <p>If the value of Tivoli Storage Manager State is Offline, check Tivoli Storage Manager Status and resolve any problems.</p> <p>Bring the component back online. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <i>site</i> &gt; <i>grid node</i> &gt; <b>ARC &gt; Target &gt; Configuration &gt; Main</b>, select <b>Tivoli Storage Manager State &gt; Online</b>, and click <b>Apply Changes</b>.</p>

Code	Name	Service	Recommended action
RDTU	Tivoli Storage Manager Status	BARC	<p>Only available for Archive Nodes with a Target Type of Tivoli Storage Manager (TSM).</p> <p>If the value of Tivoli Storage Manager Status is Configuration Error and the Archive Node has just been added to the StorageGRID system, ensure that the TSM middleware server is correctly configured.</p> <p>If the value of Tivoli Storage Manager Status is Connection Failure, or Connection Failure, Retrying, check the network configuration on the TSM middleware server, and the network connection between the TSM middleware server and the StorageGRID system.</p> <p>If the value of Tivoli Storage Manager Status is Authentication Failure, or Authentication Failure, Reconnecting, the StorageGRID system can connect to the TSM middleware server, but cannot authenticate the connection. Check that the TSM middleware server is configured with the correct user, password, and permissions, and restart the service.</p> <p>If the value of Tivoli Storage Manager Status is Session Failure, an established session has been lost unexpectedly. Check the network connection between the TSM middleware server and the StorageGRID system. Check the middleware server for 1797 errors.</p>

Code	Name	Service	Recommended action
RIRF	Inbound Replications — Failed	BLDR, BARC	<p>An Inbound Replications — Failed alarm can occur during periods of high load or temporary network disruptions. After system activity reduces, this alarm should clear. If the count of failed replications continues to increase, look for network problems and verify that the source and destination LDR and ARC services are online and available.</p> <p>To reset the count, select <b>Support &gt; Tools &gt; Grid Topology</b>, then select <i>site &gt; grid node &gt; LDR &gt; Replication &gt; Configuration &gt; Main</i>. Select <b>Reset Inbound Replication Failure Count</b>, and click <b>Apply Changes</b>.</p>
RIRQ	Inbound Replications — Queued	BLDR, BARC	<p>Alarms can occur during periods of high load or temporary network disruption. After system activity reduces, this alarm should clear. If the count for queued replications continues to increase, look for network problems and verify that the source and destination LDR and ARC services are online and available.</p>

Code	Name	Service	Recommended action
RORQ	Outbound Replications — Queued	BLDR, BARC	<p>The outbound replication queue contains object data being copied to satisfy ILM rules and objects requested by clients.</p> <p>An alarm can occur as a result of a system overload. Wait to see if the alarm clears when system activity declines. If the alarm recurs, add capacity by adding Storage Nodes.</p>
SAVP	Total Usable Space (Percent)	LDR	<p>If usable space reaches a low threshold, options include expanding the StorageGRID system or move object data to archive through an Archive Node.</p>

Code	Name	Service	Recommended action
SCAS	Status	CMN	<p>If the value of Status for the active grid task is Error, look up the grid task message. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <i>site &gt; grid node &gt; CMN &gt; Grid Tasks &gt; Overview &gt; Main</i>. The grid task message displays information about the error (for example, “check failed on node 12130011”).</p> <p>After you have investigated and corrected the problem, restart the grid task. Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <i>site &gt; grid node &gt; CMN &gt; Grid Tasks &gt; Configuration &gt; Main</i>, and select <b>Actions &gt; Run</b>.</p> <p>If the value of Status for a grid task being aborted is Error, retry aborting the grid task.</p> <p>If the problem persists, contact technical support.</p>
SCEP	Storage API Service Endpoints Certificate Expiry	CMN	<p>Triggered when the certificate used for accessing storage API endpoints is about to expire.</p> <ol style="list-style-type: none"> <li>1. Go to <b>Configuration &gt; Server Certificates</b>.</li> <li>2. In the Object Storage API Service Endpoints Server Certificate section, upload a new certificate.</li> </ol> <p><a href="#">Administer StorageGRID</a></p>

<b>Code</b>	<b>Name</b>	<b>Service</b>	<b>Recommended action</b>
SCHR	Status	CMN	<p>If the value of Status for the historical grid task is Aborted, investigate the reason and run the task again if required.</p> <p>If the problem persists, contact technical support.</p>
SCSA	Storage Controller A	SSM	<p>An alarm is triggered if there is an issue with storage controller A in a StorageGRID appliance.</p> <p>If necessary, replace the component.</p>
SCSB	Storage Controller B	SSM	<p>An alarm is triggered if there is an issue with storage controller B in a StorageGRID appliance.</p> <p>If necessary, replace the component.</p> <p>Some appliance models do not have a storage controller B.</p>
SHLH	Health	LDR	<p>If the value of Health for an object store is Error, check and correct:</p> <ul style="list-style-type: none"> <li>• problems with the volume being mounted</li> <li>• file system errors</li> </ul>

Code	Name	Service	Recommended action
SLSA	CPU Load Average	SSM	<p>The higher the value the busier the system.</p> <p>If the CPU Load Average persists at a high value, the number of transactions in the system should be investigated to determine whether this is due to heavy load at the time. View a chart of the CPU load average: Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <b>site &gt; grid node &gt; SSM &gt; Resources &gt; Reports &gt; Charts</b>.</p> <p>If the load on the system is not heavy and the problem persists, contact technical support.</p>
SMST	Log Monitor State	SSM	<p>If the value of Log Monitor State is not Connected for a persistent period of time, contact technical support.</p>

Code	Name	Service	Recommended action
SMTT	Total Events	SSM	<p>If the value of Total Events is greater than zero, check if there are known events (such as network failures) that can be the cause. Unless these errors have been cleared (that is, the count has been reset to 0), Total Events alarms can be triggered.</p> <p>When an issue is resolved, reset the counter to clear the alarm. Select <b>Nodes &gt; site &gt; grid node &gt; Events &gt; Reset event counts.</b></p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  <p>To reset event counts, you must have the Grid Topology Page Configuration permission.</p> </div> <p>If the value of Total Events is zero, or the number increases and the problem persists, contact technical support.</p>
SNST	Status	CMN	<p>An alarm indicates that there is a problem storing the grid task bundles. If the value of Status is Checkpoint Error or Quorum Not Reached, confirm that a majority of ADC services are connected to the StorageGRID system (50 percent plus one) and then wait a few minutes.</p> <p>If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
SOSS	Storage Operating System Status	SSM	<p>An alarm is triggered if SANtricity software indicates that there is a “Needs attention” issue with a component in a StorageGRID appliance.</p> <p>Select <b>Nodes</b>. Then select <b>appliance Storage Node &gt; Hardware</b>. Scroll down to view the status of each component. In SANtricity software, check other appliance components to isolate the issue.</p>
SSMA	SSM Status	SSM	<p>If the value of SSM Status is Error, select <b>Support &gt; Tools &gt; Grid Topology</b>, then select <b>site &gt; grid node &gt; SSM &gt; Overview &gt; Main</b> and <b>SSM &gt; Overview &gt; Alarms</b> to determine the cause of the alarm.</p> <p>If the problem persists, contact technical support.</p>
SSME	SSM State	SSM	<p>If the value of SSM State is Standby, continue monitoring, and if the problem persists, contact technical support.</p> <p>If the value of SSM State is Offline, restart the service. If the problem persists, contact technical support.</p>

Code	Name	Service	Recommended action
SSTS	Storage Status	BLDR	<p>If the value of Storage Status is Insufficient Usable Space, there is no more available storage on the Storage Node and data ingests are redirected to other available Storage Node. Retrieval requests can continue to be delivered from this grid node.</p> <p>Additional storage should be added. It is not impacting end user functionality, but the alarm persists until additional storage is added.</p> <p>If the value of Storage Status is Volume(s) Unavailable, a part of the storage is unavailable. Storage and retrieval from these volumes is not possible. Check the volume's Health for more information: Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <b>site &gt; grid node &gt; LDR &gt; Storage &gt; Overview &gt; Main</b>. The volume's Health is listed under Object Stores.</p> <p>If the value of Storage Status is Error, contact technical support.</p> <p><a href="#">Troubleshooting the Storage Status (SSTS) alarm</a></p>

Code	Name	Service	Recommended action
SVST	Status	SSM	<p>This alarm clears when other alarms related to a non-running service are resolved. Track the source service alarms to restore operation.</p> <p>Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <i>site</i> &gt; <b>grid node</b> &gt; <b>SSM &gt; Services &gt; Overview &gt; Main</b>. When the status of a service is shown as Not Running, its state is Administratively Down. The service's status can be listed as Not Running for the following reasons:</p> <ul style="list-style-type: none"> <li>• The service has been manually stopped (<code>/etc/init.d/&lt;service&gt; stop</code>).</li> <li>• There is an issue with the MySQL database and Server Manager shuts down the MI service.</li> <li>• A grid node has been added, but not started.</li> <li>• During installation, a grid node has not yet connected to the Admin Node.</li> </ul> <p>If a service is listed as Not Running, restart the service (<code>/etc/init.d/&lt;service&gt; restart</code>).</p> <p>This alarm might also indicate that the metadata store (Cassandra database) for a Storage Node requires rebuilding.</p> <p>If the problem persists, contact technical support.</p>

<b>Code</b>	<b>Name</b>	<b>Service</b>	<b>Recommended action</b>
TMEM	Installed Memory	SSM	Nodes running with less than 24 GiB of installed memory can lead to performance problems and system instability. The amount of memory installed on the system should be increased to at least 24 GiB.
TPOP	Pending Operations	ADC	A queue of messages can indicate that the ADC service is overloaded. Too few ADC services can be connected to the StorageGRID system. In a large deployment, the ADC service can require adding computational resources, or the system can require additional ADC services.
UMEM	Available Memory	SSM	If the available RAM gets low, determine whether this is a hardware or software issue. If it is not a hardware issue, or if available memory falls below 50 MB (the default alarm threshold), contact technical support.
VMFI	Entries Available	SSM	This is an indication that additional storage is required. Contact technical support.

Code	Name	Service	Recommended action
VMFR	Space Available	SSM	<p>If the value of Space Available gets too low (see alarm thresholds), it needs to be investigated as to whether there are log files growing out of proportion, or objects taking up too much disk space (see alarm thresholds) that need to be reduced or deleted.</p> <p>If the problem persists, contact technical support.</p>
VMST	Status	SSM	<p>An alarm is triggered if the value of Status for the mounted volume is Unknown. A value of Unknown or Offline can indicate that the volume cannot be mounted or accessed due to a problem with the underlying storage device.</p>
VPRI	Verification Priority	BLDR, BARC	<p>By default, the value of Verification Priority is Adaptive. If Verification Priority is set to High, an alarm is triggered because storage verification can slow normal operations of the service.</p>

Code	Name	Service	Recommended action
VSTU	Object Verification Status	BLDR	<p>Select <b>Support &gt; Tools &gt; Grid Topology</b>. Then select <b>site &gt; grid node &gt; LDR &gt; Storage &gt; Overview &gt; Main</b>.</p> <p>Check the operating system for any signs of block-device or file system errors.</p> <p>If the value of Object Verification Status is Unknown Error, it usually indicates a low-level file system or hardware problem (I/O error) that prevents the Storage Verification task from accessing stored content. Contact technical support.</p>
XAMS	Unreachable Audit Repositories	BADC, BARC, BCLB, BCMN, BLDR, BNMS	<p>Check network connectivity to the server hosting the Admin Node.</p> <p>If the problem persists, contact technical support.</p>

### Alarms that generate SNMP notifications (legacy system)

The following table lists the legacy alarms that generate SNMP notifications. Unlike alerts, not all alarms generate SNMP notifications. Only the alarms listed generate SNMP notifications and only at the indicated severity or higher.



While the legacy alarm system continues to be supported, the alert system offers significant benefits and is easier to use.

Code	Name	Severity
ACMS	Available Metadata Services	Critical
AITE	Retrieve State	Minor
AITU	Retrieve Status	Major
AMQS	Audit Messages Queued	Notice

<b>Code</b>	<b>Name</b>	<b>Severity</b>
AOTE	Store State	Minor
AOTU	Store Status	Major
AROQ	Objects Queued	Minor
ARRF	Request Failures	Major
ARRV	Verification Failures	Major
ARVF	Store Failures	Major
ASXP	Audit Shares	Minor
AUMA	AMS Status	Minor
AUXS	Audit Export Status	Minor
BTOF	Offset	Notice
CAHP	Java Heap Usage Percent	Major
CAQH	Number Available Destinations	Notice
CASA	Data Store Status	Major
CDLP	Metadata Used Space (Percent)	Major
CLBE	CLB State	Critical
DNST	DNS Status	Critical
ECST	Verification Status	Major
HSTE	HTTP State	Major
HTAS	Auto-Start HTTP	Notice
LOST	Lost Objects	Major
MINQ	E-mail Notifications Queued	Notice
MINS	E-mail Notifications Status	Minor

<b>Code</b>	<b>Name</b>	<b>Severity</b>
NANG	Network Auto Negotiate Setting	Notice
NDUP	Network Duplex Setting	Minor
NLNK	Network Link Detect	Minor
NRER	Receive Errors	Notice
NSPD	Speed	Notice
NTER	Transmit Errors	Notice
NTFQ	NTP Frequency Offset	Minor
NTLK	NTP Lock	Minor
NTOF	NTP Time Offset	Minor
NTSJ	Chosen Time Source Jitter	Minor
NTSU	NTP Status	Major
OPST	Overall Power Status	Major
ORSU	Outbound Replication Status	Notice
PSAS	Power Supply A Status	Major
PSBS	Power Supply B Status	Major
RDTE	Tivoli Storage Manager State	Notice
RDTU	Tivoli Storage Manager Status	Major
SAVP	Total Usable Space (Percent)	Notice
SHLH	Health	Notice
SLSA	CPU Load Average	Notice
SMTT	Total Events	Notice
SNST	Status	

Code	Name	Severity
SOSS	Storage Operating System Status	Notice
SSTS	Storage Status	Notice
SVST	Status	Notice
TMEM	Installed Memory	Minor
UMEM	Available Memory	Minor
VMST	Status	Minor
VPRI	Verification Priority	Notice
VSTU	Object Verification Status	Notice

## Log files reference

The following sections list the logs 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.

- [StorageGRID software logs](#)
- [Deployment and maintenance logs](#)
- [Logs for third-party software](#)
- [About the bycast.log](#)



The tables in this section 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 this guide.

To access these logs, you can collect log files and system data (**Support > Tools > Logs**). Or, if the primary Admin Node is unavailable or unable to reach a specific node, you can access the logs 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.

### Related information

[Collecting log files and system data](#)

## StorageGRID software logs

You can use StorageGRID logs to troubleshoot issues.

### General StorageGRID logs

File name	Notes	Found on
<code>/var/local/log/bycast.log</code>	The file <code>bycast.log</code> is the primary StorageGRID troubleshooting file. The file <code>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 <b>Support &gt; Tools &gt; Grid Topology</b> . Then select <b>Site &gt; Node &gt; SSM &gt; Events</b> .	All nodes
<code>/var/local/log/bycast-err.log</code>	The file <code>bycast.log</code> is the primary StorageGRID troubleshooting file. The file <code>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 <b>Support &gt; Tools &gt; Grid Topology</b> . Then select <b>Site &gt; Node &gt; SSM &gt; Events</b> .	All nodes
<code>/var/local/core/</code>	Contains any core dump files created if the program terminates abnormally. Possible causes include assertion failures, violations, or thread timeouts.  <b>Note:</b> The file <code>~/var/local/core/kexec_cmd</code> usually exists on appliance nodes and does not indicate an error.	All nodes

### Server Manager logs

File name	Notes	Found on
<code>/var/local/log/servermanager.log</code>	Log file for the Server Manager application running on the server.	All nodes

File name	Notes	Found on
/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

**Logs for StorageGRID services**

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/arc.errlog		Archive 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/clb.errlog	Error information for the CLB service.  <b>Note:</b> The CLB service is deprecated.	Gateway Nodes

File name	Notes	Found on
/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/cts.errlog	This log file is only created if the Target Type is <b>Cloud Tiering - Simple Storage Service (S3)</b> .	Archive 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
/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/ldr.errlog		Storage Nodes

File name	Notes	Found on
/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
/var/local/log/nginx-gw/*.log	Contains logs for the restricted admin ports on Admin Nodes and for the Load Balancer service, which provides load balancing of S3 and Swift 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

File name	Notes	Found on
/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-snmpp-firewall.*	Contain logs related to the firewall ports being managed for SNMP.	All nodes
/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

#### NMS logs

File name	Notes	Found on
/var/local/log/nms.log	<ul style="list-style-type: none"> <li>• Captures notifications from the Grid Manager and the Tenant Manager.</li> <li>• Captures events related to the operation of the NMS service, for example, alarm processing, email notifications, and configuration changes.</li> <li>• Contains XML bundle updates resulting from configuration changes made in the system.</li> <li>• Contains error messages related to the attribute downsampling done once a day.</li> <li>• Contains Java web server error messages, for example, page generation errors and HTTP Status 500 errors.</li> </ul>	Admin Nodes

File name	Notes	Found on
/var/local/log/nms.errlog	<p>Contains error messages related to MySQL database upgrades.</p> <p>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.</p>	Admin Nodes
/var/local/log/nms.request.log	Contains information about outgoing connections from the Management API to internal StorageGRID services.	Admin Nodes

### Related information

[About the bycast.log](#)

[Use S3](#)

### 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/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

## Logs for third-party software

You can use the third-party software logs to troubleshoot issues.

Category	File name	Notes	Found on
apache2 logs	<code>/var/local/log/apache2/access.log</code> <code>/var/local/log/apache2/error.log</code>  <code>/var/local/log/apache2/other_vhosts_access.log</code>	Log files for apache2.	Admin Nodes
Archiving	<code>/var/local/log/dsieurror.log</code>	Error information for TSM Client APIs.	Archive Nodes
MySQL	<code>/var/local/log/mysql.err`</code>  <code>/var/local/log/mysql1.err</code> <code>/var/local/log/mysql1-slow.log</code>	Log files generated by MySQL.  The file <code>mysql.err</code> captures database errors and events such as startups and shutdowns.  The file <code>mysql-slow.log</code> (the slow query log) captures the SQL statements that took more than 10 seconds to execute.	Admin Nodes
Operating system	<code>/var/local/log/messages</code>	This directory contains log files for the operating system. The errors contained in these logs are also displayed in the Grid Manager. Select <b>Support &gt; Tools &gt; Grid Topology</b> . Then select <b>Topology &gt; Site &gt; Node &gt; SSM &gt; Events</b> .	All nodes

Category	File name	Notes	Found on
NTP	/var/local/log/ntp.log  /var/lib/ntp/var/log/ntpstats/	The /var/local/log/ntp.log contains the log file for NTP error messages.  The /var/lib/ntp/var/log/ntpstats/ directory contains NTP timing statistics.  loopstats records loop filter statistics information.  peerstats records peer statistics information.	All nodes
Samba	/var/local/log/samba/	The Samba log directory includes a log file for each Samba process (smb, nmb, and winbind) and every client hostname/IP.	Admin Node configured to export the audit share over CIFS

### 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.

### 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.

## Related information

[Collecting 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.

This is an example of an 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.ffffff)
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. Critical messages are also displayed in the Grid Manager. Select **Support > Tools > Grid Topology**. Then select **Site > Node > SSM > Events**.

**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
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 `bycast.log`.

Error number	Error code	Meaning
001	EPERM	Operation not permitted
002	ENOENT	No such file or directory
003	ESRCH	No such process

<b>Error number</b>	<b>Error code</b>	<b>Meaning</b>
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
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

<b>Error number</b>	<b>Error code</b>	<b>Meaning</b>
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
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

<b>Error number</b>	<b>Error code</b>	<b>Meaning</b>
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
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

<b>Error number</b>	<b>Error code</b>	<b>Meaning</b>
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
078	EREMCHG	Remote address changed
079	ELIBACC	Cannot 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	Cannot 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

<b>Error number</b>	<b>Error code</b>	<b>Meaning</b>
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	Cannot assign requested address
100	ENETDOWN	Network is down
101	ENETUNREACH	Network is unreachable
102	ENETRESET	Network dropped connection because of reset
103	ECONNABORTED	Software caused connection abort
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

<b>Error number</b>	<b>Error code</b>	<b>Meaning</b>
108	ESHUTDOWN	Cannot send after transport endpoint shutdown
109	ETOOMANYREFS	Too many references: cannot 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
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


Error number	Error code	Meaning
129	EKEYREJECTED	Key was rejected by service
130	EOWNERDEAD	For robust mutexes: Owner died
131	ENOTRECOVERABLE	For robust mutexes: State not recoverable

## Troubleshoot a StorageGRID system

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

### Overview of problem determination

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



### Defining the problem

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

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

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

### Assessing the risk and impact on the system

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

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

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

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

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

### Creating 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: <ul style="list-style-type: none"> <li>When did you start the node recovery?</li> <li>When did the software upgrade complete?</li> <li>Did you interrupt the process?</li> </ul>	What happened? What did you do?	Document any relevant details about the change. For example: <ul style="list-style-type: none"> <li>Details of the network changes.</li> <li>Which hotfix was installed.</li> <li>How client workloads changed.</li> </ul> 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 configuration changes been made to the Archive Node?
- Have any other changes been made to the StorageGRID system or its environment?

### Establishing baselines

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

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

Property	Value	How to obtain
Rate of S3/Swift operations	Operations/second	Go to the Dashboard in the Grid Manager. In the Protocol Operations section, view the values for S3 rate and the Swift rate.  To see ingest and retrieval rates and counts for a specific site or node, select <b>Nodes &gt; <i>site or Storage Node</i> &gt; Objects</b> . Hover your cursor over the Ingest and Retrieve chart for S3 or Swift.
Failed S3/Swift operations	Operations	Select <b>Support &gt; Tools &gt; Grid Topology</b> . 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 <b>grid &gt; ILM</b> .  On the ILM Queue chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate a baseline value for <b>Evaluation rate</b> for your system.
ILM scan rate	Objects/second	Select <b>Nodes &gt; grid &gt; ILM</b> .  On the ILM Queue chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate a baseline value for <b>Scan rate</b> for your system.
Objects queued from client operations	Objects/second	Select <b>Nodes &gt; grid &gt; ILM</b> .  On the ILM Queue chart, find a period where the line is fairly stable. Hover your cursor over the chart to estimate a baseline value for <b>Objects queued (from client operations)</b> for your system.
Average query latency	Milliseconds	Select <b>Nodes &gt; Storage Node &gt; Objects</b> . In the Queries table, view the value for Average Latency.

## Analyzing data

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

The analysis is problem-dependent, but in general:

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

## Escalation information checklist

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

✓	Item	Notes
	Problem statement	<p>What are the problem symptoms? When did the problem start? Does it happen consistently or intermittently? If intermittently, what times has it occurred?</p> <p><a href="#">Defining the problem</a></p>
	Impact assessment	<p>What is the severity of the problem? What is the impact to the client application?</p> <ul style="list-style-type: none"> <li>• Has the client connected successfully before?</li> <li>• Can the client ingest, retrieve, and delete data?</li> </ul>
	StorageGRID System ID	<p>Select <b>Maintenance &gt; System &gt; License</b>. The StorageGRID System ID is shown as part of the current license.</p>
	Software version	<p>Click <b>Help &gt; About</b> to see the StorageGRID version.</p>
	Customization	<p>Summarize how your StorageGRID system is configured. For example, list the following:</p> <ul style="list-style-type: none"> <li>• Does the grid use storage compression, storage encryption, or compliance?</li> <li>• Does ILM make replicated or erasure coded objects? Does ILM ensure site redundancy? Do ILM rules use the Strict, Balanced, or Dual Commit ingest behaviors?</li> </ul>
	Log files and system data	<p>Collect log files and system data for your system. Select <b>Support &gt; Tools &gt; Logs</b>.</p> <p>You can collect logs for the entire grid, or for selected nodes.</p> <p>If you are collecting logs only for selected nodes, be sure to include at least one Storage Node that has the ADC service. (The first three Storage Nodes at a site include the ADC service.)</p> <p><a href="#">Collecting log files and system data</a></p>
	Baseline information	<p>Collect baseline information regarding ingest operations, retrieval operations, and storage consumption.</p> <p><a href="#">Establishing baselines</a></p>
	Timeline of recent changes	<p>Create a timeline that summarizes any recent changes to the system or its environment.</p> <p><a href="#">Creating a timeline of recent changes</a></p>

✓	Item	Notes
	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.

### Related information

[Administer StorageGRID](#)

## Troubleshooting object and storage issues

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

### Confirming object data locations

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

### What you'll need

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

### Steps

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

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

#### Object Metadata Lookup

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

Identifier

3. Click **Look Up**.

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

- System metadata, including the object ID (UUID), the object name, the name of the container, the

tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.

- Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object.
- For replicated object copies, the current storage location of each copy.
- For erasure-coded object copies, the current storage location of each fragment.
- For object copies in a Cloud Storage Pool, the location of the object, including the name of the external bucket and the object's unique identifier.
- For segmented objects and multipart objects, a list of object segments including segment identifiers and data sizes. For objects with more than 100 segments, only the first 100 segments are shown.
- All object metadata in the unprocessed, internal storage format. This raw metadata includes internal system metadata that is not guaranteed to persist from release to release.

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

### System Metadata

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

### Replicated Copies

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

### Raw Metadata

```
{
  "TYPE": "CTNT",
  "CHND": "A12E96FF-B13F-4905-9E9E-45373F6E7DA8",
  "NAME": "testobject",
  "CBID": "0x8823DE7EC7C10416",
  "PHND": "FEA0AE51-534A-11EA-9FCD-31FF00C36056",
  "PPTH": "source",
  "META": {
    "BASE": {
      "PARTS": "2",

```

## Related information

[Manage objects with ILM](#)

[Use S3](#)

[Use Swift](#)

## Object store (storage volume) failures

The underlying storage on a Storage Node is divided into object stores. These object stores are physical partitions that act as mount points for StorageGRID system's storage. Object stores are also known as storage volumes.

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

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

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

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

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

1. Select **Support > Tools > Grid Topology**.

2. Select *site* > **Storage Node** > LDR > Storage > Overview > Main.

**Overview: LDR (DC1-S1) - Storage**  
Updated: 2020-01-29 15:03:39 PST

---

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

**Utilization**

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

**Replication**

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

**Object Store Volumes**

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

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

## Related information

[Maintain & recover](#)

## Verifying object integrity

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

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

## What background verification is

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

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

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



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

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

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

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

### Alerts and alarms (legacy) related to background verification

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

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

### Changing the background verification rate

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

### What you'll need

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

### About this task

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

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

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

### Steps

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

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

Reset Missing Objects Count

**Foreground Verification**

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

**Background Verification**

Verification Rate

Reset Corrupt Objects Count

**Quarantined Objects**

Delete Quarantined Objects

Apply Changes



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

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

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

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

triggered.

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

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

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

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

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

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

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



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

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

### What foreground verification is

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

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

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

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

remaining fragments. If the missing fragment cannot be rebuilt (because too many fragments have been lost), the Corrupt Copies Detected (ECOR) attribute is incremented by one. ILM then attempts to find another copy of the object, which it can use to generate a new erasure-coded copy.

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

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

### Running foreground verification

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

### What you'll need

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

### About this task

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

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

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

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

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

If a Storage Node other than the one where foreground verification is being run goes offline, the grid task continues to run until the **% Complete** attribute reaches 99.99 percent. The **% Complete** attribute then falls

back to 50 percent and waits for the Storage Node to return to online status. When the Storage Node's state returns to online, the LDR Foreground Verification grid task continues until it completes.

### Steps

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

Reset Missing Objects Count

#### Foreground Verification

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

#### Background Verification

Verification Rate

Reset Corrupt Objects Count

Apply Changes

4. Click **Apply Changes**.

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

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

5. Monitor missing objects or missing fragments:

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

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

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

- c. Select **Storage Node > LDR > Erasure Coding**.
- d. On the Overview tab under **Verification Results**, note the value of **Missing Fragments Detected**.

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

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

6. Monitor the completion of the foreground verification grid task:
  - a. Select **Support > Tools > Grid Topology**. Then select **site > Admin Node > CMN > Grid Task > Overview > Main**.
  - b. Verify that the foreground verification grid task is progressing without errors.

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

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

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

### After you finish

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

### Related information

[Maintain & recover](#)

### Troubleshooting lost and missing object data

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

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

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

- For replicated copies, if another copy cannot be retrieved, the object is considered lost, and the alert and alarm are triggered.
- For erasure coded copies, if a copy cannot be retrieved from the expected location, the Corrupt Copies Detected (ECOR) attribute is incremented by one before an attempt is made to retrieve a copy from

another location. If no other copy is found, the alert and alarm are triggered.

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

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

### Related information

[Investigating lost objects](#)

[Resetting lost and missing object counts](#)

### Investigating lost objects

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

### What you'll need

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

### About this task

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


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

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

### Steps

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

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



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

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

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

at risk.

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

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

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

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

```

        "Location Type": "CLDI\ (Location online\)",
        "NOID\ (Node ID\)": "12448208",
        "VOLI\ (Volume ID\)": "3222345473",
        "Object File Path":
"/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila\#3udu",
        "LTIM\ (Location timestamp\)": "2020-02-
12T19:36:17.880569"
    },
    {
        "Location Type": "CLDI\ (Location online\)",
        "NOID\ (Node ID\)": "12288733",
        "VOLI\ (Volume ID\)": "3222345984",
        "Object File Path":
"/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb\#3s;L",
        "LTIM\ (Location timestamp\)": "2020-02-
12T19:36:17.934425"
    }
]
}

```

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

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

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

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

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

## Related information

[Maintain & recover](#)

[Review audit logs](#)

### Searching for and restoring potentially lost objects

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

#### What you'll need

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

#### About this task

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



Contact technical support for assistance with this procedure.

## Steps

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

For example:

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

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

For example:

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

An LLST audit message looks like this sample message.

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

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

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

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

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

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

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

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

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'
```

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

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

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

- 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 policy, which makes additional copies as specified in the policy.

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

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

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

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

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

4. If the object was successfully restored to StorageGRID, verify that new locations were created.

- Enter: `cd /proc/OBRP`
- Enter: `ObjectByUUID UUID_value`

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

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

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

```

    "*ctp(HTTP content MIME type)": "binary/octet-stream"
  },
  "BYCB(System metadata)": {
    "CSIZ(Plaintext object size)": "5242880",
    "SHSH(Supplementary Plaintext hash)": "MD5D
0xBAC2A2617C1DFF7E959A76731E6EAF5E",
    "BSIZ(Content block size)": "5252084",
    "CVER(Content block version)": "196612",
    "CTME(Object store begin timestamp)": "2020-02-
12T19:16:10.983000",
    "MTME(Object store modified timestamp)": "2020-02-
12T19:16:10.983000",
    "ITME": "1581534970983000"
  },
  "CMSM": {
    "LATM(Object last access time)": "2020-02-
12T19:16:10.983000"
  },
  "AWS3": {
    "LOCC": "us-east-1"
  }
},
"CLCO\(Locations\)": \[
  \{
    "Location Type": "CLDI\(Location online\)\"",
    "NOID\(Node ID\)": "12448208",
    "VOLI\(Volume ID\)": "3222345473",
    "Object File Path":
"/var/local/rangedb/1/p/17/11/00rH0%DkRt78Ila\#3udu",
    "LTIM\(Location timestamp\)": "2020-02-
12T19:36:17.880569"
  },
  \{
    "Location Type": "CLDI\(Location online\)\"",
    "NOID\(Node ID\)": "12288733",
    "VOLI\(Volume ID\)": "3222345984",
    "Object File Path":
"/var/local/rangedb/0/p/19/11/00rH0%DkRt78Rrb\#3s;L",
    "LTIM\(Location timestamp\)": "2020-02-
12T19:36:17.934425"
  }
]
}

```

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

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

a. Log in to the grid node:

i. Enter the following command: `ssh admin@grid_node_IP`

ii. Enter the password listed in the `Passwords.txt` file.

iii. Enter the following command to switch to root: `su -`

iv. Enter the password listed in the `Passwords.txt` file.

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

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

c. Use `grep` to extract the audit messages associated with the object to an output file. Enter: `grep uid-valueaudit_file_name > output_file_name`

For example:

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

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

For example:

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

An ORLM audit message looks like this sample message.

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

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

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

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

## Related information

[Investigating lost objects](#)

[Confirming object data locations](#)

[Resetting lost and missing object counts](#)

[Review audit logs](#)

### Resetting lost and missing object counts

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

#### What you'll need

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

#### About this task


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

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

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

#### Steps

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



4. Click **Apply Changes**.

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

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

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

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

- g. Click **Apply Changes**.

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

#### Related information

[Administer StorageGRID](#)

#### Troubleshooting the Low object data storage alert

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

#### What you'll need

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

#### About this task

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

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

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

In this condition:

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

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

## Steps

1. Select **Alerts > Current**.

The Alerts page appears.

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



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

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


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

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

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

The following values are shown:

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



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

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

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

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



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

### Related information

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

[Expand your grid](#)

[Administer StorageGRID](#)

### Troubleshooting the Storage Status (SSTS) alarm

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

#### What you'll need

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

#### About this task

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



### Storage Options Overview

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

#### Object Segmentation

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

#### Storage Watermarks

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

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

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

#### Steps

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

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

**Alarms: LDR (DC1-S3-101-195) - Storage**  
Updated: 2019-10-09 12:52:43 MDT

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

Apply Changes

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



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

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



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

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

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

### Utilization

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

### Replication

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

### Object Store Volumes

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

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


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

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



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

Attribute:	Total Usable Space	Vertical Scaling:	<input checked="" type="checkbox"/>	Start Date:	2019/10/09 12:00:00
Quick Query:	Custom Query	Raw Data:	<input type="checkbox"/>	End Date:	2019/10/09 13:10:33



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


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



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

Attribute:	Total Usable Space (Percent)	Vertical Scaling:	<input checked="" type="checkbox"/>	Start Date:	2019/10/09 12:00:00
Quick Query:	Custom Query	Raw Data:	<input type="checkbox"/>	End Date:	2019/10/09 13:10:33

Update



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

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

### Related information

[Expand your grid](#)

[Administer StorageGRID](#)

### Troubleshooting delivery of platform services messages (SMTT alarm)

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

#### About this task

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

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

For additional information about troubleshooting platform services, see the instructions for administering

StorageGRID. You might need to access the tenant from the Tenant Manager to debug a platform service error.

### Steps

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

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

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

### Related information

[Administer StorageGRID](#)

[Use a tenant account](#)

[Log files reference](#)

[Resetting event counts](#)

## Troubleshooting metadata issues

There are several tasks you can perform to help determine the source of metadata problems.

### Troubleshooting the Low metadata storage alert


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

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.

#### About this task

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



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

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

```
storagegrid_storage_utilization_metadata_bytes/storagegrid_storage_utilization_metadata_allowed_bytes
```

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

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



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

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

must add new Storage Nodes immediately.

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

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

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

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



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

## Steps

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



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

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

## Related information

[Monitoring object metadata capacity for each Storage Node](#)

[Expand your grid](#)

## Troubleshooting the Services: Status - Cassandra (SVST) alarm

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

## What you'll need

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

## About this task

If Cassandra is stopped for more than 15 days (for example, the Storage Node is powered off), Cassandra will not start when the node is brought back online. You must rebuild the Cassandra database for the affected DDS service.

You can use the Diagnostics page to obtain additional information on the current state of your grid.

## Running diagnostics



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

## Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **site > Storage Node > SSM > Services > Alarms > Main** to display alarms.

This example shows that the SVST alarm was triggered.

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

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

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

3. Try restarting Cassandra from the Storage Node:

a. Log in to the grid node:

i. Enter the following command: `ssh admin@grid_node_IP`

ii. Enter the password listed in the `Passwords.txt` file.

iii. Enter the following command to switch to root: `su -`

iv. Enter the password listed in the `Passwords.txt` file.

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

b. Enter: `/etc/init.d/cassandra status`

c. If Cassandra is not running, restart it: `/etc/init.d/cassandra restart`

4. If Cassandra does not restart, determine how long Cassandra has been down. If Cassandra has been down for longer than 15 days, you must rebuild the Cassandra database.



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

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

5. To chart Cassandra:

a. Select **Support > Tools > Grid Topology**. Then select **site > Storage Node > SSM > Services > Reports > Charts**.

b. Select **Attribute > Service: Status - Cassandra**.

c. For **Start Date**, enter a date that is at least 16 days before the current date. For **End Date**, enter the current date.

d. Click **Update**.

e. If the chart shows Cassandra as being down for more than 15 days, rebuild the Cassandra database.

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

Overview Alarms **Reports** Configuration

Charts Text


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

Attribute: Services: Status - Cassandra Vertical Scalling:

Quick Query: Last Month Update Raw Data:

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

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



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

The contents of the servermanager.log file are displayed.

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

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

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

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

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

For instructions, see “Recovering from a single Storage Node down more than 15 days” in the recovery and maintenance instructions.

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

### Related information

[Maintain & recover](#)

### Troubleshooting Cassandra Out of Memory errors (SMTT alarm)

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

#### About this task

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

#### Steps

1. To view the event, select **Nodes** > *grid node* > **Events**.
2. Verify that the Cassandra Heap Out Of Memory Errors count is 1 or greater.

You can use the Diagnostics page to obtain additional information on the current state of your grid.

#### Running diagnostics

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

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

5. Once the issue is resolved, click **Reset event counts**.



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

### Related information

[Resetting event counts](#)

### Troubleshooting certificate errors

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


#### About this task

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

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

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

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



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

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

### Steps

1. If you received an alert notification about an expired certificate, access the certificate details:
  - For a server certificate, select **Configuration > Network Settings > Server Certificates**.
  - For a client certificate, select **Configuration > Access Control > Client Certificates**.

## 2. Check the validity period of the certificate.

Some web browsers and S3 or Swift clients do not accept certificates with a validity period greater than 398 days.

## 3. 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 in the instructions for administering StorageGRID.
- For a client certificate, see the steps for configuring a client certificate in the instructions for administering StorageGRID.

## 4. 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 > Network Settings > Server Certificates** 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.

### Related information

[Administer StorageGRID](#)

## Troubleshooting Admin Node and user interface issues

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

### Troubleshooting sign-on errors

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

### What you'll need

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

### About this task

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

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

Support.

- Unable to communicate with server. Reloading page...

## Steps

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

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

2. If your StorageGRID system has more than one Admin Node, try signing in to the Grid Manager from another Admin Node.
  - If you are able to sign in, you can use the **Dashboard**, **Nodes**, **Alerts**, and **Support** options to help determine the cause of the error.
  - If you have only one Admin Node or you still cannot sign in, go to the next step.
3. Determine if the node's hardware is offline.
4. If single sign-on (SSO) is enabled for your StorageGRID system, refer to the steps for configuring single sign-on, in the instructions for administering StorageGRID.

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



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

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

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

- If the local user can sign in:
    - i. Review any displayed alarms.
    - ii. Select **Configuration > Identity Federation**.
    - iii. Click **Test Connection** to validate your connection settings for the LDAP server.
    - iv. If the test fails, resolve any configuration errors.
  - If the local user cannot sign in and you are confident that the credentials are correct, go to the next step.
6. Use Secure Shell (ssh) to log in to the Admin Node:
    - a. Enter the following command: `ssh admin@Admin_Node_IP`
    - b. Enter the password listed in the `Passwords.txt` file.
    - c. Enter the following command to switch to root: `su -`
    - d. Enter the password listed in the `Passwords.txt` file.

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

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

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

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

```

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

```

8. Confirm that the Apache web server is running: # `service apache2 status`

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

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

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

10. Review the following logs:

- `/var/local/log/bycast.log`
- `/var/local/log/bycast-err.log`
- `/var/local/log/nms.log`

- `**/*commands.txt`

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

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

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

Admin Nodes use the ADC service during the authentication process.

12. From the Admin Node, log in to each of the ADC Storage Nodes, using the IP addresses you identified.
  - 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 `#`.

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.

## Related information

[Administer StorageGRID](#)

[Log files reference](#)

## Troubleshooting user interface issues

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

### Web interface does not respond as expected

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

If you experience issues with the web interface:

- Make sure you are using a supported browser.



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

- Clear your web browser cache.

Clearing the cache removes outdated resources used by the previous version of StorageGRID software, and permits the user interface to operate correctly again. For instructions, see the documentation for your web browser.

### Related information

[Web browser requirements](#)

[Administer StorageGRID](#)

### Checking the status of an unavailable Admin Node

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

#### What you'll need

You must have specific access permissions.

#### Steps

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



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

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

Operating System: Linux  
3.16.0-4-amd64

### Services

Service	Version	Status	Threads	Load	Memory
Audit Management System (AMS)	10.4.0-20170113.2207.3ec2cd0	Running	52	0.043 %	35.7 MB
CIFS Filesharing (nmbd)	2:4.2.14+dfsg-0+deb8u2	Running	1	0 %	5.5 MB
CIFS Filesharing (smbd)	2:4.2.14+dfsg-0+deb8u2	Running	1	0 %	14.5 MB
CIFS Filesharing (winbindd)	2:4.2.14+dfsg-0+deb8u2	Not Running	0	0 %	0 B
Configuration Management Node (CMN)	10.4.0-20170113.2207.3ec2cd0	Running	52	0.055 %	41.3 MB
Database Engine	5.5.53-0+deb8u1	Running	47	0.354 %	1.33 GB
Grid Deployment Utility Server	10.4.0-20170112.2125.c4253bb	Running	3	0 %	32.8 MB
Management Application Program Interface (mgmt-api)	10.4.0-20170113.2136.07c4997	Not Running	0	0 %	0 B
NFS Filesharing	10.4.0-20161224.0333.803cd91	Not Running	0	0 %	0 B
NMS Data Cleanup	10.4.0-20161224.0333.803cd91	Running	22	0.008 %	52.4 MB
NMS Data Downsampler 1	10.4.0-20161224.0333.803cd91	Running	22	0.049 %	195 MB
NMS Data Downsampler 2	10.4.0-20161224.0333.803cd91	Running	22	0.009 %	157 MB
NMS Processing Engine	10.4.0-20161224.0333.803cd91	Running	40	0.132 %	200 MB

- Determine if alarms have been triggered.
- Take the appropriate actions to resolve the issue.

#### Related information

[Administer StorageGRID](#)

### Troubleshooting 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.

#### Troubleshooting “422: Unprocessable Entity” errors

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

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

Error message	Root cause and corrective action
<pre>422: Unprocessable Entity  Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration. Unable to authenticate, please verify your username and password: LDAP Result Code 8 "Strong Auth Required": 00002028: LdapErr: DSID-0C090256, comment: The server requires binds to turn on integrity checking if SSL\TLS are not already active on the connection, data 0, v3839</pre>	<p>This message might occur if you select the <b>Do not use TLS</b> option for Transport Layer Security (TLS) when configuring identity federation using Windows Active Directory (AD).</p> <p>Using the <b>Do not use TLS</b> option is not supported for use with AD servers that enforce LDAP signing. You must select either the <b>Use STARTTLS</b> option or the <b>Use LDAPS</b> option for TLS.</p>
<pre>422: Unprocessable Entity  Validation failed. Please check the values you entered for errors. Test connection failed. Please verify your configuration. Unable to begin TLS, verify your certificate and TLS configuration: LDAP Result Code 200 "Network Error": TLS handshake failed (EOF)</pre>	<p>This message appears if you try to use an unsupported cipher to make a Transport Layer Security (TLS) connection from StorageGRID to an external system used for identify federation or Cloud Storage Pools.</p> <p>Check the ciphers that are offered by the external system. The system must use one of the ciphers supported by StorageGRID for outgoing TLS connections, as shown in the instructions for administering StorageGRID.</p>

#### Related information

[Administer StorageGRID](#)

#### Troubleshooting the Grid Network MTU mismatch alert

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

#### About this task

The differences in MTU settings could indicate that some, but not all, eth0 networks are configured for jumbo

frames. An MTU size mismatch of greater than 1000 might cause network performance problems.

## Steps

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{interface='eth0'}`
2. Modify the MTU settings as necessary to ensure they are the same for the Grid Network interface (eth0) on all nodes.
  - For appliance nodes, see the installation and maintenance instructions for your appliance.
  - For Linux- and VMware-based nodes, use the following command: `/usr/sbin/change-mtu.py [-h] [-n node] mtu network [network...]`

**Example:** `change-mtu.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-mtu.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
<code>mtu</code>	The MTU to set. Must be in the range 1280 to 9216.
<code>network</code>	The networks to apply the MTU to. Include one or more of the following network types: <ul style="list-style-type: none"><li>• grid</li><li>• admin</li><li>• client</li></ul>

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

## Related information

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

## Troubleshooting the Network Receive Error (NRER) alarm

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

### About this task

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

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

### Steps

1. Follow the troubleshooting steps for all potential causes of the NRER alarm given your network configuration.

- If the error is caused by FEC mismatch, perform the following steps:

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

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

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

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

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

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



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



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

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



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

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

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

2. After you resolve the underlying problem, reset the error counter.
  - a. Select **Support > Tools > Grid Topology**.
  - b. Select **site > grid node > SSM > Resources > Configuration > Main**.
  - c. Select **Reset Receive Error Count** and click **Apply Changes**.

#### Related information

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

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

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[SG100 & SG1000 services appliances](#)

#### Troubleshooting time synchronization errors

You might see issues with time synchronization in your grid.

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



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

## Related information

[Maintain & recover](#)

### Linux: Network connectivity issues

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

#### MAC address cloning

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



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

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

#### Promiscuous mode

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

## Related information

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

[Install Ubuntu or Debian](#)

### Linux: Node status is “orphaned”

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

#### About this task

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

- Check logs for errors and messages.
- Attempt to start the node again.
- If necessary, use Docker 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 Docker container and any controlling storagegrid-node processes:

```
sudo docker stop --time secondscontainer-name
```

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

```
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: Troubleshooting IPv6 support

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

### About this task

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

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

## DC1-S1 (Storage Node)

Overview

Hardware

Network

Storage

Objects

ILM

Events

### Node Information

<b>Name</b>	DC1-S1
<b>Type</b>	Storage Node
<b>Software Version</b>	11.1.0 (build 20180606.2152.b3bbe9d)
<b>IP Addresses</b>	10.96.106.102 <a href="#">Show less</a> 

Interface	IP Address
eth0	10.96.106.102
eth0	fe80::250:56ff:fea7:5c83

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

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

### Steps

1. Log in to the host as root or using an account with sudo permission.
2. Run the following command: `sysctl net.ipv6.conf.all.disable_ipv6`

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

The result should be 0.

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



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

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

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

The result should be 1.

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



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

5. Exit the container: `exit`

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

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

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

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

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

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

8. Run these commands to restart the StorageGRID container:

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

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

## Review audit logs

Learn the StorageGRID system audit logs and see a list of all audit messages.

- [Audit message overview](#)
- [Audit log file and message formats](#)
- [Audit messages and the object lifecycle](#)
- [Audit messages](#)

### Audit message overview

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.

You are assumed to have a sound understanding of the nature of audited activities within the StorageGRID system. To use the text log file, you must have access to the configured audit share on the Admin Node.

### Related information

[Administer StorageGRID](#)

### 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/audit/export` 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/audit/export` 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 and Swift client requests are not 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 and Swift client requests to other nodes are not 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.

- If the `/var/local/` directory used by a Storage Node becomes 85% full, the node will start refusing S3 and Swift 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 "[Changing audit message levels.](#)"

### 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.

### Changing audit message levels

You can adjust audit levels to increase or decrease the number of audit messages recorded in the audit log for each audit message category.

#### What you'll need

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

#### About this task

The audit messages recorded in the audit log are filtered based on the settings on the **Configuration > Monitoring > Audit** page.

You can set a different audit level for each of the following categories of messages:

- **System:** By default, this level is set to Normal.
- **Storage:** By default, this level is set to Error.
- **Management:** By default, this level is set to Normal.
- **Client Reads:** By default, this level is set to Normal.
- **Client Writes:** By default, this level is set to Normal.



These defaults apply if you initially installed StorageGRID using version 10.3 or later. If you have upgraded from 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**.

#### Audit

##### Audit Levels

System	<input type="text" value="Normal"/>
Storage	<input type="text" value="Error"/>
Management	<input type="text" value="Normal"/>
Client Reads	<input type="text" value="Normal"/>
Client Writes	<input type="text" value="Normal"/>

##### Audit Protocol Headers

Header Name 1	<input type="text" value="X-Forwarded-For"/>	<b>×</b>
Header Name 2	<input type="text" value="x-amz-*"/>	<b>+ ×</b>

**Save**

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" (SUCCS).
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.

3. Under **Audit Protocol Headers**, enter the name of the HTTP request headers to be included in Client Read and Client Write audit messages. Use an asterisk (\*) as a wildcard, or use the escape sequence (\\*) as a literal asterisk. Click the plus sign to create a list of header name fields.



Audit protocol headers apply to S3 and Swift requests only.

When such 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**.

4. Click **Save**.

#### Related information

[System audit messages](#)

[Object storage audit messages](#)

[Management audit message](#)

[Client read audit messages](#)

[Administer StorageGRID](#)

#### Accessing the audit log file

The audit share contains the active `audit.log` file and any compressed audit log files. For easy access to audit logs, you can configure client access to audit shares for both NFS and CIFS (deprecated). You can also access audit log files directly from the command line of the Admin Node.

#### What you'll need

- You must 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.
2. Go to the directory containing the audit log files:

```
cd /var/local/audit/export
```
3. View the current or a saved audit log file, as required.

#### Related information

[Administer StorageGRID](#)

## Audit log file rotation

Audit logs files are saved to an Admin Node's `/var/local/audit/export` directory. The active audit log files are named `audit.log`.

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/audit/export` 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 and message formats

You can use audit logs to gather information about your system and troubleshoot issues. You should understand the format of the audit log file and the general format used for audit messages.

### 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][SAIP(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]  
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCkNB8M3MTWnt-PhoTDwB9Jok7PtyLkQmA="][SUSR(CSTR):"urn:sgws:identity::17530064241597054718:root"]  
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"bucket1"][AVER(UI32):10][ATIM(UI64):1565203410247711]  
[ATYP(FC32):SPUT][ANID(UI32):12454421][AMID(FC32):S3RQ][ATID(UI64):7074142142472611085]]
```

2019-08-07T18:43:30.783597

```
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991696][TIME(UI64):120713][SAIP(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]  
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCkNB8M3MTWnt-PhoTDwB9Jok7PtyLkQmA="][SUSR(CSTR):"urn:sgws:identity::17530064241597054718:root"]  
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"bucket1"][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(FC32):S3RQ][ATID(UI64):8439606722108456022]]
```

2019-08-07T18:43:30.784558

```
[AUDT:[RSLT(FC32):SUCS][CNID(UI64):1565149504991693][TIME(UI64):121666][SAIP(IPAD):"10.224.2.255"][S3AI(CSTR):"17530064241597054718"]  
[SACC(CSTR):"s3tenant"][S3AK(CSTR):"SGKH9100SCkNB8M3MTWnt-PhoTDwB9Jok7PtyLkQmA="][SUSR(CSTR):"urn:sgws:identity::17530064241597054718:root"]  
[SBAI(CSTR):"17530064241597054718"][SBAC(CSTR):"s3tenant"][S3BK(CSTR):"bucket1"][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(FC32):S3RQ][ATID(UI64):13489590586043706682]]
```

In their default format, the audit messages in the audit log files are not 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.

## Related information

[Using the audit-explain tool](#)

[Using the audit-sum tool](#)

## Using the audit-explain tool

You can use the `audit-explain` tool to translate the audit messages in the audit log into an easy-to-read format.

### What you'll need

- You must 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 process plain or compressed audit logs. For example:

```
audit-explain audit.log
```

```
audit-explain 2019-08-12.txt.gz
```

The `audit-explain` tool can also process multiple files at once. For example:

```
audit-explain audit.log 2019-08-12.txt.gz 2019-08-13.txt.gz
```

```
audit-explain /var/local/audit/export/*
```

Finally, the `audit-explain` tool can 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
```

Since 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:
  - a. Enter the following command: `ssh admin@primary_Admin_Node_IP`
  - b. Enter the password listed in the `Passwords.txt` file.
2. Enter the following command, where `/var/local/audit/export/audit.log` represents the name and the location of the file or files you want to analyze:

```
$ audit-explain /var/local/audit/export/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 are not shown by default. If you want to see the timestamps, use the `timestamp` (`-t`) option.

## Related information

[SPUT: S3 PUT](#)

## Using the `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.

### What you'll need

- You must 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:

Code	Description	Refer to
ARCT	Archive Retrieve from Cloud-Tier	<a href="#">ARCT: Archive Retrieve from Cloud-Tier</a>
ASCT	Archive Store Cloud-Tier	<a href="#">ASCT: Archive Store Cloud-Tier</a>
IDEL	ILM Initiated Delete: Logs when ILM starts the process of deleting an object.	<a href="#">IDEL: ILM Initiated Delete</a>
SDEL	S3 DELETE: Logs a successful transaction to delete an object or bucket.	<a href="#">SDEL: S3 DELETE</a>

Code	Description	Refer to
SGET	S3 GET: Logs a successful transaction to retrieve an object or list the objects in a bucket.	<a href="#">SGET: S3 GET</a>
SHEA	S3 HEAD: Logs a successful transaction to check for the existence of an object or bucket.	<a href="#">SHEA: S3 HEAD</a>
SPUT	S3 PUT: Logs a successful transaction to create a new object or bucket.	<a href="#">SPUT: S3 PUT</a>
WDEL	Swift DELETE: Logs a successful transaction to delete an object or container.	<a href="#">WDEL: Swift DELETE</a>
WGET	Swift GET: Logs a successful transaction to retrieve an object or list the objects in a container.	<a href="#">WGET: Swift GET</a>
WHEA	Swift HEAD: Logs a successful transaction to check for the existence of an object or container.	<a href="#">WHEA: Swift HEAD</a>
WPUT	Swift PUT: Logs a successful transaction to create a new object or container.	<a href="#">WPUT: Swift PUT</a>

The `audit-sum` tool can process plain or compressed audit logs. For example:

```
audit-sum audit.log
```

```
audit-sum 2019-08-12.txt.gz
```

The `audit-sum` tool can also process multiple files at once. For example:

```
audit-sum audit.log 2019-08-12.txt.gz 2019-08-13.txt.gz
```

```
audit-sum /var/local/audit/export/*
```

Finally, the `audit-sum` tool can also 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.
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/audit/export/audit.log` represents the name and the location of the file or files you want to analyze:

```
$ audit-sum /var/local/audit/export/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	count	min(sec)	max(sec)
average(sec)			
=====	=====	=====	=====
=====			
SPUT.bucket	1	0.125	0.125
0.125			
SPUT.object	12	0.025	1.019
0.236			

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	count	min(sec)	max(sec)
average(sec)			
=====	=====	=====	=====
=====			
SPUT.cho-non-versioning	71943	0.046	1770.563
1.571			
SPUT.cho-versioning	54277	0.047	1736.633
1.415			
SPUT.cho-west-region	80615	0.040	55.557
1.329			
SPUT.ldt002	1564563	0.011	51.569
0.361			

- 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

**Related information**

[Using the audit-explain tool](#)

**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 are not 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 " ".

### Related information

[Using the audit-explain tool](#)

[Using the audit-sum tool](#)

[Audit messages](#)

[Common elements in audit messages](#)

[Data types](#)

[Audit message examples](#)

### 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.

Type	Description
CSTR	<p>A variable-length array of UTF-8 characters. Characters can be escaped with the following conventions:</p> <ul style="list-style-type: none"> <li>• Backslash is \.</li> <li>• Carriage return is \r.</li> <li>• Double quotes is \".</li> <li>• Line feed (new line) is \n.</li> <li>• Characters can be replaced by their hexadecimal equivalents (in the format \xHH, where HH is the hexadecimal value representing the character).</li> </ul>

### 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_Nc8S01H6w3w0nCOFCGgk_E6dYzKlumRsKJA=="]
[SUSR(CSTR):"urn:sgws:identity::60025621595611246499:root"]
[SBAI(CSTR):"60025621595611246499"] [SBAC(CSTR):"account"] [S3BK(CSTR):"bucket"]
[S3KY(CSTR):"object"] [CBID(UI64):0xCC128B9B9E428347]
[UID(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.

### Related information

[Common elements in audit messages](#)

[Audit messages](#)

### 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 cannot 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).  <b>Note:</b> 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.  <b>Note:</b> 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.
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, NONE is used rather than SUCS 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 a sample 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):"bc644d381a87d6cc216adcd963fb6f95dd25a38aa2cb8c9a358e8c5087a6af5f"][S3AK(CSTR):"UJXDKKQOXB7YARDS71Q2"][S3BK(CSTR):"s3small11"][S3KY(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):1579224144102530435]]
```

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.

### Related information

[SPUT: S3 PUT](#)

[Common elements in audit messages](#)

## Audit messages and the object lifecycle

Audit messages are generated each time an object is ingested, retrieved, or deleted. You can identify these transactions in the audit log by locating API-specific (S3 or Swift) audit messages.

Audit messages are linked through identifiers specific to each protocol.

Protocol	Code
Linking S3 operations	S3BK (S3 Bucket) and/or S3KY (S3 Key)
Linking Swift operations	WCON (Swift Container) and/or WOBJ (Swift Object)
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.

### Information lifecycle management policy configuration

With the default ILM policy (Baseline 2 Copy), object data is copied once for a total of two copies. If the ILM policy requires more than two copies, there will be an additional set of CBRE, CBSE, and SCMT messages for each extra copy. For more information about ILM policies, see information about managing objects with information lifecycle management.

## Archive Nodes

The series of audit messages generated when an Archive Node sends object data to an external archival storage system is similar to that for Storage Nodes except that there is no SCMT (Store Object Commit) message, and the ATCE (Archive Object Store Begin) and ASCE (Archive Object Store End) messages are generated for each archived copy of object data.

The series of audit messages generated when an Archive Node retrieves object data from an external archival storage system is similar to that for Storage Nodes except that the ARCB (Archive Object Retrieve Begin) and ARCE (Archive Object Retrieve End) messages are generated for each retrieved copy of object data.

The series of audit messages generated when an Archive Node deletes object data from an external archival storage system is similar to that for Storage Nodes except that there is no SREM (Object Store Remove) message, and there is an AREM (Archive Object Remove) message for each delete request.

## Related information

[Manage objects with ILM](#)

## Object ingest transactions

You can identify client ingest transactions in the audit log by locating API-specific (S3 or Swift) 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	<a href="#">SPUT: S3 PUT</a>
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	<a href="#">ORLM: Object Rules Met</a>

### Swift ingest audit messages

Code	Name	Description	Trace	See
WPUT	Swift PUT transaction	A Swift PUT ingest transaction has successfully completed.	CBID, WCON, WOBJ	<a href="#">WPUT: Swift PUT</a>
ORLM	Object Rules Met	The ILM policy has been satisfied for this object.	CBID	<a href="#">ORLM: Object Rules Met</a>

## 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 stock ILM rule, Make 2 Copies.



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"]<strong
class="S3KY(CSTR):"testobject-0-
3"">[CBID(UI64):0x8EF52DF8025E63A8]</strong>[CSIZ(UI64):30720][AVER(UI32):
10]<strong
class="ATIM(UI64):150032627859669">[ATYP(FC32):SPUT]</strong>[ANID(UI32):1
2086324][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:
<strong>[CBID(UI64):0x50C4F7AC2BC8EDF7]</strong> [RULE(CSTR):"Make 2
Copies"][STAT(FC32):DONE][CSIZ(UI64):0][UUID(CSTR):"0B344E18-98ED-4F22-
A6C8-A93ED68F8D3F"]<strong class="LOCS(CSTR):*"CLDI 12828634
2148730112">[RSLT(FC32):SUCS][AVER(UI32):10] [ATYP(FC32):ORLM]</strong>
[ATIM(UI64):1563398230669][ATID(UI64):15494889725796157557][ANID(UI32):131
00453][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][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 API-specific (S3 and Swift) 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	<a href="#">SDEL: S3 DELETE</a>

### 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	<a href="#">WDEL: Swift DELETE</a>

### 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 DELETE Object 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"] <strong>[S3BK(CSTR):"example"][S3KY(CSTR):"testobject-0-
7"]</strong>[CBID(UI64):0x339F21C5A6964D89]</strong>
[CSIZ(UI64):30720][AVER(UI32):10][ATIM(UI64):150032627859669]
<strong>[ATYP(FC32):SDEL]</strong>[ANID(UI32):12086324][AMID(FC32):S3RQ][A
TID(UI64):4727861330952970593]]
```

### Object retrieve transactions

You can identify object retrieve transactions in the audit log by locating API-specific (S3 and Swift) 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	<a href="#">SGET: S3 GET</a>

#### Swift retrieval audit messages

Code	Name	Description	Trace	See
WGET	Swift GET	Request made to retrieve an object from a container.	CBID, WCON, WOBJ	<a href="#">WGET: Swift GET</a>

### 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 GET Object 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] [CSIZ (UI
64) :12] [AVER (UI32) :10] [ATIM (UI64) :1505947988782605] [ATYP (FC32) :SGET] [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 GET Object request for an object stored in a bucket that they do not 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"]
<strong>[S3AI (CSTR) : "17915054115450519830"] [SACC (CSTR) : "s3-account-
b"]</strong>[S3AK (CSTR) : "SGKHpoblWlP_kBkqSCbTi754Ls8lBUog67I2LlSiUg==" ]<st
rong
class="SUSR (CSTR) : "urn:sgws:identity::17915054115450519830:root"">[SBAI (CS
TR) : "43979298178977966408"] [SBAC (CSTR) : "s3-account-
a"]</strong>[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]]
```

### 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	<a href="#">SUPD: S3 Metadata Updated</a>

### 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**).

```
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]]
```

### Related information

[Changing audit message levels](#)

## Audit messages

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 listing 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 sample message:

```
2014-07-17T03:50:47.484627
\[AUDT:[RSLT(FC32):VRGN][AVER(UI32):10][ATIM(UI64):1405569047484627][<strong>ATYP\ (FC32\):SYSU</strong>][ANID(UI32):11627225][AMID(FC32):ARNI][ATID(UI64):9445736326500603516]]
```

### Related information

[Audit messages](#)

[Changing audit message levels](#)

### Audit message categories

You should be familiar with the various categories within which audit messages are grouped. These groups are organized based on the class of activity that the message represents.

#### System audit messages

You should be familiar with audit messages belonging to the system audit category. These are events related to the auditing system itself, grid node states, system-wide task activity (grid tasks), and service backup operations, so that you can address potential issues.

Code	Message title and description	See
ECOC	Corrupt Erasure Coded Data Fragment: Indicates that a corrupt erasure coded data fragment has been detected.	<a href="#">ECOC: Corrupt Erasure Coded Data Fragment</a>
ETAF	Security Authentication Failed: A connection attempt using Transport Layer Security (TLS) failed.	<a href="#">ETAF: Security Authentication Failed</a>

<b>Code</b>	<b>Message title and description</b>	<b>See</b>
GNRG	GNDS Registration: A service updated or registered information about itself in the StorageGRID system.	<a href="#">GNRG: GNDS Registration</a>
GNUR	GNDS Unregistration: A service has unregistered itself from the StorageGRID system.	<a href="#">GNUR: GNDS Unregistration</a>
GTED	Grid Task Ended: The CMN service finished processing the grid task.	<a href="#">GTED: Grid Task Ended</a>
GTST	Grid Task Started: The CMN service started to process the grid task.	<a href="#">GTST: Grid Task Started</a>
GTSU	Grid Task Submitted: A grid task was submitted to the CMN service.	<a href="#">GTSU: Grid Task Submitted</a>
IDEL	ILM Initiated Delete: This audit message is generated when ILM starts the process of deleting an object.	<a href="#">IDEL: ILM Initiated Delete</a>
LKCU	Overwritten Object Cleanup. This audit message is generated when an overwritten object is automatically removed to free up storage space.	<a href="#">LKCU: Overwritten Object Cleanup</a>
LLST	Location Lost: This audit message is generated when a location is lost.	<a href="#">LLST: Location Lost</a>
OLST	Object Lost: A requested object cannot be located within the StorageGRID system.	<a href="#">OLST: System Detected Lost Object</a>
ORLM	Object Rules Met: Object data is stored as specified by the ILM rules.	<a href="#">ORLM: Object Rules Met</a>
SADD	Security Audit Disable: Audit message logging was turned off.	<a href="#">SADD: Security Audit Disable</a>

Code	Message title and description	See
SADE	Security Audit Enable: Audit message logging has been restored.	<a href="#">SADE: Security Audit Enable</a>
SVRF	Object Store Verify Fail: A content block failed verification checks.	<a href="#">SVRF: Object Store Verify Fail</a>
SVRU	Object Store Verify Unknown: Unexpected object data detected in the object store.	<a href="#">SVRU: Object Store Verify Unknown</a>
SYSD	Node Stop: A shutdown was requested.	<a href="#">SYSD: Node Stop</a>
SYST	Node Stopping: A service initiated a graceful stop.	<a href="#">SYST: Node Stopping</a>
SYSU	Node Start: A service started; the nature of the previous shutdown is indicated in the message.	<a href="#">SYSU: Node Start</a>
VLST	User Initiated Volume Lost: The <code>/proc/CMSI/Volume_Lost</code> command was run.	<a href="#">VLST: User Initiated Volume Lost</a>

### Related information

[LKCU: Overwritten Object Cleanup](#)

### Object storage audit messages

You should be familiar with audit messages belonging to the object storage audit category. These are 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.

Code	Description	See
APCT	Archive Purge from Cloud-Tier: Archived object data is deleted from an external archival storage system, which connects to the StorageGRID through the S3 API.	<a href="#">APCT: Archive Purge from Cloud-Tier</a>
ARCB	Archive Object Retrieve Begin: The ARC service begins the retrieval of object data from the external archival storage system.	<a href="#">ARCB: Archive Object Retrieve Begin</a>

<b>Code</b>	<b>Description</b>	<b>See</b>
ARCE	Archive Object Retrieve End: Object data has been retrieved from an external archival storage system, and the ARC service reports the status of the retrieval operation.	<a href="#">ARCE: Archive Object Retrieve End</a>
ARCT	Archive Retrieve from Cloud-Tier: Archived object data is retrieved from an external archival storage system, which connects to the StorageGRID through the S3 API.	<a href="#">ARCT: Archive Retrieve from Cloud-Tier</a>
AREM	Archive Object Remove: A content block was successfully or unsuccessfully deleted from the external archival storage system.	<a href="#">AREM: Archive Object Remove</a>
ASCE	Archive Object Store End: A content block has been written to the external archival storage system, and the ARC service reports the status of the write operation.	<a href="#">ASCE: Archive Object Store End</a>
ASCT	Archive Store Cloud-Tier: Object data is stored to an external archival storage system, which connects to the StorageGRID through the S3 API.	<a href="#">ASCT: Archive Store Cloud-Tier</a>
ATCE	Archive Object Store Begin: Writing a content block to an external archival storage has started.	<a href="#">ATCE: Archive Object Store Begin</a>
AVCC	Archive Validate Cloud-Tier Configuration: The account and bucket settings provided were successfully or unsuccessfully validated.	<a href="#">AVCC: Archive Validate Cloud-Tier Configuration</a>
CBSE	Object Send End: The source entity completed a grid-node to grid-node data transfer operation.	<a href="#">CBSE: Object Send End</a>

Code	Description	See
CBRE	Object Receive End: The destination entity completed a grid-node to grid-node data transfer operation.	<a href="#">CBRE: Object Receive End</a>
SCMT	Object Store Commit: A content block was completely stored and verified, and can now be requested.	<a href="#">SCMT: Object Store Commit</a>
SREM	Object Store Remove: A content block was deleted from a grid node, and can no longer be requested directly.	<a href="#">SREM: Object Store Remove</a>

#### Client read audit messages

Client read audit messages are logged when an S3 or Swift client application makes a request to retrieve an object.

Code	Description	Used by	See
SGET	S3 GET: Logs a successful transaction to retrieve an object or list the objects in a bucket.  <b>Note:</b> If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	<a href="#">SGET: S3 GET</a>
SHEA	S3 HEAD: Logs a successful transaction to check for the existence of an object or bucket.	S3 client	<a href="#">SHEA: S3 HEAD</a>
WGET	Swift GET: Logs a successful transaction to retrieve an object or list the objects in a container.	Swift client	<a href="#">WGET: Swift GET</a>
WHEA	Swift HEAD: Logs a successful transaction to check for the existence of an object or container.	Swift client	<a href="#">WHEA: Swift HEAD</a>

## Client write audit messages

Client write audit messages are logged when an S3 or Swift 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 clients Swift clients	<a href="#">OVWR: Object Overwrite</a>
SDEL	S3 DELETE: Logs a successful transaction to delete an object or bucket.  <b>Note:</b> If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	<a href="#">SDEL: S3 DELETE</a>
SPOS	S3 POST: Logs a successful transaction to restore an object from AWS Glacier storage to a Cloud Storage Pool.	S3 client	<a href="#">SPOS: S3 POST</a>
SPUT	S3 PUT: Logs a successful transaction to create a new object or bucket.  <b>Note:</b> If the transaction operates on a subresource, the audit message will include the field S3SR.	S3 client	<a href="#">SPUT: S3 PUT</a>
SUPD	S3 Metadata Updated: Logs a successful transaction to update the metadata for an existing object or bucket.	S3 client	<a href="#">SUPD: S3 Metadata Updated</a>
WDEL	Swift DELETE: Logs a successful transaction to delete an object or container.	Swift client	<a href="#">WDEL: Swift DELETE</a>

Code	Description	Used by	See
WPUT	Swift PUT: Logs a successful transaction to create a new object or container.	Swift client	<a href="#">WPUT: Swift PUT</a>

#### 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.	<a href="#">MGAU: Management audit message</a>

#### Audit messages

When system events occur, the StorageGRID system generates audit messages and records them in the audit log.

#### APCT: Archive Purge from Cloud-Tier

This message is generated when archived object data is deleted from an external archival storage system, which connects to the StorageGRID through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was deleted.
CSIZ	Content Size	The size of the object in bytes. Always returns 0.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the cloud-tier from which the object was deleted.

#### ARCB: Archive Object Retrieve Begin

This message is generated when a request is made to retrieve archived object data and the retrieval process begins. Retrieval requests are processed immediately, but can be reordered to improve efficiency of retrieval from linear media such as tape.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival storage system.
RSLT	Result	Indicates the result of starting the archive retrieval process. Currently defined value is:SUCS: The content request was received and queued for retrieval.

This audit message marks the time of an archive retrieval. It allows you to match the message with a corresponding ARCE end message to determine the duration of archive retrieval, and whether the operation was successful.

#### **ARCE: Archive Object Retrieve End**

This message is generated when an attempt by the Archive Node to retrieve object data from an external archival storage system completes. If successful, the message indicates that the requested object data has been completely read from the archive location, and was successfully verified. After the object data has been retrieved and verified, it is delivered to the requesting service.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival storage system.
VLID	Volume Identifier	The identifier of the volume on which the data was archived.If an archive location for the content is not found, a Volume ID of 0 is returned.
RSLT	Retrieval Result	The completion status of the archive retrieval process: <ul style="list-style-type: none"> <li>• SUCS: successful</li> <li>• VRFL: failed (object verification failure)</li> <li>• ARUN: failed (external archival storage system unavailable)</li> <li>• CANC: failed (retrieval operation canceled)</li> <li>• GERR: failed (general error)</li> </ul>

Matching this message with the corresponding ARCB message can indicate the time taken to perform the

archive retrieval. This message indicates whether the retrieval was successful, and in the case of failure, the cause of the failure to retrieve the content block.

**ARCT: Archive Retrieve from Cloud-Tier**

This message is generated when archived object data is retrieved from an external archival storage system, which connects to the StorageGRID through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was retrieved.
CSIZ	Content Size	The size of the object in bytes. The value is only accurate for successful retrieves.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the external archival storage system.
TIME	Time	Total processing time for the request in microseconds.

**AREM: Archive Object Remove**

The Archive Object Remove audit message indicates that a content block was successfully or unsuccessfully deleted from an Archive Node. If the result is successful, the Archive Node has successfully informed the external archival storage system that StorageGRID has released an object location. Whether the object is removed from the external archive storage system depends on the type of system and its configuration.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the Content Block to be retrieved from the external archival media system.
VLID	Volume Identifier	The identifier of the volume on which the object data was archived.

Code	Field	Description
RSLT	Result	<p>The completion status of the archive removal process:</p> <ul style="list-style-type: none"> <li>• SUCS: successful</li> <li>• ARUN: failed (external archival storage system unavailable)</li> <li>• GERR: failed (general error)</li> </ul>

**ASCE: Archive Object Store End**

This message indicates that writing a content block to an external archival storage system has ended.

Code	Field	Description
CBID	Content Block Identifier	The identifier of the content block stored on the external archival storage system.
VLID	Volume Identifier	The unique identifier of the archive volume to which the object data is written.
VREN	Verification Enabled	<p>Indicates if verification is performed for content blocks. Currently defined values are:</p> <ul style="list-style-type: none"> <li>• VENA: verification is enabled</li> <li>• VDSA: verification is disabled</li> </ul>
MCLS	Management Class	A string identifying the TSM Management Class to which the content block is assigned if applicable.

Code	Field	Description
RSLT	Result	<p>Indicates the result of the archive process. Currently defined values are:</p> <ul style="list-style-type: none"> <li>• SUCS: successful (archiving process succeeded)</li> <li>• OFFL: failed (archiving is offline)</li> <li>• VRFL: failed (object verification failed)</li> <li>• ARUN: failed (external archival storage system unavailable)</li> <li>• GERR: failed (general error)</li> </ul>

This audit message means that the specified content block has been written to the external archival storage system. If the write fails, the result provides basic troubleshooting information about where the failure occurred. More detailed information about archive failures can be found by examining Archive Node attributes in the StorageGRID system.

**ASCT: Archive Store Cloud-Tier**

This message is generated when archived object data is stored to an external archival storage system, which connects to StorageGRID through the S3 API.

Code	Field	Description
CBID	Content Block ID	The unique identifier for the content block that was retrieved.
CSIZ	Content Size	The size of the object in bytes.
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	Unique identifier (UUID) of the cloud-tier the content was stored to.
TIME	Time	Total processing time for the request in microseconds.

**ATCE: Archive Object Store Begin**

This message indicates that writing a content block to an external archival storage has started.

Code	Field	Description
CBID	Content Block ID	The unique identifier of the content block to be archived.
VLID	Volume Identifier	The unique identifier of the volume to which the content block is written. If the operation fails, a volume ID of 0 is returned.
RSLT	Result	Indicates the result of the transfer of the content block. Currently defined values are: <ul style="list-style-type: none"> <li>• SUCS: success (content block stored successfully)</li> <li>• EXIS: ignored (content block was already stored)</li> <li>• ISFD: failed (insufficient disk space)</li> <li>• STER: failed (error storing the CBID)</li> <li>• OFFL: failed (archiving is offline)</li> <li>• GERR: failed (general error)</li> </ul>

**AVCC: Archive Validate Cloud-Tier Configuration**

This message is generated when the configuration settings are validated for a Cloud Tiering - Simple Storage Service (S3) target type.

Code	Field	Description
RSLT	Result Code	Returns successful (SUCS) or the error reported by the backend.
SUID	Storage Unique Identifier	UUID associated with the external archival storage system being validated.

**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.
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.

<b>Code</b>	<b>Field</b>	<b>Description</b>
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.
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.

Code	Field	Description
RSLT	Transfer Result	<p>The result of the transfer operation (from the perspective of the sending entity):</p> <p>SUCS: transfer successfully completed; all requested sequence counts were sent.</p> <p>CONL: connection lost during transfer</p> <p>CTMO: connection timed-out during establishment or transfer</p> <p>UNRE: destination node ID unreachable</p> <p>CRPT: transfer ended due to reception of corrupt or invalid data (might indicate tampering)</p>

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.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	<p>Indicates if the CBID transfer was push-initiated or pull-initiated:</p> <p>PUSH: The transfer operation was requested by the sending entity.</p> <p>PULL: The transfer operation was requested by the receiving entity.</p>

Code	Field	Description
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 (might indicate tampering)

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.

**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.
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.

Code	Field	Description
RSLT	Reason Code	<p>The reason for the failure:</p> <p>SCNI: Secure connection establishment failed.</p> <p>CERM: Certificate was missing.</p> <p>CERT: Certificate was invalid.</p> <p>CERE: Certificate was expired.</p> <p>CERR: Certificate was revoked.</p> <p>CSGN: Certificate signature was invalid.</p> <p>CSGU: Certificate signer was unknown.</p> <p>UCRM: User credentials were missing.</p> <p>UCRI: User credentials were invalid.</p> <p>UCRU: User credentials were disallowed.</p> <p>TOUT: Authentication timed out.</p>

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: <ul style="list-style-type: none"> <li>• SUCS: Successful</li> <li>• SUNV: Service Unavailable</li> <li>• GERR: Other failure</li> </ul>
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: <ul style="list-style-type: none"> <li>• SUCS: Successful</li> <li>• SUNV: Service Unavailable</li> <li>• GERR: Other failure</li> </ul>
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	<p>This field uniquely identifies a generated grid task and allows the grid task to be managed over its lifecycle.</p> <p><b>Note:</b> 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.</p>
RSLT	Result	<p>The final status result of the grid task:</p> <ul style="list-style-type: none"> <li>• SUCS: The grid task completed successfully.</li> <li>• ABRT: The grid task was aborted without a rollback error.</li> <li>• ROLF: The grid task was aborted and was unable to complete the rollback process.</li> <li>• CANC: The grid task was canceled by the user before it was started.</li> <li>• EXPR: The grid task expired before it was started.</li> <li>• IVLD: The grid task was invalid.</li> <li>• AUTH: The grid task was unauthorized.</li> <li>• DUPL: The grid task was rejected as a duplicate.</li> </ul>

**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	<p>This field uniquely identifies a generated grid task and allows the task to be managed over its lifecycle.</p> <p><b>Note:</b> 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.</p>
RSLT	Result	<p>The result. This field has only one value:</p> <ul style="list-style-type: none"> <li>• SUCS: The grid task was started successfully.</li> </ul>

#### 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	<p>Uniquely identifies a generated grid task and allows the task to be managed over its lifecycle.</p> <p><b>Note:</b> 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.</p>
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.

Code	Field	Description
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: <ul style="list-style-type: none"> <li>• TXTB: The grid task was submitted through the StorageGRID system as a signed text block.</li> <li>• GRID: The grid task was submitted through the internal Grid Task Submission Service.</li> </ul>
ACTV	Activation Type	The type of activation: <ul style="list-style-type: none"> <li>• AUTO: The grid task was submitted for automatic activation.</li> <li>• PEND: The grid task was submitted into the pending table. This is the only possibility for the TXTB source.</li> </ul>
RSLT	Result	The result of the submission: <ul style="list-style-type: none"> <li>• SUCS: The grid task was submitted successfully.</li> <li>• FAIL: The task has been moved directly to the historical table.</li> </ul>

**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 or Swift containers.** This message is generated when ILM

starts the process of deleting an object because no placement instructions in the active ILM policy currently apply to the object.

<b>Code</b>	<b>Field</b>	<b>Description</b>
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 do not include this field.
CSIZ	Content size	The size of the object in bytes.

Code	Field	Description
LOCS	Locations	<p>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).</p> <p>CLEC: for erasure-coded objects, the erasure coding profile ID and the erasure coding group ID that is applied to the object's data.</p> <p>CLDI: for replicated objects, the LDR node ID and the volume ID of the object's location.</p> <p>CLNL: ARC node ID of the object's location if the object data is archived.</p>
PATH	S3 Bucket/Key or Swift Container/Object ID	The S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
RSLT	Result	<p>The result of the ILM operation.</p> <p>SUCS: The ILM operation was successful.</p>
RULE	Rules Label	<ul style="list-style-type: none"> <li>• If an object in a compliant S3 bucket is being deleted automatically because its retention period has expired, this field is blank.</li> <li>• 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.</li> </ul>
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 do not 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 or Swift 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.
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 or Swift Container/Object ID	The S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
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.

### LLST: Location Lost

This message is generated whenever a location for an object copy (replicated or erasure coded) cannot be found.

Code	Field	Description
CBIL	CBID	The affected CBID.
NOID	Source Node ID	The node ID on which the locations were lost.
UUID	Universally Unique ID	The identifier of the affected object in the StorageGRID system.
ECPR	Erasure Coding Profile	For erasure-coded object data. The ID of the Erasure Coding profile used.

Code	Field	Description
LTyp	Location Type	CLDI (Online): For replicated object data  CLEC (Online): For erasure-coded object data  CLNL (Nearline): For archived replicated object data
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

#### MGAU: Management audit message

The Management category logs user requests to the Management API. Every request that is not a GET or HEAD request to the API logs a response with the username, IP, and type of request to the API.

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.

Code	Field	Description
MRBD	Request body	<p>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:</p> <ul style="list-style-type: none"> <li>• Username and account ID in <b>POST authorize</b></li> <li>• New subnets configuration in <b>POST /grid/grid-networks/update</b></li> <li>• New NTP servers in <b>POST /grid/ntp-servers/update</b></li> <li>• Decommissioned server IDs in <b>POST /grid/servers/decommission</b></li> </ul> <p><b>Note:</b> Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).</p>
MRMD	Request method	<p>The HTTP request method:</p> <ul style="list-style-type: none"> <li>• POST</li> <li>• PUT</li> <li>• DELETE</li> <li>• PATCH</li> </ul>
MRSC	Response code	The response code.
MRSP	Response body	<p>The content of the response (the response body) is logged by default.</p> <p><b>Note:</b> Sensitive information is either deleted (for example, an S3 access key) or masked with asterisks (for example, a password).</p>
MSIP	Source IP address	The client (source) IP address.

Code	Field	Description
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 cannot 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 nearline 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 or Swift Container/Object ID	If available, the S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
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 or Archive 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
CBID	Content Block Identifier	The CBID of the object.
CSIZ	Content size	The size of the object in bytes.
LOCS	Locations	<p>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).</p> <p>CLEC: for erasure-coded objects, the erasure coding profile ID and the erasure coding group ID that is applied to the object's data.</p> <p>CLDI: for replicated objects, the LDR node ID and the volume ID of the object's location.</p> <p>CLNL: ARC node ID of the object's location if the object data is archived.</p>
PATH	S3 Bucket/Key or Swift Container/Object ID	The S3 bucket name and S3 key name, or the Swift container name and Swift object identifier.
RSLT	Result	<p>The result of the ILM operation.</p> <p>SUCS: The ILM operation was successful.</p>
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 if the object is segmented.

Code	Field	Description
STAT	Status	<p>The status of ILM operation.</p> <p>DONE: ILM operations against the object have completed.</p> <p>DFER: The object has been marked for future ILM re-evaluation.</p> <p>PRGD: The object has been deleted from the StorageGRID system.</p> <p>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.</p>
UUID	Universally Unique Identifier	The identifier of the object within the StorageGRID system.

The ORLM audit message can be issued a number of times for a single object. For instance, it is issued whenever one of the following events take place:

- 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.

Code	Field	Description
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 or Swift Object Path	The S3 or Swift object path used for both the previous and new object
RSLT	Result Code	Result of the Object Overwrite transaction. Result is always:  SUCS: Successful

**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. 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 do not 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 do not 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 do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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
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 do not 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.
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.

Code	Field	Description
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 do not 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. 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 do not 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 retrieved object in bytes. Operations on buckets do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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).

<b>Code</b>	<b>Field</b>	<b>Description</b>
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.
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 do not 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.
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 do not 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 do not include this field.
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 do not include this field.

Code	Field	Description
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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 do not 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 do not include this field.

#### SPOS: S3 POST

When an S3 client issues a POST Object restore request, a request is made to restore an object from AWS Glacier storage to a Cloud Storage Pool. 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.

Code	Field	Description
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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 POST Object restore 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 do not 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.
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: urn:sgws:identity::03393893651506583485:root  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 do not 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. 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 do not include this field.

Code	Field	Description
CMPS	Compliance Settings	The compliance settings used when creating the bucket, if present in the PUT Bucket 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 do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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).
LKEN	Object Lock Enabled	Value of the request header x-amz-bucket-object-lock-enabled, if present in the PUT Bucket request.
LKLH	Object Lock Legal Hold	Value of the request header x-amz-object-lock-legal-hold, if present in the PUT Object request.
LKMD	Object Lock Retention Mode	Value of the request header x-amz-object-lock-mode, if present in the PUT Object request.
LKRU	Object Lock Retain Until Date	Value of the request header x-amz-object-lock-retain-until-date, if present in the PUT Object request.

<b>Code</b>	<b>Field</b>	<b>Description</b>
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	S3KY	The S3 key name, not including the bucket name. Operations on buckets do not 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.
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).

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.
ULID	Upload ID	Included only in SPUT messages for Complete Multipart Upload 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 do not include this field.
VSST	Versioning State	The new versioning state of a bucket. Two states are used: "enabled" or "suspended." Operations on objects do not 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.

Code	Field	Description
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 do not 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 do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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
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 do not 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.

Code	Field	Description
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 do not 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.

Code	Field	Description
RSLT	Result Code	<p>Verification failure type:</p> <p>CRCF: Cyclic redundancy check (CRC) failed.</p> <p>HMAC: Hash-based message authentication code (HMAC) check failed.</p> <p>EHSB: Unexpected encrypted content hash.</p> <p>PHSH: Unexpected original content hash.</p> <p>SEQC: Incorrect data sequence on disk.</p> <p>PERR: Invalid structure of disk file.</p> <p>DERR: Disk error.</p> <p>FNAM: Bad file name.</p>

**Note:** This message should be monitored closely. Content verification failures can indicate attempts to tamper with content or 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 'SUCS' so that this message is not filtered.

**Note:** 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 attempts to tamper with content or 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 prior to 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 cannot 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 cannot 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.

#### **VLST: User Initiated Volume Lost**

This message is issued whenever the `/proc/CMSI/Volume_Lost` command is run.

Code	Field	Description
VOLL	Volume Identifier Lower	The lower end of the affected volume range or a single volume.
VOLU	Volume Identifier Upper	The upper end of the affected volume range. Equal to VOLL if a single volume.
NOID	Source Node ID	The node ID on which the locations were lost.
LTYP	Location Type	'CLDI' (Online) or 'CLNL' (Nearline). If not specified, defaults to 'CLDI'.
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.

#### **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 do not include this field.
CSIZ	Content Size	The size of the deleted object in bytes. Operations on containers do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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.
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.

Code	Field	Description
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers do not 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 do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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
SAIP	IP address of requesting client	The IP address of the client application that made the request.

Code	Field	Description
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 do not include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers do not 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 do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on accounts and containers do not include this field.

Code	Field	Description
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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 do not include this field.
WOBJ	Swift Object	The Swift object identifier. Operations on accounts and containers do not 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 do not include this field.
CSIZ	Content Size	The size of the retrieved object in bytes. Operations on containers do not include this field.
HTRH	HTTP Request Header	List of logged HTTP request header names and values as selected during configuration.  <b>Note:</b> 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.

<b>Code</b>	<b>Field</b>	<b>Description</b>
WCON	Swift Container	The Swift container name.
WOBJ	Swift Object	The Swift object identifier. Operations on containers do not include this field.
WUSR	Swift Account User	The Swift account username that uniquely identifies the client performing the transaction.

# Maintain

## Expand your grid

Learn how to expand a StorageGRID system without interrupting system operations.

- [Planning a StorageGRID expansion](#)
- [Preparing for an expansion](#)
- [Overview of expansion procedure](#)
- [Adding storage volumes to Storage Nodes](#)
- [Adding grid nodes to an existing site or adding a new site](#)
- [Configuring your expanded StorageGRID system](#)
- [Contacting technical support](#)

### Planning a StorageGRID expansion

You can expand StorageGRID to increase storage capacity, to add metadata capacity, to add redundancy or new capabilities, or to add a new site. The number, type, and location of the nodes you need to add depends on the reason for the expansion.

- [Adding storage capacity](#)
- [Adding metadata capacity](#)
- [Adding grid nodes to add capabilities to your system](#)
- [Adding a new site](#)

### Adding storage capacity

When existing Storage Nodes become full, you must increase the storage capacity of your StorageGRID system.

To increase storage capacity, you must first understand where data is stored currently and then add capacity in all required locations. For example, if you currently store copies of object data at several sites, you might need to increase the storage capacity of each site.

- [Guidelines for adding object capacity](#)
- [Adding storage capacity for replicated objects](#)
- [Adding storage capacity for erasure-coded objects](#)
- [Considerations for rebalancing erasure-coded data](#)

### 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 or erasure-coded objects. See the instructions for managing objects with information lifecycle management.
- You cannot 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 16 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 cannot add storage volumes to any other storage appliance.
- You cannot increase the size of an existing storage volume.
- You cannot 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 expansion shelves to an SG6060 storage appliance, see the instructions for SG6000 appliance installation and maintenance.

### [SG6000 storage appliances](#)

- For a software-based node, follow the instructions for adding storage volumes to Storage Nodes.

### [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 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 the steps for checking data repair jobs in the recovery and maintenance instructions.
- 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 the 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.



Learn about the ADC quorum in the recovery and maintenance instructions.

- 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 cannot 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 cannot add the ADC service to a node after it is deployed.

## Related information

[Manage objects with ILM](#)

[SG6000 storage appliances](#)

[Adding storage volumes to Storage Nodes](#)

[Maintain & recover](#)

[Performing the expansion](#)

### Adding 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 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 as described in the instructions for monitoring and troubleshooting StorageGRID.

When planning the timing of an expansion, remember to consider how long it might take to procure and install additional storage.

## Related information

[Manage objects with ILM](#)

[Monitor & troubleshoot](#)

### Adding 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.

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 lead to this recommendation or to develop a more precise plan for your site, review the next section. For a custom recommendation optimized for your situation, contact your NetApp account representative.

### Calculating the number of expansion Storage Nodes to add for erasure-coded objects

To optimize how you expand a deployment that stores erasure-coded objects, you must consider many factors:

- Erasure-coding scheme in use
- Characteristics of the storage pool used for erasure coding, including the number of nodes at each site and the amount of free space on each node
- Whether the grid was previously expanded (because the amount of free space per Storage Node might not be approximately the same on all nodes)
- Exact nature of the ILM policy, such as whether ILM rules make both replicated and erasure-coded objects

The following examples can help you understand the impact of the erasure-coding scheme, the number of nodes in the storage pool, and the amount of free space on each node.

Similar considerations affect the calculations for an ILM policy that stores both replicated and erasure-coded data and the calculations for a grid that has been previously expanded.



The examples in this section represent the best practices for adding storage capacity to a StorageGRID system. If you are unable to add the recommended number of nodes, you might need to run the EC rebalance procedure to allow additional erasure-coded objects to be stored.

[Considerations for rebalancing erasure-coded data](#)

### Example 1: Expanding a one-site grid that uses 2+1 erasure coding

This example shows how to expand a simple grid that includes only three Storage Nodes.



This example uses only three Storage Nodes for simplicity. However, using only three Storage Nodes is not recommended: an actual production grid should use a minimum of  $k+m + 1$  Storage Nodes for redundancy, which equals four Storage Nodes (2+1+1) for this example.

Assume the following:


- All data is stored using the 2+1 erasure-coding scheme. With the 2+1 erasure coding scheme, every object is stored as three fragments, and each fragment is saved on a different Storage Node.
- You have one site with three Storage Nodes. Each Storage Node has a total capacity of 100 TB.
- You want to expand by adding new 100 TB Storage Nodes.
- You want to eventually balance erasure-coded data across the old and new nodes.

You have a number of options, based on how full the Storage Nodes are when you perform the expansion.

- **Add three 100 TB Storage Nodes when the existing nodes are 100% full**

In this example, the existing nodes are 100% full. Because there is no free capacity, you must immediately add three nodes to continue 2+1 erasure coding.

After the expansion is complete, when objects are erasure-coded, all fragments will be placed on the new nodes.




This expansion adds  $k+m$  nodes. Adding four nodes is recommended for redundancy. If you add only  $k+m$  expansion Storage Nodes when existing nodes are 100% full, all new objects must be stored on the expansion nodes. If any of the new nodes become unavailable, even temporarily, StorageGRID cannot meet ILM requirements.

- **Add two 100 TB Storage Nodes, when the existing Storage Nodes are 67% full**

In this example, the existing nodes are 67% full. Because there are 100 TB of free capacity on the existing

nodes (33 TB per node), you only need to add two nodes if you perform the expansion now.


Adding 200 TB of additional capacity will allow you to continue 2+1 erasure coding and to eventually balance erasure-coded data across all nodes.



- **Add one 100 TB Storage Node when the existing Storage Nodes are 33% full**

In this example, the existing nodes are 33% full. Because there are 200 TB of free capacity on the existing nodes (67 TB per node), you only need to add one node if you perform the expansion now.

Adding 100 TB of additional capacity will allow you to continue 2+1 erasure coding and to eventually balance erasure-coded data across all nodes.



### Example 2: Expanding a three-site grid that uses 6+3 erasure coding

This example shows how to develop an expansion plan for a multi-site grid that has an erasure-coding scheme with a larger number of fragments. Despite the differences between these examples, the recommended expansion plan is very similar.

Assume the following:

- All data is stored using the 6+3 erasure coding scheme. With the 6+3 erasure coding scheme, every object is stored as 9 fragments, and each fragment is saved to a different Storage Node.
- You have three sites, and each site has four Storage Nodes (12 nodes in total). Each node has a total capacity of 100 TB.
- You want to expand by adding new 100 TB Storage Nodes.
- You want to eventually balance erasure-coded data across the old and new nodes.

You have a number of options, based on how full the Storage Nodes are when you perform the expansion.

- **Add nine 100 TB Storage Nodes (three per site), when existing nodes are 100% full**

In this example, the 12 existing nodes are 100% full. Because there is no free capacity, you must immediately add nine nodes (900 TB of additional capacity) to continue 6+3 erasure coding.

After the expansion is complete, when objects are erasure-coded, all fragments will be placed on the new nodes.



This expansion adds  $k+m$  nodes. Adding 12 nodes (four per site) is recommended for redundancy. If you add only  $k+m$  expansion Storage Nodes when existing nodes are 100% full, all new objects must be stored on the expansion nodes. If any of the new nodes become unavailable, even temporarily, StorageGRID cannot meet ILM requirements.

- **Add six 100 TB Storage Nodes (two per site), when existing nodes are 75% full**

In this example, the 12 existing nodes are 75% full. Because there are 300 TB of free capacity (25 TB per node), you only need to add six nodes if you perform the expansion now. You would add two nodes to each of the three sites.

Adding 600 TB of storage capacity will allow you to continue 6+3 erasure coding and to eventually balance erasure-coded data across all nodes.

- **Add three 100 TB Storage Nodes (one per site), when existing nodes are 50% full**

In this example, the 12 existing nodes are 50% full. Because there are 600 TB of free capacity (50 TB per node), you only need to add three nodes if you perform the expansion now. You would add one node to each of the three sites.

Adding 300 TB of storage capacity will allow you to continue 6+3 erasure coding and to eventually balance erasure-coded data across all nodes.

## Related information

[Manage objects with ILM](#)

[Monitor & troubleshoot](#)

[Considerations for rebalancing erasure-coded data](#)

### Considerations for rebalancing erasure-coded data

If you are performing an expansion to add Storage Nodes and your ILM policy includes one or more ILM rules to erasure code data, you might need to perform the EC rebalance procedure after the expansion is complete.

For example, if you cannot add the recommended number of Storage Nodes in an expansion, you might need to run the EC rebalance procedure to allow additional erasure-coded objects to be stored.

### 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 expanded Storage Nodes at a site.

When the EC rebalance procedure runs:

- It only moves erasure-coded object data. It does not move replicated object data.
- It redistributes the data within a site. It does not move data between sites.
- It redistributes data among all Storage Nodes at a site. It does not redistribute data within storage volumes.

When the EC rebalance procedure is complete:


- Erasure-coded data is moved from Storage Nodes with less available space to Storage Nodes with more available space.
- Used (%) values might remain different between Storage Nodes because the EC rebalance procedure does not move replicated object copies.
- The data protection of erasure-coded objects will be unchanged.

When the EC rebalance procedure is running, the performance of ILM operations and S3 and Swift client operations are likely to be impacted. For this reason, you should only perform this procedure in limited cases.

### When not to perform an EC rebalance


As an example of when you do not need to perform an EC rebalance, consider the following:

- 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 0.2 MB and a 2-copy replication rule for smaller objects.
- All Storage Nodes have become completely full, and the **Low Object Storage** alert has been triggered at the major severity level. The recommended action is to perform an expansion procedure to add Storage Nodes.




To expand the site in this example, it is recommended that you add three or more new Storage Nodes. StorageGRID requires three Storage Nodes for 2+1 erasure coding so that it can place the two data fragments and the one parity fragment on different nodes.

After you add the three Storage Nodes, the original Storage Nodes remain full, but objects can continue to be ingested into the 2+1 erasure coding scheme on the new nodes. Running the EC rebalance procedure is not recommended for this case: running the procedure will temporarily decrease performance, which might impact client operations.




### When to perform an EC rebalance

As an example of when you should perform the EC rebalance procedure, consider the same example, but assume that you can only add two Storage Nodes. Because 2+1 erasure coding requires at least three Storage Nodes, the new nodes cannot be used for erasure-coded data.



To resolve this issue and make use of the new Storage Nodes, you can run the EC rebalance procedure. When this procedure runs, StorageGRID redistributes 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.



### Considerations for EC rebalancing

In general, you should only run the EC rebalance procedure in limited cases. Specifically, you should perform EC rebalancing only 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 the recommended number of new Storage Nodes for the erasure-coding scheme in use.

## Adding storage capacity for erasure-coded objects

- Your S3 and Swift clients can tolerate lower performance for their write and read operations while the EC rebalance procedure is running.

### How the EC rebalance procedure interacts with other maintenance tasks

You cannot 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.  <ul style="list-style-type: none"><li>• You are prevented from starting a decommission procedure or an EC data repair while the EC rebalance procedure is running.</li><li>• You are prevented from starting the EC rebalance procedure while a Storage Node decommission procedure or an EC data repair is running.</li></ul>
Expansion procedure	No.  If you need to add new Storage Nodes in an expansion, you should wait to run the EC rebalance procedure until after you have added all new nodes. If an EC rebalance procedure is in progress when you add new Storage Nodes, data will not be moved to those nodes.
Upgrade procedure	No.  If you need to upgrade StorageGRID software, you should 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, you should wait to run the EC rebalance procedure until after you have added the new node. If an EC rebalance procedure is in progress when you add new Storage Nodes, data will not be moved to those nodes.
Hotfix procedure	Yes.  You can apply a StorageGRID hotfix while the EC rebalance procedure is running.

Procedure	Allowed during EC rebalance procedure?
Other maintenance procedures	No.  You must terminate the EC rebalance procedure before running other maintenance procedures.

### How the 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, do not start using an ILM rule that has a different Erasure Coding profile. If you need to make such ILM changes, you should abort the EC rebalance procedure.

#### Related information

[Rebalancing erasure-coded data after adding Storage Nodes](#)

### Adding 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. See the instructions for monitoring and troubleshooting StorageGRID for details.

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.

#### 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. See the instructions for administering StorageGRID for more information.
- You cannot increase metadata capacity by adding storage volumes to existing Storage Nodes, because metadata is stored only on volume 0.
- You cannot 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.

## 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.



### Related information

[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

### Adding 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 additional Gateway Nodes to support the creation of High Availability groups of Gateway Nodes, 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
- Archive Nodes

When preparing to add grid nodes, be aware of the following limitations:

- The primary Admin Node is deployed during the initial installation. You cannot 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.

#### [Adding storage capacity](#)

- If you are adding Archive Nodes, note that each Archive Node only supports tape through Tivoli Storage Manager (TSM) middleware.
- If the **New Node Client Network Default** option is set to **Untrusted** on the Untrusted Client Networks page, client applications that connect to expansion nodes using the Client Network must connect using a load balancer endpoint port (**Configuration > Network Settings > Untrusted Client Network**). See the instructions for administering StorageGRID to change the setting for the new node and to configure load balancer endpoints.

### Related information

#### [Administer StorageGRID](#)

### Adding 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 cannot 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 information on monitoring object metadata capacity, see the instructions for monitoring and troubleshooting StorageGRID.

- 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 are not 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 in the instructions for managing objects with information lifecycle management.

## Related information

[Manage objects with ILM](#)

## Preparing for an expansion

You must prepare for the StorageGRID expansion by obtaining required materials and installing and configuring any new hardware and networks.

### Gathering required materials

Before you perform an expansion operation, you must gather the materials listed in the following table.

Item	Notes
StorageGRID installation archive	<p>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.</p> <p>For details, see the instructions for downloading and extracting the StorageGRID installation files.</p> <p><b>Note:</b> You do not need to download files if you are adding new storage volumes to existing Storage Nodes or installing a new StorageGRID appliance.</p>
Service laptop	<p>The service laptop must meet the following requirements:</p> <ul style="list-style-type: none"><li>• Network port</li><li>• SSH client (for example, PuTTY)</li><li>• Supported browser</li></ul>
Provisioning passphrase	<p>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.</p>
StorageGRID documentation	<ul style="list-style-type: none"><li>• <i>Administering StorageGRID</i></li><li>• <i>StorageGRID Release Notes</i></li><li>• Installation instructions for your platform</li></ul>

Item	Notes
Current documentation for your platform	For supported versions, see the Interoperability Matrix.

### Related information

[Administer StorageGRID](#)

[Release notes](#)

[Install VMware](#)

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

[Install Ubuntu or Debian](#)

[NetApp Interoperability Matrix Tool](#)

### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

### Downloading and extracting 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 the NetApp Downloads page for StorageGRID.

[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 check box, 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
VMware	<code>StorageGRID-Webscale-version-VMware-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-VMware-uniqueID.tgz</code>
Red Hat Enterprise Linux or CentOS	<code>StorageGRID-Webscale-version-RPM-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-RPM-uniqueID.tgz</code>
Ubuntu or Debian and Appliance	<code>StorageGRID-Webscale-version-DEB-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-DEB-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 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.

<b>Path and file name</b>	<b>Description</b>
<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-archive.ovf</code> <code>./vsphere/vsphere-archive.mf</code>	The template file (.ovf) and manifest file (.mf) for deploying Archive 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.
<b>Deployment scripting tool</b>	<b>Description</b>
<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>	A sample 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.
<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 is enabled.
<code>./vsphere/configure-storagegrid.sample.json</code>	A sample 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.

9. If you are expanding a Red Hat Enterprise Linux or CentOS 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 or CentOS hosts.
<code>./rpms/StorageGRID-Webscale-Service-version-SHA.rpm</code>	RPM package for installing the StorageGRID host service on your RHEL or CentOS 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>	A sample 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.
<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 or CentOS hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.

10. 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.
<code>./debs/NLF000000.txt</code>	A non-production NetApp License File that you can use for testing and proof of concept deployments.

Path and file name	Description
<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.
<code>./debs/configure-storagegrid.sample.json</code>	A sample 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.

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>	Checksum of the DEB installation package used by the StorageGRID Appliance Installer to validate that the package is intact after upload.



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.

## Verifying hardware and networking

Before you begin the expansion of your StorageGRID system, you must ensure that you have installed and configured the necessary hardware to support the new grid nodes or new site.

For information about supported versions, see the Interoperability Matrix.

You must also verify network connectivity between servers at the site, and confirm that the primary Admin Node can communicate with all expansion servers that are intended to host the StorageGRID system.

If you are performing an expansion activity that includes adding a new subnet, you must add the new Grid subnet before you start the expansion procedure.

Do not 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.

### Related information

[NetApp Interoperability Matrix Tool](#)

[Updating subnets for the Grid Network](#)

## Overview of expansion procedure

The basic steps for performing a StorageGRID expansion vary for the different types of expansion: adding storage volumes to a Storage Node, adding new nodes to an existing site, or adding a new site. In all cases, you can perform expansions without interrupting the operation of your current system.


The type of node that you are adding to the grid or the reason that you are adding nodes does not affect the basic expansion procedure. But as shown in the workflow diagram below, the steps for adding nodes vary slightly depending on whether you are adding StorageGRID appliances or hosts running VMware or Linux.



NetApp-provided virtual machine disk files and scripts for new installations or expansions of StorageGRID on OpenStack are no longer supported. To expand an existing deployment on OpenStack, refer to the steps for your Linux distribution.



“Linux” refers to a Red Hat® Enterprise Linux®, Ubuntu®, CentOS, or Debian® deployment. Use the NetApp Interoperability Matrix Tool to get a list of supported versions.



### Related information

[NetApp Interoperability Matrix Tool](#)

[Planning a StorageGRID expansion](#)

[Preparing for an expansion](#)

[Adding storage volumes to Storage Nodes](#)

[Adding grid nodes to an existing site or adding a new site](#)

### Adding storage volumes to Storage Nodes

You can expand the storage capacity of Storage Nodes that have 16 or fewer storage volumes by adding additional storage volumes. You might need to add storage volumes to more than one Storage Node to satisfy ILM requirements for replicated or erasure-coded copies.

### What you'll need

Before adding storage volumes, review the guidelines for adding storage capacity to ensure that you know where to add volumes to meet the requirements of your ILM policy.

## Adding storage capacity



These instructions apply to software-based Storage Nodes only. See the installation and maintenance instructions for the SG6060 appliance to learn how to add storage volumes to SG6060 by installing expansion shelves. Other appliance Storage Nodes cannot be expanded.

[SG6000 storage appliances](#)

### About this task

The underlying storage of a Storage Node is divided into a number of 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 16 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. That is, 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 software (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	0.03%	0 bytes/s	4 KB/s	
cvloc(8:2,sda2)	N/A	0.37%	0 bytes/s	29 KB/s	
sdc(8:16,sdb)	N/A	0.00%	0 bytes/s	0 bytes/s	
sdd(8:32,sdc)	N/A	0.00%	0 bytes/s	183 bytes/s	
sde(8:48,sdd)	N/A	0.00%	0 bytes/s	12 bytes/s	

Volumes					
Mount Point	Device	Status	Size	Available	Write Cache Status
/	croot	Online	10.50 GB	3.46 GB	Unknown
/var/local	cvloc	Online	96.59 GB	94.99 GB	Unknown
/var/local/rangedb/0	sdc	Online	53.66 GB	53.57 GB	Enabled
/var/local/rangedb/1	sdd	Online	53.66 GB	53.57 GB	Enabled
/var/local/rangedb/2	sde	Online	53.66 GB	53.57 GB	Enabled

Object Stores						
ID	Size	Available	Object Data	Object Data (%)	Health	
0000	53.66 GB	48.21 GB	976.25 KB	0.00%	No Errors	
0001	53.66 GB	53.57 GB	0 bytes	0.00%	No Errors	
0002	53.66 GB	53.57 GB	0 bytes	0.00%	No Errors	

3. Follow the instructions for your platform to add new storage volumes to the Storage Node.

- [VMware: Adding storage volumes to a Storage Node](#)
- [Linux: Adding direct-attached or SAN volumes to a Storage Node](#)

### VMware: Adding storage volumes to a Storage Node

If a Storage Node includes fewer than 16 storage volumes, you can increase its capacity by using VMware vSphere to add volumes.

#### What you'll need

- You must have access to the instructions for installing StorageGRID for VMware deployments.
- You must have the `Passwords.txt` file.
- You must have specific access permissions.



Do not 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 may 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`



Do not 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 storage volumes are formatted.
- e. When asked, enter **y** to stop storage services.

The storage services are stopped, and the `setup_rangedbs.sh` script runs automatically. After the volumes are ready for use as rangedbs, the services start again.

5. Check that the services start correctly:
  - a. View a listing 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 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**.
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 now use the expanded capacity of the Storage Nodes to save object data.

## Related information

[Install VMware](#)

## Linux: Adding direct-attached or SAN volumes to a Storage Node

If a Storage Node includes fewer than 16 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.

## What you'll need

- You must have access to the instructions for installing StorageGRID for your Linux platform.
- You must have the `Passwords.txt` file.
- You must have specific access permissions.



Do not 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 disk 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 into 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`. Do not 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-snl-var-local
BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-snl-rangedb-0
BLOCK_DEVICE_RANGEDB_01 = /dev/mapper/sgws-snl-rangedb-1
BLOCK_DEVICE_RANGEDB_02 = /dev/mapper/sgws-snl-rangedb-2
BLOCK_DEVICE_RANGEDB_03 = /dev/mapper/sgws-snl-rangedb-3
<strong>BLOCK_DEVICE_RANGEDB_04 = /dev/mapper/sgws-snl-rangedb-
4</strong>
<strong>BLOCK_DEVICE_RANGEDB_05 = /dev/mapper/sgws-snl-rangedb-
5</strong>
<strong>BLOCK_DEVICE_RANGEDB_06 = /dev/mapper/sgws-snl-rangedb-
6</strong>
<strong>BLOCK_DEVICE_RANGEDB_07 = /dev/mapper/sgws-snl-rangedb-
7</strong>
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 listing 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 storage volumes are formatted.

- d. When asked, enter **y** to stop storage services.

The storage services are stopped, and the `setup_rangedbs.sh` script runs automatically. After the volumes are ready for use as rangedbs, the services start again.

14. Check that the services start correctly:

- a. View a listing 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
```

15. Verify that the Storage Node is online:

- a. Sign in to the Grid Manager using a supported 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**.
16. 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.

## Related information

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

[Install Ubuntu or Debian](#)

## Adding grid nodes to an existing site or adding a new site

You can follow this procedure to add grid nodes to existing sites or to add a new site, but you cannot perform both types of expansion at the same time.

### What you'll need

- You must have root or maintenance permissions. For details, see information about controlling system access with administration user accounts and groups.
- All existing nodes in the grid must be up and running across all sites.
- Any previous expansion, upgrade, decommissioning, or recovery procedures must be 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. [Updating subnets for the Grid Network](#)
2. [Deploying new grid nodes](#)
3. [Performing the expansion](#)

### Updating subnets for the 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.

- You must have the provisioning passphrase.
- You must 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 the new Grid subnet before you start the expansion procedure.

### Steps

1. Select **Maintenance > Network > Grid Network**.

#### Grid Network

Configure the subnets that are used on the Grid Network. These entries typically include the subnets for the Grid Network (eth0) for each site in your StorageGRID system as well as any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

#### Subnets

Subnet 1            +

#### Passphrase

Provisioning  
Passphrase

Save

2. In the Subnets list, click the plus sign to add a new subnet in CIDR notation.

For example, enter 10.96.104.0/22.

3. Enter the provisioning passphrase, and click **Save**.

The subnets you have specified are configured automatically for your StorageGRID system.

### Deploying 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 the grid, the nodes you add do not have to match the existing node types. You can add VMware nodes, Linux container-based nodes, or appliance nodes.

#### VMware: Deploying 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 grid node as a 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.

### [Deploying a StorageGRID node as a virtual machine](#)

2. After you have deployed all new VMware nodes, return to these instructions to perform the expansion procedure.

### [Performing the expansion](#)

#### **Linux: Deploying 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.

#### **What you'll need**

- You have the instructions for installing StorageGRID for your version of Linux, and you have reviewed the hardware and storage requirements.
- 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, do not 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 AFF 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.

#### **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:  
`CLI:sudo storagegrid node start [<node name\>]`

#### **After you finish**

After deploying all new grid nodes, you can perform the expansion.

#### **Related information**

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

## [Install Ubuntu or Debian](#)

## [Performing 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.

#### **What you'll need**

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- You have used the StorageGRID Appliance Installer to complete all of the “configuring the hardware” steps in the appliance installation and maintenance instructions.

Configuring 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 version on the replacement appliance matches the software version of your StorageGRID system. (If the versions do not match, you must upgrade the StorageGRID Appliance Installer firmware.)

For instructions, see the appliance installation and maintenance instructions.

- [SG100 & SG1000 services appliances](#)
- [SG5600 storage appliances](#)
- [SG5700 storage appliances](#)
- [SG6000 storage appliances](#)
- 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	<ol style="list-style-type: none"><li>a. Unselect the <b>Enable Admin Node discovery</b> check box.</li><li>b. Enter the IP address manually.</li><li>c. Click <b>Save</b>.</li><li>d. Wait for the connection state for the new IP address to become ready.</li></ol>
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"><li>a. Select the <b>Enable Admin Node discovery</b> check box.</li><li>b. Wait for the list of discovered IP addresses to be displayed.</li><li>c. Select the primary Admin Node for the grid where this appliance Storage Node will be deployed.</li><li>d. Click <b>Save</b>.</li><li>e. Wait for the connection state for the new IP address to become ready.</li></ol>

4. In the **Node name** field, enter the name you want to use for this appliance node, and click **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 installation and maintenance instructions for your appliance.

6. From the StorageGRID Appliance Installer home page, click **Start Installation**.

## 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

Connection state

Connection to 172.16.4.210 ready



### Node name

Node name




### Installation

Current state

Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

The Current state changes to “Installation is in progress,” and the Monitor Installation page is displayed.

- 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.

- If you need to manually access the Monitor Installation page, click **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 are not re-run. If you are re-running an installation, any tasks that do not 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 software 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.

[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...

```

### Related information

[SG5700 storage appliances](#)[SG5600 storage appliances](#)[SG6000 storage appliances](#)[SG100 & SG1000 services appliances](#)

### Performing the expansion

When you perform the expansion, the new grid nodes are added to your existing

## StorageGRID deployment.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.
- You must have deployed all of the grid nodes that are being added in this expansion.
- If you are adding Storage Nodes, you must have confirmed that all data-repair operations performed as part of a recovery are complete. See the steps for checking data repair jobs in the recovery and maintenance instructions.
- If you are adding a new site, you must review and update ILM rules before starting the expansion procedure to ensure that object copies are not 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 the ILM rule 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. See the instructions for managing objects with information lifecycle management.

### About this task

Performing the expansion includes these phases:

1. You configure the expansion by specifying whether you are adding new grid nodes or a new site and approving the grid nodes you want to add.
2. You start the expansion.
3. While the expansion process is running, you download a new Recovery Package file.
4. You monitor the status of the grid configuration tasks, which run automatically. The set of tasks depends on what types of grid nodes are being added and on whether a new site is being added.



Some tasks 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 relatively empty. However, if the Cassandra database includes a large amount of object metadata, this stage might take several hours or longer. You can look at the “streamed” percentage shown during the “Starting Cassandra and streaming data” stage to determine how complete the Cassandra streaming operation is.

### Steps

1. Select **Maintenance > Maintenance Tasks > Expansion**.

The Grid Expansion page appears. The Pending Nodes section lists all 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.

<a href="#">+ Approve</a> <a href="#">x Remove</a>		<input type="text" value="Search"/>				
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address	
<input type="radio"/>	00:50:56:87:68:1a	DC2-ADM1-184	Admin Node	VMware VM	172.17.3.184/21	
<input type="radio"/>	00:50:56:87:f1:fc	DC2-S1-185	Storage Node	VMware VM	172.17.3.185/21	
<input type="radio"/>	00:50:56:87:54:1e	DC2-S2-186	Storage Node	VMware VM	172.17.3.186/21	
<input type="radio"/>	00:50:56:87:6f:0c	DC2-S3-187	Storage Node	VMware VM	172.17.3.187/21	
<input type="radio"/>	00:50:56:87:b6:83	DC2-S4-188	Storage Node	VMware VM	172.17.3.188/21	
<input type="radio"/>	00:50:56:87:b3:7d	DC2-ARC1-189	Archive Node	VMware VM	172.17.3.189/21	

### 2. Click **Configure Expansion**.

The Site Selection dialog box appears.

### Site Selection

You can add grid nodes to a new site or to existing sites, but you cannot perform both types of expansion at the same time.

Site  New  Existing

Site Name

### 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 grid nodes to an existing site, select **Existing**.

### 4. Click **Save**.

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

As required, you can hover your cursor over a node's **Grid Network MAC Address** to see details about that node.

+ Approve
\* Remove

Grid Network MAC	
<input type="radio"/>	00:50:56:87:68:1a
<input type="radio"/>	00:50:56:87:54:1e
<input type="radio"/>	00:50:56:87:6f:0c
<input type="radio"/>	00:50:56:87:b6:83
<input type="radio"/>	00:50:56:87:b3:7d

### DC2-S3-187

**Storage Node**

Address	Name
<b>Network</b>	
Grid Network	172.17.3.187/21 172.17.0.1
Admin Network	
Client Network	10.224.3.187/21 10.224.0.1

**Hardware**

VMware VM 8 CPUs 8 GB RAM

**Disks**

107 GB 107 GB 107 GB 107 GB 107 GB



If a grid node is missing, confirm that it was deployed successfully.

6. From the list of pending nodes, approve the grid nodes for this expansion.
  - a. Select the radio button next to the first pending grid node you want to approve.
  - b. Click **Approve**.

The grid node configuration form appears.

## Storage Node Configuration

### General Settings

Site	<input type="text" value="Site A"/>
Name	<input type="text" value="DC2-S3-187"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

Select "Yes" if this node will replace another node at this site that has the ADC service.

### Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.17.3.187/21"/>
Gateway	<input type="text" value="172.17.0.1"/>

### Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text"/>
Gateway	<input type="text"/>
Subnets (CIDR)	<input type="text"/> +

### Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text"/>
Gateway	<input type="text"/>

Cancel

Save

c. As required, modify the general settings:

- **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 hostname that will be assigned to the node, and the name that will be displayed in the Grid Manager.
- **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.



Assign the Primary NTP role to at least two nodes at each site. This provides redundant system access to external timing sources.

- **ADC Service** (Storage Nodes only): 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 cannot add the ADC service to a node after it is deployed.
  - If you are adding this node to replace a Storage Node, select **Yes** if the node you are replacing includes the ADC service. Because you cannot 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.
  - Otherwise, select **Automatic** to let the system determine whether this node requires the ADC service.

Learn about the ADC quorum in the recovery and maintenance instructions.

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
- **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. Click **Save**.

The approved grid node moves to the Approved Nodes list.

#### 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.

<input type="button" value="Edit"/> <input type="button" value="Reset"/> <input type="button" value="Remove"/> <span style="float: right;">Search <input type="text"/></span>							
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address	
<input type="radio"/>	00:50:56:87:f1:fc	DC2-S1-185	Site A	Storage Node	VMware VM	172.17.3.185/21	
<input type="radio"/>	00:50:56:87:6f:0c	DC2-S3-187	Site A	Storage Node	VMware VM	172.17.3.187/21	

#### Passphrase

Enter the provisioning passphrase to change the grid topology of your StorageGRID system.

Provisioning Passphrase

- To modify the properties of an approved grid node, select its radio button, and click **Edit**.
- To move an approved grid node back to the Pending Nodes list, select its radio button, and click **Reset**.

- To permanently remove an approved grid node, power the node off. Then, select its radio button, and click **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 click **Expand**.

After a few minutes, this page updates to display the status of the expansion procedure. When tasks that affect individual grid node are in progress, the Grid Node Status section lists the current status for each grid node.



During this process, for appliances the StorageGRID Appliance Installer shows installation moving from Stage 3 to Stage 4, Finalize Installation. When Stage 4 completes, the controller is rebooted.

### Grid Expansion

A new Recovery Package has been generated as a result of the configuration change. Go to the [Recovery Package](#) page to download it.

### 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
<b>Grid Node Status</b>							
Lists the installation and configuration status of each grid node included in the expansion.							
Search <input type="text"/>							
Name	Site	Grid Network IPv4 Address	Progress	Stage			
DC2-ADM1-184	Site A	172.17.3.184/21	<div style="width: 25%;"></div>	Waiting for NTP to synchronize			
DC2-S1-185	Site A	172.17.3.185/21	<div style="width: 25%;"></div>	Waiting for Dynamic IP Service peers			
DC2-S2-186	Site A	172.17.3.186/21	<div style="width: 25%;"></div>	Waiting for NTP to synchronize			
DC2-S3-187	Site A	172.17.3.187/21	<div style="width: 25%;"></div>	Waiting for NTP to synchronize			
DC2-S4-188	Site A	172.17.3.188/21	<div style="width: 25%;"></div>	Waiting for Dynamic IP Service peers			
DC2-ARC1-189	Site A	172.17.3.189/21	<div style="width: 25%;"></div>	Waiting for NTP to synchronize			
2. Initial Configuration							Pending
3. Distributing the new grid node's certificates to the StorageGRID system.							Pending
4. Starting services on the new grid nodes							Pending
5. Cleaning 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. Click the download link.
- b. Enter the provisioning passphrase, and click **Start Download**.
- c. When the download completes, open the .zip file and confirm it includes a `gpt-backup` directory and a `_SAID.zip` file. Then, extract the `_SAID.zip` file, go to the `/GID*_REV*` directory, and confirm you can open 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 one or more Storage Nodes, monitor the progress of the “Starting Cassandra and streaming data” stage by reviewing the percentage shown in the status message.

4. 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
DC1-S4	Data Center 1	10.96.99.55/23	<div style="width: 90%;"></div>	Starting Cassandra and streaming data (90.0% streamed)
DC1-S5	Data Center 1	10.96.99.56/23	<div style="width: 100%;"></div>	Complete
DC1-S6	Data Center 1	10.96.99.57/23	<div style="width: 100%;"></div>	Complete

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.



Do not reboot any Storage Nodes during Step 4 (Starting services on the new grid nodes). The “Starting Cassandra and streaming data” stage might take hours to complete for each new Storage Node, especially if existing Storage Nodes contain a large amount of object metadata.

10. 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, you must perform additional integration and configuration steps.

### Related information

[Manage objects with ILM](#)

## Configuring your expanded StorageGRID system

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 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 grid nodes added during the expansion.

### Steps

1. If you added a Storage Node, complete the following configuration tasks.

Storage Node configuration tasks	For information
<p>Review the storage pools used in your ILM rules to ensure the new storage will be used.</p> <ul style="list-style-type: none"> <li>• If you added a site, create a storage pool for the site and update ILM rules to use the new storage pool.</li> <li>• If you added a Storage Node to an existing site, confirm that the new node uses the correct storage grade.</li> </ul> <p><b>Note:</b> By default, a new Storage Node is assigned to the All Storage Nodes storage grade and added to storage pools that use that grade for the site. If you want a new node to use a custom storage grade, you must assign it to the custom grade manually (<b>ILM &gt; Storage Grades</b>).</p>	<p><a href="#">Manage objects with ILM</a></p>
<p>Verify that the Storage Node is ingesting objects.</p>	<p><a href="#">Verifying that the Storage Node is active</a></p>
<p>Rebalance erasure-coded data (only if you were unable to add the recommended number of Storage Nodes).</p>	<p><a href="#">Rebalancing erasure-coded data after adding Storage Nodes</a></p>

2. If you added a Gateway Node, complete the following configuration tasks.

Gateway Node configuration tasks	For information
<p>If High Availability Groups are used for client connections, add the Gateway Nodes to an HA group. Select <b>Configuration &gt; Network Settings &gt; High Availability Groups</b> to review the list of existing HA groups and to add the new nodes.</p>	<p><a href="#">Administer StorageGRID</a></p>

3. If you added an Admin Node, complete the following configuration tasks.

Admin Node configuration tasks	For information
If single sign-on is enabled for your StorageGRID system, you must create a relying party trust in Active Directory Federation Services (AD FS) for the new Admin Node. You cannot sign in to the node until you create this relying party trust.	<a href="#">Configuring single sign-on</a>
If you plan to use the Load Balancer service on Admin Nodes, you might need to add the Admin Nodes to High Availability groups. Select <b>Configuration &gt; Network Settings &gt; High Availability Groups</b> to review the list of existing HA groups and to add the new nodes.	<a href="#">Administer StorageGRID</a>
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.	<a href="#">Copying the Admin Node database</a>
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.	<a href="#">Copying Prometheus metrics</a>
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.	<a href="#">Copying audit logs</a>
Optionally, configure access to the system for auditing purposes through an NFS or a CIFS file share.  <b>Note:</b> Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.	<a href="#">Administer StorageGRID</a>
Optionally, change the preferred sender for notifications. You can make the expansion Admin Node the preferred sender. Otherwise, an existing Admin Node configured as the preferred sender continues to send notifications, including AutoSupport messages, SNMP notifications, alert emails, and alarm emails (legacy system).	<a href="#">Administer StorageGRID</a>

4. If you added an Archive Node, complete the following configuration tasks.

Archive Node configuration tasks	For information
Configure the Archive Node's connection to the targeted external archival storage system. When you complete the expansion, Archive Nodes are in an alarm state until you configure connection information through the <b>ARC &gt; Target</b> component.	<a href="#">Administer StorageGRID</a>
Update the ILM policy to archive object data through the new Archive Node.	<a href="#">Manage objects with ILM</a>

Archive Node configuration tasks	For information
Configure custom alarms for the attributes that are used to monitor the speed and efficiency of object data retrieval from Archive Nodes.	<a href="#">Administer StorageGRID</a>

- 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 > Network Settings > Untrusted Client Network**.

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 the instructions for administering StorageGRID for more information.

- Configure the Domain Name System (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 information about modifying the DNS configuration for a single grid node in the recovery and maintenance instructions.

The best practice is for the grid-wide DNS server list to contain some DNS servers that are accessible locally from each site. If you just added a new site, add new DNS servers for the site to the grid-wide DNS configuration.



Provide two to six IPv4 addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see the information about modifying the DNS configuration in the recovery and maintenance instructions.

- If you added a new site, confirm that Network Time Protocol (NTP) servers are accessible from that site.



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.

For more information, see the recovery and maintenance instructions.

## Related information

[Manage objects with ILM](#)

[Verifying that the Storage Node is active](#)

[Copying the Admin Node database](#)

[Copying Prometheus metrics](#)

[Copying audit logs](#)

[Upgrade software](#)

[Maintain & recover](#)

## Verifying that the 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 browser.
2. Select **Nodes > Expansion Storage Node > Storage**.
3. Hover 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.

## Copying the 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.

### What you'll need

- You must have completed the required expansion steps to add an Admin Node.
- You must have the `Passwords.txt` file.
- You must 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
- Alarm history
- Historical attribute data, which is used in the charts and text reports available from the **Support > Tools > Grid Topology** 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`

### Copying 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.

#### What you'll need

- The new Admin Node must be installed and running.
- You must have the `Passwords.txt` file.
- You must 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
```

## Copying 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. 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.

### What you'll need

- You must have completed the required expansion steps to add an Admin Node.
- You must have the `Passwords.txt` file.

### About this task

To make the historical audit messages from other Admin Nodes available on the expansion Admin Node, you must copy the audit log files manually from the primary Admin Node, or another existing Admin Node, to the expansion Admin Node.

### 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. Stop the AMS service to prevent it from creating a new file: `service ams stop`

3. Rename the `audit.log` file to ensure that it does not overwrite the file on the expansion Admin Node you are copying it to:

```
cd /var/local/audit/export
ls -l
mv audit.log new_name.txt
```

4. Copy all audit log files to the expansion Admin Node:

```
scp -p * IP_address:/var/local/audit/export
```

5. 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.

6. Restore the original `audit.log` file:

```
mv new_name.txt audit.log
```

7. Start the AMS service:

```
service ams start
```

8. Log out from the server:

```
exit
```

9. 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 `#`.

10. Update the user and group settings for the audit log files:

```
cd /var/local/audit/export
chown ams-user:bycast *
```

11. Log out from the server:

```
exit
```

## Rebalancing erasure-coded data after adding Storage Nodes

In some cases, you might need to rebalance erasure-coded data after you add new Storage Nodes.

### What you'll need

- You must have completed the expansion steps to add the new Storage Nodes.
- You must have reviewed the considerations for rebalancing erasure-coded data.

### [Considerations for rebalancing erasure-coded data](#)



Only perform this procedure if the **Low Object Storage** alert has been triggered for one or more Storage Nodes at a site and you were unable to add the recommended number of new Storage Nodes.

- You must have the `Passwords.txt` file.

### About this task

When the EC rebalance procedure is running, the performance of ILM operations and S3 and Swift client operations are likely to be impacted. For this reason, you should only perform this procedure in limited cases.



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.



S3 and Swift API operations to upload objects (or object parts) might fail during the EC rebalancing procedure if they require more than 24 hours to complete. Long-duration PUT operations will fail if the applicable ILM rule uses Strict or Balanced placement on ingest. The following error will be reported:

```
500 Internal Server Error
```

## 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. Hover 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. Enter the following command:

```
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 the status is `In progress`, the EC rebalance operation is still running. You should periodically monitor the procedure until it completes.
- If the status is `Failure`, perform the [failure steps](#).
- If the status is `Success`, perform the [success step](#).

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 abort --job-id job-id
```



When you terminate an EC rebalance procedure, any data fragments that have already been moved remain in the new location. Data is not moved back to the original location.

9. If the status of the EC rebalance procedure is `Failure`, follow these steps:

- 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.

For information about specific alerts, see the monitoring and troubleshooting instructions.

c. Restart the EC rebalance procedure:

```
rebalance-data start --job-id job-id
```

d. If the status of the EC rebalance procedure is still `Failure`, contact technical support.

10. If the status of the EC rebalance procedure 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.



Replicated object data is not moved by the EC rebalance procedure.

11. If you are using erasure coding at more than one site, run this procedure for all other affected sites.

## Related information

[Considerations for rebalancing erasure-coded data](#)

[Monitor & troubleshoot](#)

## Contacting technical support

If you encounter errors during the grid expansion process that you are unable to resolve, or if a grid task fails, contact technical support.

### About this task

When you contact technical support, you must provide the required log files to assist in troubleshooting the errors you are encountering.

### 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 Docker container 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.

Once you are logged 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	<ul style="list-style-type: none"><li>• <code>/var/log/daemon.log</code></li><li>• <code>/var/log/storagegrid/daemon.log</code></li><li>• <code>/var/log/storagegrid/nodes/&lt;node-name&gt;.log</code></li></ul>
Linux	<ul style="list-style-type: none"><li>• <code>/var/log/storagegrid/daemon.log</code></li><li>• <code>/etc/storagegrid/nodes/&lt;node-name&gt;.conf</code> (for each failed node)</li><li>• <code>/var/log/storagegrid/nodes/&lt;node-name&gt;.log</code> (for each failed node; might not exist)</li></ul>

## Maintain & recover

Learn how to apply a hotfix; recover a failed grid node; decommission grid nodes and sites; and recover objects in the case of a system failure.

- [Introduction to StorageGRID recovery and maintenance](#)
- [StorageGRID hotfix procedure](#)

- [Grid node recovery procedures](#)
- [How site recovery is performed by technical support](#)
- [Decommission procedure](#)
- [Network maintenance procedures](#)
- [Host-level and middleware procedures](#)
- [Grid node procedures](#)
- [Appliance node cloning](#)

## Introduction to StorageGRID recovery and maintenance

The recovery and maintenance procedures for StorageGRID include applying a software hotfix, recovering grid nodes, recovering a failed site, decommissioning grid nodes or an entire site, performing network maintenance, performing host-level and middleware maintenance procedures, and performing grid node procedures.

All recovery and maintenance activities require a broad understanding of the StorageGRID system. You should review your StorageGRID system's topology to ensure that you understand the grid configuration.

You must follow all instructions exactly and heed all warnings.

Maintenance procedures not described are not supported or require a services engagement.

For hardware procedures, see the installation and maintenance instructions for your StorageGRID appliance.



“Linux” refers to a Red Hat® Enterprise Linux®, Ubuntu®, CentOS, or Debian® deployment. Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

### Related information

[Grid primer](#)

[Network guidelines](#)

[Administer StorageGRID](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[NetApp Interoperability Matrix Tool](#)

### Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	87
Microsoft Edge	87
Mozilla Firefox	84

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

## Downloading the Recovery Package

The Recovery Package file allows you to restore the StorageGRID system if a failure occurs.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the provisioning passphrase.
- You must 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:
  - a. Open the `.zip` file.
  - b. Confirm it includes a `gpt-backup` directory and an inner `.zip` file.
  - c. Extract the inner `.zip` file.
  - d. Confirm you can open 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.

## Related information

[Administer StorageGRID](#)

## 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](#)
- [How your system is affected when you apply a hotfix](#)
- [Obtaining the required materials for a hotfix](#)
- [Downloading the hotfix file](#)
- [Checking the system's condition before applying a hotfix](#)
- [Applying the hotfix](#)

### Considerations for applying a hotfix

When you apply a hotfix, a cumulative series of software updates is applied to the nodes in your StorageGRID system.

You cannot apply a StorageGRID hotfix when another maintenance procedure is running. For example, you cannot 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 do not need to deploy the major release first.

4. You then apply hotfix 11.A.B.y to the primary Admin Node.

### Related information

[Configuring the 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.

#### 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. To configure high availability groups, see the instructions for administering StorageGRID.

#### 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:

- Do not 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.
- Do not update the ILM configuration until the hotfix has been applied to all nodes.

## Obtaining the required materials for a hotfix

Before applying a hotfix, you must obtain all required materials.

Item	Notes
StorageGRID hotfix file	You must download the StorageGRID hotfix file.
<ul style="list-style-type: none"><li>• Network port</li><li>• Supported web browser</li><li>• SSH client (for example, PuTTY)</li></ul>	See “Web browser requirements.”
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 included in the SAID package, which 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.

### Related information

[Downloading the hotfix file](#)

[Downloading the Recovery Package](#)

### Downloading the hotfix file

You must download the hotfix file before you can apply the hotfix.

#### Steps

1. Go to the NetApp Downloads page for StorageGRID.

[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.



Do not 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**.

#### Related information

[Configuring the replacement primary Admin Node](#)

#### Checking the system's condition before applying a hotfix

You must verify the system is ready to accommodate the hotfix.

1. Sign in to the Grid Manager using a supported 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.

For information about specific alerts, see the instructions for monitoring and troubleshooting StorageGRID.

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 cannot apply a StorageGRID hotfix when another maintenance procedure is running. For example, you cannot 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.

#### Related information

### Applying the 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.

### What you'll need

- You have reviewed all of the considerations and completed all of the steps in “Hotfix planning and preparation.”
- You must have the provisioning passphrase.
- You must have Root Access or the Maintenance permission.
- 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 cannot perform a StorageGRID software upgrade or a SANtricity OS upgrade until you have completed the hotfix process.

### Steps

1. Sign in to the Grid Manager using a supported 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.

- To perform a major version upgrade of StorageGRID, see the [instructions for upgrading StorageGRID](#), and then select **StorageGRID Upgrade**.
- To apply a hotfix to all nodes in your system, see “Hotfix procedure” in the [recovery and maintenance instructions](#), and then select **StorageGRID Hotfix**.
- To upgrade SANtricity OS software on a storage controller, see “Upgrading SANtricity OS Software on the storage controllers” in the installation and maintenance instructions for your storage appliance, and then select **SANtricity OS**.

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)



3. Select **StorageGRID 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 

Browse

### Passphrase

Provisioning Passphrase 

Start

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.



Do not change the file name since it is part of the verification process.

## 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 

Browse

✔ hotfix-install-11.5.0.1

Details 

hotfix-install-11.5.0.1

### Passphrase

Provisioning Passphrase 

Start

5. Enter the provisioning passphrase in the text box.

The **Start** button becomes enabled.

## 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    hotfix-install-11.5.0.1

Details  hotfix-install-11.5.0.1

### Passphrase

Provisioning Passphrase 

Start

#### 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.

### Warning

#### Connection Might be Temporarily Lost

When the hotfix is applied, your browser's connection might be lost temporarily as services on the primary Admin Node are stopped and restarted. Are you sure you want to start the hotfix installation process?

Cancel

OK

#### 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
- Storage Nodes

- Archive Nodes



The progress bar reaches completion, and then the primary Admin Node is shown first with stage “Complete.”

### Hotfix Installation Progress

Approve All
Remove All

Admin Nodes - 1 out of 1 completed

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.



Do not 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.

10. If you need to remove a node or all nodes from the hotfix queue, select **Remove** or **Remove All**.

As shown in the example, 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

11. 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.

12. 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.

### Related information

[Hotfix planning and preparation](#)

[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

## Grid node recovery procedures

If a grid node fails, you can recover it by replacing the failed physical or virtual server, reinstalling StorageGRID software, and restoring recoverable data.

Grid nodes can fail if a hardware, virtualization, operating system, or software fault renders the node inoperable or unreliable. There are many kinds of failure that can trigger the need to recover a grid node.

The steps to recover a grid node vary, depending on the platform where the grid node is hosted and on the type of grid node. Each type of grid node has a specific recovery procedure, which you must follow exactly.

Generally, you try to preserve data from the failed grid node where possible, repair or replace the failed node, use the Grid Manager to configure the replacement node, and restore the node's data.



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.

## Related information

[How site recovery is performed by technical support](#)

## 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. Do not use disk snapshots to restore grid nodes. Instead, refer to the recovery and maintenance procedures for each type of node.

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, Archive 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 hardware installation and maintenance for verifying and upgrading the StorageGRID Appliance Installer version.
  - [SG100 & SG1000 services appliances](#)
  - [SG5600 storage appliances](#)
  - [SG5700 storage appliances](#)
  - [SG6000 storage appliances](#)
- 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.

## 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

Do not 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.

#### Gathering 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	<p>If you need to recover a grid node, you need the StorageGRID installation archive for your platform.</p> <p><b>Note:</b> You do not need to download files if you are recovering failed storage volumes on a Storage Node.</p>
Recovery Package .zip file	<p>Obtain a copy of the most recent Recovery Package .zip file: <code>sgws-recovery-package-id-revision.zip</code></p> <p>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.</p> <p>If the primary Admin Node is operating normally, you can download the Recovery Package from the Grid Manager. Select <b>Maintenance &gt; System &gt; Recovery Package</b>.</p> <p>If you cannot 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.</p>
Passwords.txt file	<p>Contains the passwords required to access grid nodes on the command line. Included in the Recovery Package.</p>
Provisioning passphrase	<p>The passphrase is created and documented when the StorageGRID system is first installed. The provisioning passphrase is not in the Passwords.txt file.</p>

Item	Notes
Current documentation for your platform	<p>For the current supported versions of your platform, see the Interoperability Matrix Tool.</p> <p><a href="#">NetApp Interoperability Matrix Tool</a></p> <p>Go to the platform vendor's website for documentation.</p>

## Related information

[Downloading and extracting the StorageGRID installation files](#)

[Web browser requirements](#)

### Downloading and extracting the StorageGRID installation files

Before you can recover StorageGRID grid nodes, you must download the software and extract the files.

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 Grid Manager, go to **Help > About**.
2. Go to the NetApp Downloads page for StorageGRID.

[NetApp Downloads: 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 check box, 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
VMware	<code>StorageGRID-Webscale-version-VMware-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-VMware-uniqueID.tgz</code>
Red Hat Enterprise Linux or CentOS	<code>StorageGRID-Webscale-version-RPM-uniqueID.zip</code>
	<code>StorageGRID-Webscale-version-RPM-uniqueID.tgz</code>

Platform	Installation archive
Ubuntu or Debian or Appliances	StorageGRID-Webscale- <i>version</i> -DEB- <i>uniqueID</i> .zip  StorageGRID-Webscale- <i>version</i> -DEB- <i>uniqueID</i> .tgz
OpenStack or other hypervisor	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.

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 VMware system, select the appropriate files.

Path and file name	Description
./vsphere/README	A text file that describes all of the files contained in the StorageGRID download file.
./vsphere/NLF000000.txt	A free license that does not provide any support entitlement for the product.
./vsphere/NetApp-SG- <i>version</i> -SHA.vmdk	The virtual machine disk file that is used as a template for creating grid node virtual machines.
./vsphere/vsphere-primary-admin.ovf ./vsphere/vsphere-primary-admin.mf	The Open Virtualization Format template file (.ovf) and manifest file (.mf) for deploying the primary Admin Node.
./vsphere/vsphere-non-primary-admin.ovf ./vsphere/vsphere-non-primary-admin.mf	The template file (.ovf) and manifest file (.mf) for deploying non-primary Admin Nodes.
./vsphere/vsphere-archive.ovf ./vsphere/vsphere-archive.mf	The template file (.ovf) and manifest file (.mf) for deploying Archive Nodes.
./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.

<b>Path and file name</b>	<b>Description</b>
<b>Deployment scripting tool</b>	<b>Description</b>
<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>	A sample 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.
<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 is enabled.
<code>./vsphere/configure-storagegrid.sample.json</code>	A sample 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.

10. If you are recovering a Red Hat Enterprise Linux or CentOS system, select the appropriate files.

<b>Path and file name</b>	<b>Description</b>
<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 or CentOS hosts.
<code>./rpms/StorageGRID-Webscale-Service-version-SHA.rpm</code>	RPM package for installing the StorageGRID host service on your RHEL or CentOS hosts.
<b>Deployment scripting tool</b>	<b>Description</b>
<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.

<b>Path and file name</b>	<b>Description</b>
<code>./rpms/configure-storagegrid.sample.json</code>	A sample 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.
<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 or CentOS hosts for StorageGRID container deployment. You can customize the role or playbook as necessary.

11. If you are recovering an Ubuntu or Debian system, select the appropriate files.

<b>Path and file name</b>	<b>Description</b>
<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.
<b>Deployment scripting tool</b>	<b>Description</b>
<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.

Path and file name	Description
<code>./debs/configure-storagegrid.sample.json</code>	A sample 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.

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>	Checksum of the DEB installation package used by the StorageGRID Appliance Installer to validate that the package is intact after upload.

**Note:** 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.

### Related information

[Install VMware](#)

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

[Install Ubuntu or Debian](#)

### Selecting a 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	<p>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.</p> <p><a href="#">How site recovery is performed by technical support</a></p>

Grid node	Recovery procedure
A single Storage Node	<p>The Storage Node recovery procedure depends on the type and duration of the failure.</p> <p><a href="#">Recovering from Storage Node failures</a></p>
Admin Node	<p>The Admin Node procedure depends on whether you need to recover the primary Admin Node or a non-primary Admin Node.</p> <p><a href="#">Recovering from Admin Node failures</a></p>
Gateway Node	<p><a href="#">Recovering from Gateway Node failures.</a></p>
Archive Node	<p><a href="#">Recovering from Archive Node failures.</a></p>



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.

### Recovering 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
<ul style="list-style-type: none"> <li>• More than one Storage Node has failed.</li> <li>• A second Storage Node has failed less than 15 days after a Storage Node failure or recovery.</li> </ul> <p>This includes the case where a Storage Node fails while recovery of another Storage Node is still in progress.</p>	<p>You must contact technical support.</p>	<p>If all failed Storage Nodes are at the same site, it might be necessary to perform a site recovery procedure.</p> <p>Technical support will assess your situation and develop a recovery plan.</p> <p><a href="#">How site recovery is performed by technical support</a></p> <p>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.</p> <p>Technical support can determine when it is safe to begin recovery of a second Storage Node.</p> <p><b>Note:</b> 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.</p>
<p>A Storage Node has been offline for more than 15 days.</p>	<p><a href="#">Recovering a Storage Node that has been down more than 15 days</a></p>	<p>This procedure is required to ensure Cassandra database integrity.</p>
<p>An appliance Storage Node has failed.</p>	<p><a href="#">Recovering a StorageGRID appliance Storage Node</a></p>	<p>The recovery procedure for appliance Storage Nodes is the same for all failures.</p>
<p>One or more storage volumes have failed, but the system drive is intact</p>	<p><a href="#">Recovering from storage volume failure where the system drive is intact</a></p>	<p>This procedure is used for software-based Storage Nodes.</p>
<p>The system drive has failed.</p>	<p><a href="#">Recovering from system drive failure</a></p>	<p>The node replacement procedure depends on the deployment platform and on whether any storage volumes have also failed.</p>



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.

## Recovering a Storage Node that has been down more than 15 days

If a single Storage Node has been offline and not connected to other Storage Nodes for more than 15 days, you must rebuild Cassandra on the node.

### What you'll need

- 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** > **Maintenance Tasks** > **Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **Maintenance** > **Maintenance Tasks** > **Expansion**.)

### About this task

Storage Nodes have a Cassandra database that includes object metadata. If a Storage Node has not been able to communicate with other Storage Nodes for more than 15 days, StorageGRID assumes that node's Cassandra database is stale. The Storage Node cannot rejoin the grid until Cassandra has been rebuilt using information from other Storage Nodes.

Use this procedure to rebuild Cassandra only if a single Storage Node is down. Contact technical support if additional Storage Nodes are offline or if Cassandra has been rebuilt on another Storage Node within the last 15 days; for example, Cassandra might have been rebuilt as part of the procedures to recover failed storage volumes or to recover a failed Storage Node.



If more than one Storage Node has failed (or is offline), contact technical support. Do not 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. Do not perform the following recovery procedure. Data loss could occur.



If more than one Storage Node at a site has failed, a site recovery procedure might be required. Contact technical support.

### How site recovery is performed by technical support

#### Steps

1. If necessary, power on the Storage Node that needs to be recovered.
2. 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 `#`.



If you are unable to log in to the grid node, the system disk might not be intact. Go to the procedure for recovering from system drive failure. [Recovering from system drive failure](#)

3. Perform the following checks on the Storage Node:

- a. Issue this command: `nodetool status`

The output should be `Connection refused`

- b. In the Grid Manager, select **Support > Tools > Grid Topology**.
- c. Select *site* > **Storage Node > SSM > Services**. Verify that the Cassandra service displays `Not Running`.
- d. Select **Storage Node > SSM > Resources**. Verify that there is no error status in the Volumes section.
- e. Issue this command: `grep -i Cassandra /var/local/log/servermanager.log`

You should see the following message in the output:

```
Cassandra not started because it has been offline for more than 15
day grace period - rebuild Cassandra
```

4. Issue this command, and monitor the script output: `check-cassandra-rebuild`

- If storage services are running, you will be prompted to stop them. Enter: **y**
- Review the warnings in the script. If none of them apply, confirm that you want to rebuild Cassandra. Enter: **y**



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. After the rebuild completes, perform the following checks:

- a. In the Grid Manager, select **Support > Tools > Grid Topology**.
- b. Select *site* > **recovered Storage Node > SSM > Services**.
- c. Confirm that all services are running.
- d. Select **DDS > Data Store**.
- e. Confirm that the **Data Store Status** is “Up” and the **Data Store State** is “Normal.”

### Related information

[Recovering from system drive failure](#)


### Recovering a StorageGRID appliance Storage Node

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.

### About this task

You must prepare the appliance and reinstall software, configure the node to rejoin the grid, reformat storage,

and restore object data.



If more than one Storage Node has failed (or is offline), contact technical support. Do not 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. Contact technical support.

#### How site recovery is performed by technical support



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.



If you encounter a Services: Status - Cassandra (SVST) alarm during recovery, see the monitoring and troubleshooting instructions to recover from the alarm by rebuilding Cassandra. After Cassandra is rebuilt, alarms should clear. If alarms do not clear, contact technical support.



For hardware maintenance procedures, such as instructions for replacing a controller or reinstalling SANtricity OS, see the installation and maintenance instructions for your storage appliance.

## Related information

[Monitor & troubleshoot](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

## Steps

- [Preparing an appliance Storage Node for reinstallation](#)
- [Starting StorageGRID appliance installation](#)
- [Monitoring StorageGRID appliance installation](#)
- [Selecting Start Recovery to configure an appliance Storage Node](#)
- [Remounting and reformatting appliance storage volumes \(“Manual Steps”\)](#)
- [Restoring object data to a storage volume for an appliance](#)
- [Checking the storage state after recovering an appliance Storage Node](#)

## Preparing an appliance Storage Node for reinstallation

When recovering an appliance Storage Node, you must first prepare the appliance for reinstallation of StorageGRID software.

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.

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.

## Starting StorageGRID appliance installation

To install StorageGRID on an appliance Storage Node, you use the StorageGRID Appliance Installer, which is included on the appliance.

### What you'll need

- 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 and maintenance instructions for your storage appliance:
  - [SG5600 storage appliances](#)
  - [SG5700 storage appliances](#)
  - [SG6000 storage appliances](#)
- 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 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 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	<ol style="list-style-type: none"><li>a. Unselect the <b>Enable Admin Node discovery</b> check box.</li><li>b. Enter the IP address manually.</li><li>c. Click <b>Save</b>.</li><li>d. Wait while the connection state for the new IP address becomes "ready."</li></ol>
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"><li>a. Select the <b>Enable Admin Node discovery</b> check box.</li><li>b. From the list of discovered IP addresses, select the primary Admin Node for the grid where this appliance Storage Node will be deployed.</li><li>c. Click <b>Save</b>.</li><li>d. Wait while the connection state for the new IP address becomes "ready."</li></ol>

4. In the **Node Name** field, enter the same 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 installation and maintenance instructions for your appliance.

6. From the StorageGRID Appliance Installer home page, click **Start Installation**.

## 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

Connection state

Connection to 172.16.4.210 ready

### Node name

Node name

### Installation

Current state

Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

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.

#### Related information

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

## Monitoring StorageGRID appliance installation

The StorageGRID Appliance Installer provides status until installation is complete. When the software installation is complete, the appliance is rebooted.

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 <span style="float: right;">Running</span>		
Step	Progress	Status
Connect to storage controller	<div style="width: 100%; height: 10px; background-color: green;"></div>	Complete
Clear existing configuration	<div style="width: 100%; height: 10px; background-color: green;"></div>	Complete
Configure volumes	<div style="width: 30%; height: 10px; background-color: blue;"></div>	Creating volume StorageGRID-obj-00
Configure host settings	<div style="width: 0%; height: 10px; background-color: blue;"></div>	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 are not re-run. If you are re-running an installation, any tasks that do not 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 software 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.

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...

```

- Go to the procedure to configure the appliance Storage Node.

### Selecting Start Recovery to configure an 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.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.

- You must have deployed a recovery appliance Storage Node.
- You must know the start date of any repair jobs for erasure-coded data.
- You must have verified that the Storage Node has not been rebuilt within the last 15 days.

## Steps

1. From the Grid Manager, select **Maintenance > 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 cannot 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

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	<div style="width: 20%; background-color: #0070c0;"></div>	Waiting For Manual Steps

Reset



At any point during the recovery, you can click **Reset** to start a new recovery. An Info 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.

### Remounting and reformatting 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.

#### What you'll need

- 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** > **Maintenance Tasks** > **Decommission**.)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **Maintenance** > **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. Do not 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.



Do not 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:

- 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 Webscale will attempt to restore data redundancy by
making additional replicated copies or EC fragments, according to the
rules in the active ILM policy.

Do not continue to the next step if you believe that the data
remaining on this volume cannot 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 Webscale will attempt to restore data redundancy by  
making additional replicated copies or EC fragments, according to the  
rules in the active ILM policy.
```

```
Do not continue to the next step if you believe that the data  
remaining on this volume cannot 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 uninitialized or the disk's superblock was corrupted. When the script cannot 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 do not 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 are not 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 cannot 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).



Do not run the `sn-recovery-postinstall.sh` script if you believe that the data remaining on a failed storage volume cannot 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.
- Do not 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

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	<div style="width: 50%; background-color: #0070C0;"></div>	Recovering Cassandra

6. Return to the Monitor Install page of the StorageGRID Appliance Installer by entering `http://Controller_IP:8080`, using the IP address of the compute controller.

The Monitor Install page shows the installation progress while the script is running.

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, as described in the next procedure.

## Related information

[Reviewing warnings for Storage Node system drive recovery](#)

[Restoring object data to a storage volume for an appliance](#)

## Restoring object data to a storage volume for an appliance

After recovering storage volumes for the appliance Storage Node, you can restore the object data that was lost when the Storage Node failed.

### What you'll need

- You must 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, an Archive Node, or a Cloud Storage Pool, assuming that the grid's ILM rules were configured such that object copies are 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. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.



If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive Node. Due to the latency associated with retrievals from external archival storage systems, restoring object data to a Storage Node from an Archive Node takes longer than restoring copies from other Storage Nodes.

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. You use different options with the `repair-data` script, based on whether you are restoring replicated data or erasure coded data, as follows:

- 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
```

- 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
```

Repairs of erasure coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available. You can track repairs of erasure coded data with this command:

```
repair-data show-ec-repair-status
```



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.

For more information on using the `repair-data` script, enter `repair-data --help` from the command line 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. 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`
3. If all storage volumes have failed, repair the entire node. (If only some volumes have failed, go to the next step.)



You cannot run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

- If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option 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 cannot 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 the instructions for monitoring and troubleshooting StorageGRID.

- If your grid contains erasure coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option 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.

- If your grid has both replicated and erasure coded data, run both commands.

#### 4. If only some of the volumes have failed, repair the affected volumes.

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 are not 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.

- If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` option to identify 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 cannot 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 the instructions for monitoring and troubleshooting StorageGRID.

- If your grid contains erasure coded data, use the `start-ec-volume-repair` command with the `--nodes` option to identify 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.

- If your grid has both replicated and erasure coded data, run both commands.

#### 5. Monitor the repair of replicated data.

- Select **Nodes > Storage Node being repaired > ILM**.
- Use the attributes in the Evaluation section to determine if repairs are complete.

When repairs are complete, the Awaiting - All attribute indicates 0 objects.

- To monitor the repair in more detail, select **Support > Tools > Grid Topology**.
- Select **grid > Storage Node being repaired > LDR > Data Store**.
- 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 are not 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 do not 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.

6. Monitor the repair of erasure coded data, and retry any requests that might have failed.

a. Determine the status of erasure coded data repairs:

- 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.

```

root@DC1-ADM1:~ # repair-data show-ec-repair-status

Repair ID Scope Start Time End Time State Est Bytes
Affected/Repaired Retry Repair
=====
=====
 949283 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:27:06.9 Success
17359 17359 No
 949292 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:37:06.9 Failure
17359 0 Yes
 949294 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:47:06.9 Failure
17359 0 Yes
 949299 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:57:06.9 Failure
17359 0 Yes

```

- b. 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 83930030303133434:

```
repair-data start-ec-node-repair --repair-id 83930030303133434
```

This command retries a failed volume repair, using the repair ID 83930030303133434:

```
repair-data start-ec-volume-repair --repair-id 83930030303133434
```

## Related information

[Monitor & troubleshoot](#)

### Checking the storage state after recovering an 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.

#### What you'll need

- You must be signed in to the Grid Manager using a supported 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.

#### Recovering from storage volume failure where the 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.

#### About this task

This recovery procedure applies to software-based Storage Nodes only. If storage volumes have failed on an appliance Storage Node, use the procedure for “Recovering a StorageGRID appliance Storage Node.”



#### Related information

[Recovering a StorageGRID appliance Storage Node](#)

#### Steps

- [Reviewing warnings about storage volume recovery](#)
- [Identifying and unmounting failed storage volumes](#)
- [Recovering failed storage volumes and rebuilding the Cassandra database](#)
- [Restoring object data to a storage volume where the system drive is intact](#)
- [Checking the storage state after recovering storage volumes](#)

#### Reviewing warnings about storage volume recovery

Before recovering failed storage volumes for a Storage Node, you must 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. Do not 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. Contact technical support.

#### [How site recovery is performed by technical support](#)



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.



If you encounter a Services: Status - Cassandra (SVST) alarm during recovery, see the monitoring and troubleshooting instructions to recover from the alarm by rebuilding Cassandra. After Cassandra is rebuilt, alarms should clear. If alarms do not clear, contact technical support.

#### **Related information**

[Monitor & troubleshoot](#)

[Warnings and considerations for grid node recovery](#)

#### **Identifying and unmounting 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.

#### **What you'll need**

You must be signed in to the Grid Manager using a supported 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. Once a volume has been reformatted, data on the volume cannot 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:

```

/dev/sdc /etc/fstab file ext3 errors=remount-ro,barri
/dev/sdd /var/local ext3 errors=remount-ro,barri
/dev/sde swap swap defaults 0
proc /proc proc defaults 0
sysfs /sys sysfs noauto 0
debugfs /sys/kernel/debug debugfs noauto 0
devpts /dev/pts devpts node=0620,gid=5 0
/dev/td0 /media/Eloppy auto noauto,user,sync 0
/dev/cdrom /cdrom iso9660 ro,noauto 0 0
/dev/disk/by-uuid/384c4687-8511-47a7-9700-7b31b495a0b8 /var/local/mysql_ibda
/dev/mapper/fsgvg-fsglv /fsg xfs daapi,mtpt=/fsg,noalign,nobarrier,ikeep 0 2
/dev/disk/by-uuid/822b0547-3b2b-472e-ad5e-e1cf1809faba /var/local/rangedb/0
  
```

Mount Point	Device	Status	Size	Space Available	Total Entries	Entries Available	Write Cache
/	croot	Online	10.4 GB	4.53 GB	655,360	559,513	Unknown
/var/local	cyloc	Online	96.6 GB	92.8 GB	94,369,792	94,369,445	Unknown
/var/local/rangedb/0	sdc	Online	4,396 GB	4,379 GB	858,993,408	858,983,455	Unavailable
/var/local/rangedb/1	sdd	Online	4,396 GB	4,362 GB	858,993,408	858,973,530	Unavailable
/var/local/rangedb/2	sde	Online	4,396 GB	4,370 GB	858,993,408	858,982,305	Unavailable

## 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 `sd` 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	sd	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 stop the storage services and 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 storage services on the Storage Node.



If the storage services are already stopped, you are not prompted. The Cassandra service is stopped only for volume 0.

```
root@Storage-180:~ # sn-unmount-volume 0
Storage services (ldr, chunk, dds, cassandra) are not down.
Storage services must be stopped before running this script.
Stop storage services [y/N]? y
Shutting down storage services.
Storage services stopped.
Unmounting /var/local/rangedb/0
/var/local/rangedb/0 is unmounted.
```

In a few seconds, the storage services are stopped and the volume is unmounted. Messages appear indicating each step of the process. The final message indicates that the volume is unmounted.

## Recovering failed storage volumes and rebuilding the 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.

- You must have the `Passwords.txt` file.
- The system drives on the server must be intact.
- The cause of the failure must have been identified and, if necessary, replacement storage hardware must already have been acquired.
- The total size of the replacement storage must be 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 > Maintenance Tasks > Decommission.**)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **Maintenance > Maintenance Tasks > Expansion.**)
- You have reviewed the warnings about storage volume recovery.

### Reviewing warnings about storage volume recovery

1. As needed, replace failed physical or virtual storage associated with the failed storage volumes that you identified and unmounted earlier.

After you replace the storage, make sure you rescan or reboot to make sure that it is recognized by the operating system, but do not remount the volumes. The storage is remounted and added to `/etc/fstab` in a later step.

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 `#`.

1. 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.

2. Reformat any failed storage volumes and rebuild the Cassandra database if it is necessary. Enter:

```
reformat_storage_block_devices.rb
```

- If storage services are running, you will be prompted to stop them. Enter: **y**
- 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 do not 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
Storage services must be stopped before running this script.
Stop storage services [y/N]? **y**
Shutting down storage services.
Storage services stopped.
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 c817f87f-f989-4a21-
8f03-b6f42180063f
Skipping in use device /dev/sdg
All devices processed
Running: /usr/local/ldr/setup_rangedb.sh 12075630
Cassandra does not need rebuilding.
Starting services.

Reformatting done. Now do manual steps to
restore copies of data.
```

## Related information

[Reviewing warnings about storage volume recovery](#)

## Restoring object data to a storage volume where the system drive is intact

After recovering a storage volume on a Storage Node where the system drive is intact, you can restore the object data that was lost when the storage volume failed.

### What you'll need

- You must 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, an Archive Node, or a Cloud Storage Pool, assuming that the grid's ILM rules were configured such that object copies are 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. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.



If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive Node. Due to the latency associated with retrievals from external archival storage systems, restoring object data to a Storage Node from an Archive Node takes longer than restoring copies from other Storage Nodes.

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. You use different options with the `repair-data` script, based on whether you are restoring replicated data or erasure coded data, as follows:

- **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
```

- **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
```

Repairs of erasure coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available. You can track repairs of erasure coded data with this command:

```
repair-data show-ec-repair-status
```



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.

For more information on using the `repair-data` script, enter `repair-data --help` from the command line 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. 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`
3. If all storage volumes have failed, repair the entire node. (If only some volumes have failed, go to the next step.)



You cannot run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

- If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option 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 cannot 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 the instructions for monitoring and troubleshooting StorageGRID.

- If your grid contains erasure coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option 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.

- If your grid has both replicated and erasure coded data, run both commands.

4. If only some of the volumes have failed, repair the affected volumes.

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 are not 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.

- If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` option to identify 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 cannot 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 the instructions for monitoring and troubleshooting StorageGRID.

- If your grid contains erasure coded data, use the `start-ec-volume-repair` command with the `--nodes` option to identify 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.

- If your grid has both replicated and erasure coded data, run both commands.

5. Monitor the repair of replicated data.

- Select **Nodes > Storage Node being repaired > ILM**.
- Use the attributes in the Evaluation section to determine if repairs are complete.

When repairs are complete, the Awaiting - All attribute indicates 0 objects.

- To monitor the repair in more detail, select **Support > Tools > Grid Topology**.
- Select **grid > Storage Node being repaired > LDR > Data Store**.
- 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 are not 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 do not 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.

6. Monitor the repair of erasure coded data, and retry any requests that might have failed.

- Determine the status of erasure coded data repairs:
  - 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.

```
root@DC1-ADM1:~ # repair-data show-ec-repair-status

Repair ID Scope Start Time End Time State Est Bytes
Affected/Repaired Retry Repair
=====
=====
949283 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:27:06.9 Success
17359 17359 No
949292 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:37:06.9 Failure
17359 0 Yes
949294 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:47:06.9 Failure
17359 0 Yes
949299 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:57:06.9 Failure
17359 0 Yes
```

- b. 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 83930030303133434:

```
repair-data start-ec-node-repair --repair-id 83930030303133434
```

This command retries a failed volume repair, using the repair ID 83930030303133434:

```
repair-data start-ec-volume-repair --repair-id 83930030303133434
```

## Related information

[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

## Checking the 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported 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.

### Recovering from system drive failure

If the system drive on a software-based Storage Node has failed, the Storage Node is not available to the StorageGRID system. You must complete a specific set of tasks to recover from a system drive failure.


#### About this task

Use this procedure to recover from a system drive failure on a software-based Storage Node. This procedure includes the steps to follow if any storage volumes also failed or cannot be remounted.



This procedure applies to software-based Storage Nodes only. You must follow a different procedure to recover an appliance Storage Node.

### [Recovering a StorageGRID appliance Storage Node](#)



### Steps

- Reviewing warnings for Storage Node system drive recovery
- Replacing the Storage Node
- Selecting Start Recovery to configure a Storage Node
- Remounting and reformatting storage volumes ("Manual Steps")
- Restoring object data to a storage volume, if required

- [Checking the storage state after recovering a Storage Node system drive](#)

## Reviewing warnings for Storage Node system drive recovery

Before recovering a failed system drive of a Storage Node, you must review the following 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. Do not 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. Contact technical support.

### [How site recovery is performed by technical support](#)



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 for recovering from loss of storage volumes where the 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.



If you encounter a Services: Status - Cassandra (SVST) alarm during recovery, see the monitoring and troubleshooting instructions to recover from the alarm by rebuilding Cassandra. After Cassandra is rebuilt, alarms should clear. If alarms do not clear, contact technical support.

## Related information

[Monitor & troubleshoot](#)

[Warnings and considerations for grid node recovery](#)

[Recovering from storage volume failure where the system drive is intact](#)

## Replacing 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.

### Recovering a StorageGRID appliance Storage Node

**Linux:** If you are not sure if your system drive has failed, follow the instructions to replace the node to determine which recovery steps are required.

Platform	Procedure
VMware	<a href="#">Replacing a VMware node</a>
Linux	<a href="#">Replacing a Linux node</a>
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.

### Selecting Start Recovery to configure a 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.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.
- You must have deployed and configured the replacement node.
- You must know the start date of any repair jobs for erasure-coded data.
- You must 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 > 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 cannot 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

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. An Info 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 stage “Waiting for Manual Steps” stage, go to the next task in the recovery procedure to remount and reformat 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	<div style="width: 20%; background-color: #0070C0;"></div>	Waiting For Manual Steps

Reset

### Related information

[Preparing an appliance for reinstallation \(platform replacement only\)](#)

### Remounting and reformatting 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.

#### What you’ll need

- 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 > Maintenance Tasks > Decommission.**)
- You have checked that an expansion is not in progress. (In the Grid Manager, select **Maintenance > Maintenance Tasks > Expansion.**)
- You have reviewed the warnings for Storage Node system drive recovery.

#### [Reviewing 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. Do not 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.



Do not reboot a Storage Node during recovery before running `sn-recovery-postinstall.sh` (see the step for [post-install script](#)) 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:
  - 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 Webscale will attempt to restore data redundancy by
making
additional replicated copies or EC fragments, according to the rules
in
the active ILM policy.

Do not continue to the next step if you believe that the data
remaining on
this volume cannot 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 Webscale will attempt to restore data redundancy by making additional replicated copies or EC fragments, according to the rules in the active ILM policy.

Do not continue to the next step if you believe that the data remaining on this volume cannot 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 uninitialized or the disk's superblock was corrupted. When the script cannot 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 do not 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 are not 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 cannot 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).



Do not run the `sn-recovery-postinstall.sh` script if you believe that the data remaining on a failed storage volume cannot 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.
- Do not 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

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	<div style="width: 50%; background-color: #0070C0;"></div>	Recovering Cassandra

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, as described in that procedure.

### Related information

[Reviewing warnings for Storage Node system drive recovery](#)

[Restoring object data to a storage volume, if required](#)

### Restoring object data to a storage volume, if required

If the `sn-recovery-postinstall.sh` script is needed to reformat one or more failed storage volumes, you must restore object data to the reformatted storage volume from other Storage Nodes and Archive Nodes. These steps are not required unless one or more storage volumes were reformatted.

### What you'll need

- You must 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, an Archive Node, or a Cloud Storage Pool, assuming that the grid's ILM rules were configured such that object copies are 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. Before performing this procedure, contact technical support for help in estimating the recovery time frame and the associated costs.



If the only remaining copy of an object is on an Archive Node, object data is retrieved from the Archive Node. Due to the latency associated with retrievals from external archival storage systems, restoring object data to a Storage Node from an Archive Node takes longer than restoring copies from other Storage Nodes.

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. You use different options with the `repair-data` script, based on whether you are restoring replicated data or erasure coded data, as follows:

- **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
```

- **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
```

Repairs of erasure coded data can begin while some Storage Nodes are offline. Repair will complete after all nodes are available. You can track repairs of erasure coded data with this command:

```
repair-data show-ec-repair-status
```



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.

For more information on using the `repair-data` script, enter `repair-data --help` from the command line 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. 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`
3. If all storage volumes have failed, repair the entire node. (If only some volumes have failed, go to the next step.)



You cannot run `repair-data` operations for more than one node at the same time. To recover multiple nodes, contact technical support.

- If your grid includes replicated data, use the `repair-data start-replicated-node-repair` command with the `--nodes` option 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 cannot 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 the instructions for monitoring and troubleshooting StorageGRID.

- If your grid contains erasure coded data, use the `repair-data start-ec-node-repair` command with the `--nodes` option 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.

- If your grid has both replicated and erasure coded data, run both commands.

4. If only some of the volumes have failed, repair the affected volumes.

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 are not 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.

- If your grid contains replicated data, use the `start-replicated-volume-repair` command with the `--nodes` option to identify 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 cannot 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 the instructions for monitoring and troubleshooting StorageGRID.

- If your grid contains erasure coded data, use the `start-ec-volume-repair` command with the `--nodes` option to identify 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.

- If your grid has both replicated and erasure coded data, run both commands.

5. Monitor the repair of replicated data.

- Select **Nodes > Storage Node being repaired > ILM**.
- Use the attributes in the Evaluation section to determine if repairs are complete.

When repairs are complete, the Awaiting - All attribute indicates 0 objects.

- To monitor the repair in more detail, select **Support > Tools > Grid Topology**.
- Select **grid > Storage Node being repaired > LDR > Data Store**.
- 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 are not 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 do not 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.

6. Monitor the repair of erasure coded data, and retry any requests that might have failed.

- Determine the status of erasure coded data repairs:
  - 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.

```
root@DC1-ADM1:~ # repair-data show-ec-repair-status

Repair ID Scope Start Time End Time State Est Bytes
Affected/Repaired Retry Repair
=====
=====
949283 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:27:06.9 Success
17359 17359 No
949292 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:37:06.9 Failure
17359 0 Yes
949294 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:47:06.9 Failure
17359 0 Yes
949299 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:57:06.9 Failure
17359 0 Yes
```

- b. 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 83930030303133434:

```
repair-data start-ec-node-repair --repair-id 83930030303133434
```

This command retries a failed volume repair, using the repair ID 83930030303133434:

```
repair-data start-ec-volume-repair --repair-id 83930030303133434
```

## Related information

[Administer StorageGRID](#)

[Monitor & troubleshoot](#)

## Checking the storage state after recovering a 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported 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.

### Recovering from Admin Node failures

The recovery process for an Admin Node depends on whether it is the primary Admin Node or a non-primary Admin Node.

#### About this task

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.

**Related information**

### Choices

- [Recovering from primary Admin Node failures](#)
- [Recovering from non-primary Admin Node failures](#)

### Recovering 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.

#### About this task

A failed primary Admin Node should be replaced promptly. 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 cannot 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.



You must repair or replace a failed primary Admin Node within approximately a month 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.

### Steps

- [Copying audit logs from the failed primary Admin Node](#)
- [Replacing the primary Admin Node](#)
- [Configuring the replacement primary Admin Node](#)
- [Restoring the audit log on the recovered primary Admin Node](#)
- [Resetting the preferred sender on the recovered primary Admin Node](#)
- [Restoring the Admin Node database when recovering a primary Admin Node](#)
- [Restoring Prometheus metrics when recovering a primary Admin Node](#)

### Copying audit logs from the 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.

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 are not 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 are not 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. Rename the `audit.log` file so that it does not overwrite the existing file when you copy it to the recovered Admin Node.

Rename `audit.log` to a unique numbered file name such as `yyyy-mm-dd.txt.1`. For example, you can rename the `audit.log` file to `2015-10-25.txt.1`

```
cd /var/local/audit/exporttls -l`mv audit.log 2015-10-25.txt.1
```

4. Restart the AMS service: `service ams start`
5. 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.

6. Copy all audit log files: `scp -p * admin@grid_node_IP:/var/local/tmp/saved-audit-logs`

When prompted, enter the password for admin.

7. Log out as root: `exit`

## Replacing the 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	<a href="#">Replacing a VMware node</a>

Replacement platform	Procedure
Linux	<a href="#">Replacing a Linux node</a>
SG100 and SG1000 services appliances	<a href="#">Replacing a services appliance</a>
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.

### Configuring the replacement primary Admin Node

The replacement node must be configured as the primary Admin Node for your StorageGRID system.

#### What you'll need

- For primary Admin Nodes hosted on virtual machines, the virtual machine must be 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 guide for your appliance.

#### [SG100 & SG1000 services appliances](#)

- You must have the latest backup of the Recovery Package file (`sgws-recovery-package-id-revision.zip`).
- You must have the provisioning passphrase.

#### Steps


1. Open your web browser and navigate to `https://primary_admin_node_ip`.

Install

## Welcome

Use this page to install a new StorageGRID system, or recover a failed primary Admin Node for an existing system.

**Note:** You must have access to a StorageGRID license, network configuration and grid topology information, and NTP settings to complete the installation. You must have the latest version of the Recovery Package file to complete a primary Admin Node recovery.



2. Click **Recover a failed primary Admin Node**.
3. 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**.
4. Enter the provisioning passphrase.
5. 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.

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 Server 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 the instructions for administering StorageGRID. 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.

7. Determine if you need to apply a hotfix.
  - a. Sign in to the Grid Manager using a supported browser.
  - b. Select **Nodes**.
  - c. From the list on the left, select the primary Admin Node.
  - d. On the Overview tab, note the version displayed in the **Software Version** field.
  - e. Select any other grid node.

- f. 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 do not 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.

#### Related information

[Administer StorageGRID](#)

[StorageGRID hotfix procedure](#)

#### Restoring the audit log on the 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.

- The recovered Admin Node must be installed and running.
- You must have copied the audit logs to another location after the original Admin Node failed.

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 cannot 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.

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/audit/export`
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`

You must also restore any pre-existing client access to the audit share. For more information, see the instructions for administering StorageGRID.

#### Related information

[Administer StorageGRID](#)

#### Resetting the preferred sender on the recovered primary Admin Node

If the primary Admin Node you are recovering is currently set as the preferred sender of alert notifications, alarm notifications, and AutoSupport messages, you must reconfigure this setting.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- The recovered Admin Node must be installed and running.

#### Steps

1. Select **Configuration > System Settings > Display Options**.
2. Select the recovered Admin Node from the **Preferred Sender** drop-down list.
3. Click **Apply Changes**.

#### Related information

[Administer StorageGRID](#)

#### Restoring the Admin Node database when recovering a primary Admin Node

If you want to retain the historical information about attributes, alarms, 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.

- The recovered Admin Node must be installed and running.
- The StorageGRID system must include at least two Admin Nodes.
- You must have the `Passwords.txt` file.
- You must have the provisioning passphrase.

If an Admin Node fails, the historical information stored in its Admin Node database is lost. This database includes the following information:

- Alert history
- Alarm history
- Historical attribute data, which is used in the charts and text reports available from the **Support > Tools > Grid Topology** 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 cannot 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.

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`

## Restoring Prometheus metrics when recovering a 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.

- The recovered Admin Node must be installed and running.
- The StorageGRID system must include at least two Admin Nodes.
- You must have the `Passwords.txt` file.
- You must have the provisioning passphrase.

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 cannot 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.

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`

### Recovering 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.

### Related information

[SG100 & SG1000 services appliances](#)

### Steps

- [Copying audit logs from the failed non-primary Admin Node](#)
- [Replacing a non-primary Admin Node](#)
- [Selecting Start Recovery to configure a non-primary Admin Node](#)
- [Restoring the audit log on the recovered non-primary Admin Node](#)
- [Resetting the preferred sender on the recovered non-primary Admin Node](#)
- [Restoring the Admin Node database when recovering a non-primary Admin Node](#)
- [Restoring Prometheus metrics when recovering a non-primary Admin Node](#)

### Copying audit logs from the 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 are not 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 are not 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. Rename the `audit.log` file so that it does not overwrite the existing file when you copy it to the recovered Admin Node.

Rename `audit.log` to a unique numbered file name such as `yyyy-mm-dd.txt.1`. For example, you can rename the `audit.log` file to `2015-10-25.txt.1`

```
cd /var/local/audit/exportls -l`mv audit.log 2015-10-25.txt.1
```

4. Restart the AMS service: `service ams start`
5. 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.

6. Copy all audit log files: `scp -p * admin@grid_node_IP:/var/local/tmp/saved-audit-logs`

When prompted, enter the password for admin.

7. Log out as root: `exit`

## Replacing a 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	<a href="#">Replacing a VMware node</a>
Linux	<a href="#">Replacing a Linux node</a>
SG100 and SG1000 services appliances	<a href="#">Replacing a services appliance</a>

Replacement 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.

### Selecting Start Recovery to configure a 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.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.
- You must have deployed and configured the replacement node.

#### Steps

1. From the Grid Manager, select **Maintenance > 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 cannot 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

<input type="text" value="Search"/>				
	Name	IPv4 Address	State	Recoverable
<input checked="" type="radio"/>	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. An Info 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.
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 Server 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 the instructions for administering StorageGRID. 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.

### Related information

[Administer StorageGRID](#)

[Preparing an appliance for reinstallation \(platform replacement only\)](#)

### Restoring the audit log on the 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.

- The recovered Admin Node must be installed and running.
- You must have copied the audit logs to another location after the original Admin Node failed.

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 cannot 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.

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/audit/export
```

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`

You must also restore any pre-existing client access to the audit share. For more information, see the instructions for administering StorageGRID.

## Related information

[Administer StorageGRID](#)

## Resetting the preferred sender on the recovered non-primary Admin Node

If the non-primary Admin Node you are recovering is currently set as the preferred sender of alert notifications, alarm notifications, and AutoSupport messages, you must reconfigure this setting in the StorageGRID system.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions.
- The recovered Admin Node must be installed and running.

### Steps

1. Select **Configuration > System Settings > Display Options**.
2. Select the recovered Admin Node from the **Preferred Sender** drop-down list.
3. Click **Apply Changes**.

### Related information

[Administer StorageGRID](#)

## Restoring the Admin Node database when recovering a non-primary Admin Node

If you want to retain the historical information about attributes, alarms, and alerts on a non-primary Admin Node that has failed, you can restore the Admin Node database from the primary Admin Node.

- The recovered Admin Node must be installed and running.
- The StorageGRID system must include at least two Admin Nodes.
- You must have the `Passwords.txt` file.
- You must have the provisioning passphrase.

If an Admin Node fails, the historical information stored in its Admin Node database is lost. This database includes the following information:

- Alert history
- Alarm history
- Historical attribute data, which is used in the charts and text reports available from the **Support > Tools > Grid Topology** 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.

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`

### Restoring Prometheus metrics when recovering a non-primary Admin Node

Optionally, you can retain the historical metrics maintained by Prometheus on a non-primary Admin Node that has failed.

- The recovered Admin Node must be installed and running.
- The StorageGRID system must include at least two Admin Nodes.
- You must have the `Passwords.txt` file.
- You must have the provisioning passphrase.

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.

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`

## Recovering from Gateway Node failures

You must complete a sequence of tasks in exact order to recover from a Gateway Node

failure.



#### Related information

[SG100 & SG1000 services appliances](#)

#### Steps

- [Replacing a Gateway Node](#)
- [Selecting Start Recovery to configure a Gateway Node](#)

#### Replacing a 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	<a href="#">Replacing a VMware node</a>
Linux	<a href="#">Replacing a Linux node</a>
SG100 and SG1000 services appliances	<a href="#">Replacing a services appliance</a>
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.

### Selecting Start Recovery to configure a 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.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.
- You must have deployed and configured the replacement node.

#### Steps

1. From the Grid Manager, select **Maintenance > 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 cannot 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

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. An Info 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.

## Related information

[Preparing an appliance for reinstallation \(platform replacement only\)](#)

## Recovering from Archive Node failures

You must complete a sequence of tasks in exact order to recover from an Archive Node failure.



## About this task

Archive Node recovery is affected by the following issues:

- If the ILM policy is configured to replicate a single copy.

In a StorageGRID system that is configured to make a single copy of objects, an Archive Node failure might result in an unrecoverable loss of data. If there is a failure, all such objects are lost; however, you must still perform recovery procedures to “clean up” your StorageGRID system and purge lost object information from the database.

- If an Archive Node failure occurs during Storage Node recovery.

If the Archive Node fails while processing bulk retrievals as part of a Storage Node recovery, you must repeat the procedure to recover copies of object data to the Storage Node from the beginning to ensure that all object data retrieved from the Archive Node is restored to the Storage Node.

### Steps

- [Replacing an Archive Node](#)
- [Selecting Start Recovery to configure an Archive Node](#)
- [Resetting Archive Node connection to the cloud](#)

### Replacing an Archive Node

To recover an Archive Node, you must first replace the 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.

Platform	Procedure
VMware	<a href="#">Replacing a VMware node</a>
Linux	<a href="#">Replacing a Linux node</a>
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.

### Selecting Start Recovery to configure an Archive Node

After replacing an Archive Node, you must select Start Recovery in the Grid Manager to configure the new node as a replacement for the failed node.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.
- You must have deployed and configured the replacement node.

### Steps

1. From the Grid Manager, select **Maintenance > 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 cannot 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

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. An Info 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`

## Resetting Archive Node connection to the cloud

After you recover an Archive Node that targets the cloud through the S3 API, you need to modify configuration settings to reset connections. An Outbound Replication Status (ORSU) alarm is triggered if the Archive Node is unable to retrieve object data.



If your Archive Node connects to external storage through TSM middleware, then the node resets itself automatically and you do not need to reconfigure.

### What you'll need

You must be signed in to the Grid Manager using a supported browser.

### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target**.
3. Edit the **Access Key** field by entering an incorrect value and click **Apply Changes**.
4. Edit the **Access Key** field by entering the correct value and click **Apply Changes**.

### All grid node types: Replacing a VMware node

When you recover a failed StorageGRID node that was hosted on VMware, you must remove the failed node and deploy a recovery node.

### What you'll need

You must have determined that the virtual machine cannot 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, Gateway Nodes, and Archive 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.

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.

For instructions:

[Install VMware](#) > Deploying a StorageGRID node as a virtual machine

8. Complete the node recovery procedure, based on the type of node you are recovering.

Type of node	Go to
Primary Admin Node	<a href="#">Configuring the replacement primary Admin Node</a>
Non-primary Admin Node	<a href="#">Selecting Start Recovery to configure a non-primary Admin Node</a>
Gateway Node	<a href="#">Selecting Start Recovery to configure a Gateway Node</a>
Storage Node	<a href="#">Selecting Start Recovery to configure a Storage Node</a>
Archive Node	<a href="#">Selecting Start Recovery to configure an Archive Node</a>

### All grid node types: Replacing a Linux node

If a failure requires that you deploy one or more new physical or virtual hosts or reinstall Linux on an existing host, you must 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®, CentOS, or Debian® deployment. Use the NetApp Interoperability Matrix Tool to get a list of supported versions.

This procedure is only performed as one step in the process of recovering software-based Storage Nodes, primary or non-primary Admin Nodes, Gateway Nodes, or Archive 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.

1. [Deploying new Linux hosts](#)
2. [Restoring grid nodes to the host](#)
3. [What's next: Performing additional recovery steps, if required](#)

### Related information

[NetApp Interoperability Matrix Tool](#)

## Deploying 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.

This procedure includes steps to accomplish the following tasks:

1. Install Linux.
2. Configure the host network.
3. Configure host storage.
4. Install Docker.
5. Install the StorageGRID host service.



Stop after you complete the “Install StorageGRID host service” task in the installation instructions. Do not 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 disk 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 AFF 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.

## Related information

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

[Install Ubuntu or Debian](#)

## Restoring grid nodes to the host

To restore a failed grid node to a new Linux host, you restore the node configuration file using the appropriate commands.

When doing a fresh 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.

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.

### Steps

- [Restoring and validating grid nodes](#)
- [Starting the StorageGRID host service](#)
- [Recovering nodes that fail to start normally](#)

### Restoring and validating 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 recreate 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 on the location of the `/var/local` volume for a node.

### Steps

1. At the command line of the recovered host, list all currently configured StorageGRID grid 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 are not 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 “Performing additional recovery steps, if required.”

3. For grid nodes that do not have a `/var/local` volume, recreate the node’s configuration file to restore it to the host.

Follow the guidelines in “Creating node configuration files” in the installation instructions.



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 recreating the configuration file for a node, be sure to follow all of the guidelines in “Fixing 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.

## Related information

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

[Install Ubuntu or Debian](#)

[Fixing missing network interface errors](#)

[Fixing missing block device errors](#)

## Fixing 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 “Deploying 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.

### Related information

[Deploying new Linux hosts](#)

## Fixing 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 “Deploying 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 recreate 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 from the tables in the “Storage requirements” section of the installation instructions for your Linux operating system. Review the recommendations in “Configuring host storage” before proceeding with the block device replacement.



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. Do not 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
```

#### Related information

[Deploying new Linux hosts](#)

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

[Install Ubuntu or Debian](#)

#### Starting 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.

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
```

For any node that returns a status of Not-Running or Stopped, run the following command:

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

3. 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
```

## Recovering nodes that fail to start normally

If a StorageGRID node does not rejoin the grid normally and does not show up as recoverable, it may be corrupted. You can force the node into recovery mode.

To force the node into recovery mode:

```
sudo storagegrid node force-recovery node-name
```



Before issuing this command, confirm that the node's network configuration is correct; it may have failed to rejoin the grid due to incorrect network interface mappings or an incorrect Grid Network IP address or gateway.



After issuing the `storagegrid node force-recovery node-name` command, you must perform additional recovery steps for *node-name*.

## Related information

[What's next: Performing additional recovery steps, if required](#)

**What's next: Performing 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 may 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	<a href="#">Configuring the replacement primary Admin Node</a>
Non-primary Admin Node	<a href="#">Selecting Start Recovery to configure a non-primary Admin Node</a>

Type of recovery	Next step
Gateway Node	<a href="#">Selecting Start Recovery to configure a Gateway Node</a>
Archive Node	<a href="#">Selecting Start Recovery to configure an Archive Node</a>
Storage Node (software-based): <ul style="list-style-type: none"> <li>• 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></li> <li>• If you had to do a full node reinstall, or you needed to restore <code>/var/local</code></li> </ul>	<a href="#">Selecting Start Recovery to configure a Storage Node</a>
Storage Node (software-based): <ul style="list-style-type: none"> <li>• If you added a new block device.</li> <li>• If, for any <code>&lt;PURPOSE&gt;</code>, the value of the <code>BLOCK_DEVICE_&lt;PURPOSE&gt;</code> configuration file variable refers to a block device that does not contain the same data it did before the host failure.</li> </ul>	<a href="#">Recovering from storage volume failure where the system drive is intact</a>

## Replacing a failed node with a services appliance

You can use an SG100 or SG1000 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.

### What you'll need

- You must have determined that one of the following situations is true:
  - The virtual machine hosting the node cannot 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 must make sure that the StorageGRID Appliance Installer version on the services appliance matches the software version of your StorageGRID system, as described in hardware installation and maintenance for verifying and upgrading the StorageGRID Appliance Installer version.

### [SG100 & SG1000 services appliances](#)



Do not deploy both an SG100 and an SG1000 service appliance in the same site. Unpredictable performance might result.

### About this task

You can use an SG100 or SG1000 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)

### Steps

- [Installing a services appliance \(platform change only\)](#)
- [Preparing an appliance for reinstallation \(platform replacement only\)](#)
- [Starting software installation on a services appliance](#)
- [Monitoring services appliance installation](#)

### Installing a 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 an SG100 or SG1000 services appliance for the replacement node, you must first install the new appliance hardware using the same node name as the failed node.

You must 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.
- **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.

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.

1. Follow the instructions for installing a new SG100 or SG1000 services appliance.
2. When prompted for a node name, use the node name of the failed node.

### Related information

[SG100 & SG1000 services appliances](#)

### Preparing an 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. Do not follow these steps if the failed node was originally hosted on VMware or a Linux host.

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.

### Starting software installation on a services appliance

To install a Gateway Node or Admin Node on an SG100 or SG1000 services appliance, you use the StorageGRID Appliance Installer, which is included on the appliance.

#### What you'll need

- The appliance must be installed in a rack, connected to your networks, and powered on.
- Network links and IP addresses must be 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 must be defined in the Grid Network Subnet List on the primary Admin Node.

For instructions for completing these prerequisite tasks, see the installation and maintenance instructions for an SG100 or SG1000 services appliance.

- You must be using a supported web browser.
- You must know 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 do not need the installation files.

#### About this task

To install StorageGRID software on an SG100 or SG1000 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 SG100 or SG1000 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.

The screenshot shows the 'NetApp StorageGRID Appliance Installer' interface. The top navigation bar includes 'Home', 'Configure Networking', 'Configure Hardware', 'Monitor Installation', and 'Advanced'. The 'Advanced' menu is expanded to show 'Upload StorageGRID Software'. Below the navigation bar, the page title is 'Upload StorageGRID Software'. A paragraph explains that this page is used to upload the StorageGRID software installation package. Below this, there are two sections: 'Current StorageGRID Installation Software' and 'Upload StorageGRID Installation Software'. The 'Current StorageGRID Installation Software' section shows 'Version: None' and 'Package Name: None'. The 'Upload StorageGRID Installation Software' section has two rows: 'Software Package' with a 'Browse' button, and 'Checksum File' with a 'Browse' button.

- 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:

- a. In the This Node section, for **Node Type**, select **Gateway** or **Non-Primary Admin**, depending on the type of node you are restoring.
- 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 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	<ol style="list-style-type: none"> <li>a. Unselect the <b>Enable Admin Node discovery</b> check box.</li> <li>b. Enter the IP address manually.</li> <li>c. Click <b>Save</b>.</li> <li>d. Wait while the connection state for the new IP address becomes “ready.”</li> </ol>
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"> <li>a. Select the <b>Enable Admin Node discovery</b> check box.</li> <li>b. From the list of discovered IP addresses, select the primary Admin Node for the grid where this services appliance will be deployed.</li> <li>c. Click <b>Save</b>.</li> <li>d. Wait while the connection state for the new IP address becomes “ready.”</li> </ol>

4. 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 installation and maintenance instructions for your appliance.

5. 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.

#### Related information

[SG100 & SG1000 services appliances](#)

#### Monitoring services appliance installation

The StorageGRID Appliance Installer provides status until installation is complete. When the software installation is complete, the appliance is rebooted.

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		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 are not re-run. If you are re-running an installation, any tasks that do not 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	<div style="width: 25%; background-color: #00a0e3; border: 1px solid #ccc;"></div>	Do not refresh. You will be redirected when the installer is ready

4. Go on 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	<a href="#">Selecting Start Recovery to configure a Gateway Node</a>
Non-primary Admin Node	<a href="#">Selecting Start Recovery to configure a non-primary Admin Node</a>
Primary Admin Node	<a href="#">Configuring the replacement primary Admin Node</a>

## How site recovery is performed by technical support

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, it is important to 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 policy:** The number, type, and location of object copies in your grid is controlled by your active ILM policy. The specifics of your ILM policy 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 level 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 your consistency level 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 policy. Technical support must assess the recent history of your grid as well as its current situation before beginning a site recovery.

## Overview of site recovery

This is a general overview of the process that technical support uses to recover a failed site.



Site recovery can only be performed by technical support.



**Caution:** Do not use the recovery procedures designed for a single failed Storage Node. Data loss will occur.

### 1. Contact technical support.

Technical support does a detailed assessment of the failure and works with you to review your business objectives. Based on this information, technical support develops a recovery plan tailored for your situation.

### 2. Technical support recovers the primary Admin Node if it has failed.

### 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, specialized commands are required to successfully restore objects and object metadata.

4. Technical support recovers other failed nodes.

After object metadata and data have been recovered, failed Gateway Nodes, non-primary Admin Nodes, or Archive Nodes can be recovered using standard procedures.

#### Related information

[Site decommissioning](#)

## Decommission procedure

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 **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 **connected site decommission** to remove a site in which all nodes are connected to StorageGRID.
- Perform a **disconnected site decommission** to remove a site in which all nodes are disconnected from StorageGRID.



Before performing a disconnected site decommission, 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.

If a site contains a mixture of connected (✓) and disconnected nodes (⬜ or ⬛), you must bring all offline nodes back online.

#### Related information

[Grid node decommissioning](#)

[Site decommissioning](#)

### Grid node decommissioning

You can use the node decommission procedure to remove one or more Storage Nodes, Gateway Nodes, or non-primary Admin Nodes at one or more sites. You cannot decommission the primary Admin Node or an Archive Node.


In general, you should decommission grid nodes only while they are connected to the StorageGRID system

and all nodes are in normal health (have green icons on the **Nodes** pages and on the **Decommission Nodes** page). However, if required, you can decommission a grid node that is disconnected. Before removing a disconnected node, make sure you understand the implications and restrictions of that process.

Use the node decommission procedure when any of the following are true:

- You have added a larger Storage Node to the system and you want to remove one or more smaller Storage Nodes, while at the same time preserving objects.
- 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 cannot recover or bring back online.

The flowchart shows the high-level steps for decommissioning grid nodes.



### Steps

- Preparing to decommission grid nodes
- Gathering required materials
- Accessing the Decommission Nodes page
- Decommissioning disconnected grid nodes
- Decommissioning connected grid nodes
- Pausing and resuming the decommission process for Storage Nodes
- Troubleshooting node decommissioning

## Preparing to decommission grid nodes

You must review the considerations for removing grid nodes and confirm no repair jobs are active for erasure-coded data.

### Steps

- [Considerations for decommissioning Storage Nodes](#)
- [Checking data repair jobs](#)

## Considerations for decommissioning grid nodes

Before you start this procedure to decommission one or more nodes, you must understand the implications of removing each type of node. Upon the successful decommissioning of a node, its services will be disabled and the node will be automatically shut down.

You cannot decommission a node if doing so will leave the StorageGRID in an invalid state. The following rules are enforced:

- You cannot decommission the primary Admin Node.
- You cannot decommission Archive Nodes.
- You cannot decommission an Admin Node or a Gateway Node if one of its network interfaces is part of a high availability (HA) group.
- You cannot decommission a Storage Node if its removal would affect the ADC quorum.
- You cannot decommission a Storage Node if it is required for the active ILM policy.
- You should not decommission more than 10 Storage Nodes in a single Decommission Node procedure.
- You cannot decommission a connected node if your grid includes any disconnected nodes (nodes whose health is Unknown or Administratively Down). You must decommission or recover the disconnected nodes first.
- If your grid contains multiple disconnected nodes, the software requires you to decommission them at all the same time, which increases the potential for unexpected results.
- If a disconnected node cannot be removed (for example, a Storage Node that is required for the ADC quorum), no other disconnected node can be removed.
- If you want to replace an older appliance with a newer appliance, consider using the appliance node cloning procedure instead of decommissioning the old node and adding the new node in an expansion.

### [Appliance node cloning](#)



Do not remove a grid node's virtual machine or other resources until instructed to do so in decommission procedures.

## Considerations for decommissioning Admin Nodes or a Gateway Nodes

Review the following considerations before decommissioning an Admin Node or a Gateway Node.

- The decommission procedure requires exclusive access to some system resources, so you must confirm

that no other maintenance procedures are running.

- You cannot decommission the primary Admin Node.
- You cannot decommission an Admin Node or 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 administering StorageGRID.
- As required, you can safely change the ILM policy while decommissioning a Gateway Node or 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).

## Related information

[Administer StorageGRID](#)

## Considerations for decommissioning Storage Nodes

If you plan to decommission a Storage Node, you must understand how StorageGRID manages the object data and metadata on that node.

The following considerations and restrictions apply when decommissioning Storage Nodes:

- 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.
- If the Storage Node is disconnected when you decommission it, the system must reconstruct the data using data from the connected Storage Nodes, which can result in data loss.
- When you remove a Storage Node, large volumes of object data must be transferred over the network. Although these transfers should not affect normal system operations, they can have an impact on 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.
- Decommission procedures that involve Storage Nodes can be paused during certain stages to allow other maintenance procedures to run if needed, and resumed once they are complete.
- You cannot run data repair operations on any grid nodes when a decommission task is running.
- You should not make any changes to the ILM policy while a Storage Node is being decommissioned.
- When you remove a Storage Node, data on the node is migrated to other grid nodes; however, this data is not completely removed from the decommissioned grid node. To permanently and securely remove data, you must wipe the decommissioned grid node's drives after the decommission procedure is complete.
- When you decommission a Storage Node, the following alerts and alarms might be raised and you might receive related email and SNMP notifications:

- **Unable to communicate with node** alert. This alert is triggered when you decommission a Storage Node that includes the ADC service. The alert is resolved when the decommission operation completes.
- VSTU (Object Verification Status) alarm. This notice-level alarm indicates that the Storage Node is going into maintenance mode during the decommission process.
- CASA (Data Store Status) alarm. This major-level alarm indicates that the Cassandra database is going down because services have stopped.

### Related information

[Restoring object data to a storage volume, if required](#)

[Understanding the ADC quorum](#)

[Reviewing the ILM policy and storage configuration](#)

[Decommissioning disconnected Storage Nodes](#)

[Consolidating Storage Nodes](#)

[Decommissioning multiple Storage Nodes](#)

### Understanding the ADC quorum

You might not be able to decommission certain Storage Nodes at a data center site if too few Administrative Domain Controller (ADC) services would remain after the decommissioning. This 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 cannot 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 data center site must have the ADC service. If a data center site has more than three Storage Nodes with the ADC service, a simple majority of those nodes must remain available after the decommissioning ( $((0.5 * \text{Storage Nodes with ADC}) + 1)$ ).

For example, suppose a data center 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, you must add a new Storage Node in an expansion and specify that it should have an ADC service. Then, you can decommission the existing Storage Node.

### Related information

[Expand your grid](#)

## Reviewing the 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 the 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 "Consolidating 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 policy?



An ILM policy that cannot be satisfied will lead to backlogs and alarms, and can 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 factors listed in the table.

Area to assess	Notes
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?

Area to assess	Notes
Storage type	Will there be enough storage of the appropriate type after decommissioning is complete? For example, ILM rules might dictate that content be moved from one type of storage to another as content ages. If so, you must ensure that enough storage of the appropriate type is available in the final configuration of the StorageGRID system.

### Related information

[Consolidating Storage Nodes](#)

[Manage objects with ILM](#)

[Expand your grid](#)

### Decommissioning disconnected Storage Nodes

You must understand what can happen if you decommission a Storage Node while it is disconnected (health is Unknown or Administratively Down).

When you decommission a Storage Node that is disconnected from the grid, StorageGRID uses data from other Storage Nodes to reconstruct the object data and metadata that was on the disconnected node. It does this by automatically starting data repair jobs at the end of the decommissioning process.

Before decommissioning a disconnected Storage Node, be aware of the following:

- You should never decommission a disconnected node unless you are sure it cannot be brought online or recovered.



Do not 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.

- If a disconnected Storage Node contains the only copy of an object, that object will be lost when you decommission the node. 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.
- 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 are not monitored by the decommission procedure. You must manually monitor these jobs and restart them as needed. See the instructions on monitoring data repair.

#### Checking data repair jobs

- If you decommission more than one disconnected Storage Node at a time, data loss might occur. The system might not be able to reconstruct data if too few copies of object data, metadata, or erasure-coded fragments remain available.



If you have more than one disconnected Storage Node that you cannot recover, contact technical support to determine the best course of action.

## Consolidating 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 to add 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.

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.

### Related information

[Expand your grid](#)

## Decommissioning multiple Storage Nodes

If you need to remove more than one Storage Node, you can decommission them either sequentially or in parallel.

- 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.

## Checking 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.

If you need to decommission a disconnected Storage Node, you will also complete these steps after the decommission procedure completes in order 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.

1. Log in to the primary Admin Node:

- a. Enter the following command: `ssh admin@grid_node_IP`

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

- 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. 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 do not 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. Go to the next step.

```

root@DC1-ADM1:~ # repair-data show-ec-repair-status

Repair ID Scope Start Time End Time State Est/Affected Bytes Repaired
Retry Repair
=====
=====
949283 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:27:06.9 Success 17359
17359 No
949292 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:37:06.9 Failure 17359 0
Yes
949294 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:47:06.9 Failure 17359 0
Yes
949299 DC1-S-99-10 (Volumes: 1,2) 2016-11-30T15:57:06.9 Failure 17359 0
Yes

```

- 3. If the State for all repairs is `Success`, you do not need to restart any repair jobs.
- 4. If the State for any repair is `Failure`, you must restart that repair.
  - a. Obtain the repair ID for the failed repair from the output.
  - b. 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`

- c. Continue to track the status of EC data repairs until the State for all repairs is `Success`.

**Gathering 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.

Item	Notes
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](#)

[Downloading the Recovery Package](#)

### Accessing the 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.

### Steps

1. Select **Maintenance > Maintenance Tasks > Decommission**.

The Decommission page appears.

#### Decommission

Select **Decommission Nodes** to remove one or more nodes from a single site. Select **Decommission Site** to remove an entire data center site.

Learn important details about removing grid nodes and sites in the "Decommission procedure" section of the [recovery and maintenance instructions](#).



2. Click the **Decommission Nodes** button.

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.  
For example, this page shows all grid nodes in a single data center. The Decommission Possible column indicates that you can decommission the non-primary Admin Node, the Gateway Node, and two of the five Storage Nodes.

### Decommission Nodes

Before decommissioning a grid node, review the health of all nodes. If possible, resolve any issues or alarms before proceeding.

Select the checkbox for each grid node you want to decommission. If decommission is not possible for a node, see the Recovery and Maintenance Guide to learn how to proceed.

#### Grid Nodes

Search <input type="text" value=""/>							
	Name	Site	Type	Has ADC	Health	Decommission Possible	
	DC1-ADM1	Data Center 1	Admin Node	-		No, primary Admin Node decommissioning is not supported.	
<input type="checkbox"/>	DC1-ADM2	Data Center 1	Admin Node	-			
<input type="checkbox"/>	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.	
<input type="checkbox"/>	DC1-S4	Data Center 1	Storage Node	No			
<input type="checkbox"/>	DC1-S5	Data Center 1	Storage Node	No			

#### Passphrase

Provisioning  
Passphrase

[Start Decommission](#)

### 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-most column includes a check box. If a node cannot be decommissioned, this column describes the issue. If there is more than one reason a node cannot be decommissioned, the most critical reason is shown.

Decommission Possible reason	Description	Steps to resolve
No, node type decommissioning is not supported.	You cannot decommission the primary Admin Node or an Archive Node.	None.

Decommission Possible reason	Description	Steps to resolve
<p>No, at least one grid node is disconnected.</p> <p><b>Note:</b> This message is shown for connected grid nodes only.</p>	<p>You cannot decommission a connected grid node if any grid node is disconnected.</p> <p>The <b>Health</b> column includes one of these icons for grid nodes that are disconnected:</p> <ul style="list-style-type: none"> <li>•  (gray): Administratively Down</li> <li>•  (blue): Unknown</li> </ul>	<p>Go to the <a href="#">step that lists the decommission procedure choices</a>.</p>
<p>No, one or more required nodes is currently disconnected and must be recovered.</p> <p><b>Note:</b> This message is shown for disconnected grid nodes only.</p>	<p>You cannot 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).</p>	<ol style="list-style-type: none"> <li>a. Review the Decommission Possible messages for all disconnected nodes.</li> <li>b. Determine which nodes cannot be decommissioned because they are required. <ul style="list-style-type: none"> <li>◦ If the Health of a required node is Administratively Down, bring the node back online.</li> <li>◦ If the health of a required node is Unknown, perform a node recovery procedure to recover the required node.</li> </ul> </li> </ol>
<p>No, member of HA group(s): <math>x</math>. Before you can decommission this node, you must remove it from all HA groups.</p>	<p>You cannot decommission an Admin Node or a Gateway Node if a node interface belongs to a high availability (HA) group.</p>	<p>Edit the HA group to remove the node's interface or remove the entire HA group. See the instructions for administering StorageGRID.</p>
<p>No, site <math>x</math> requires a minimum of <math>n</math> Storage Nodes with ADC services.</p>	<p><b>Storage Nodes only.</b> You cannot decommission a Storage Node if insufficient nodes would remain at the site to support ADC quorum requirements.</p>	<p>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.</p>

Decommission Possible reason	Description	Steps to resolve
<p>No, one or more Erasure Coding profiles need at least <math>n</math> Storage Nodes. If the profile is not used in an ILM rule, you can deactivate it.</p>	<p><b>Storage Nodes only.</b> You cannot decommission a Storage Node unless enough nodes would remain for the existing Erasure Coding profiles.</p> <p>For example, if an Erasure Coding profile exists for 4+2 erasure coding, at least 6 Storage Nodes must remain.</p>	<p>For each affected Erasure Coding profile, perform one of the following steps, based on how the profile is being used:</p> <ul style="list-style-type: none"> <li>• <b>Used in the active ILM policy:</b> Perform an expansion. Add enough new Storage Nodes to allow erasure coding to continue. See the instructions for expanding StorageGRID.</li> <li>• <b>Used in an ILM rule but not in the active ILM policy:</b> Edit or delete the rule and then deactivate the Erasure Coding profile.</li> <li>• <b>Not used in any ILM rule:</b> Deactivate the Erasure Coding profile.</li> </ul> <p><b>Note:</b> 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.</p> <p>Learn about deactivating an Erasure Coding profile in the instructions for managing objects with information lifecycle management.</p>

4. If decommissioning is possible for the node, determine which procedure you need to perform:

If your grid includes...	Go to...
Any disconnected grid nodes	<a href="#">Decommissioning disconnected grid nodes</a>
Only connected grid nodes	<a href="#">Decommissioning connected grid nodes</a>

#### Related information

[Checking data repair jobs](#)

[Understanding the ADC quorum](#)

[Manage objects with ILM](#)

[Expand your grid](#)

[Administer StorageGRID](#)

#### Decommissioning 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).

## What you'll need

- You understand the requirements and considerations for decommissioning grid nodes.

### Considerations for decommissioning grid nodes

- You have obtained all prerequisite items.
- You have ensured that no data repair jobs are active.

### Checking 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 must have the provisioning passphrase.

You can identify disconnected nodes by looking for Unknown (blue) or Administratively Down (gray) icons in the **Health** column. In the example, the Storage Node named DC1-S4 is disconnected; all of the other nodes are connected.

## Decommission Nodes

Before decommissioning a grid node, review the health of all nodes. If possible, resolve any issues or alarms before proceeding.

 A grid node is disconnected (has a blue or gray health icon). Try to bring it back online or recover it. Data loss might occur if you decommission a node that is disconnected.

See the [Recovery and Maintenance Guide](#) for details. Contact Support if you cannot recover a node and do not want to decommission it.

Select the checkbox for each grid node you want to decommission. If decommission is not possible for a node, see the [Recovery and Maintenance Guide](#) to learn how to proceed.

### Grid Nodes

Name	Site	Type	Has ADC	Health	Decommission Possible
DC1-ADM1	Data Center 1	Admin Node	-		No, primary Admin Node decommissioning is not supported.
DC1-ADM2	Data Center 1	Admin Node	-		No, at least one grid node is disconnected.
DC1-G1	Data Center 1	API Gateway Node	-		No, at least one grid node is disconnected.
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.
<input type="checkbox"/> DC1-S4	Data Center 1	Storage Node	No		

### Passphrase

Provisioning  
Passphrase

Start Decommission

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 at all the same time, which increases the potential for unexpected results.



Be very careful when decommissioning more than one disconnected grid node at a time, especially if you are selecting multiple disconnected Storage Nodes.

- If a disconnected node cannot be removed (for example, a Storage Node that is required for the ADC quorum), no other disconnected node can be removed.

Before decommissioning a disconnected **Storage Node**, note the following

- You should never decommission a disconnected Storage Node unless you are sure it cannot be brought online or recovered.



If you believe that object data can still be recovered from the node, do not perform this procedure. Instead, contact technical support to determine if node recovery is possible.

- If you decommission more than one disconnected Storage Node, data loss might occur. The system might not be able to reconstruct data if not enough object copies, erasure-coded fragments, or object metadata remain available.



If you have more than one disconnected Storage Node that you cannot 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 are not monitored by the decommission procedure. You must manually monitor these jobs and restart them as needed. See the instructions on monitoring data repair.

### Checking 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.

Before decommissioning a disconnected **Admin Node** or **Gateway Node**, note the following:

- 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.
- You can safely decommission a Gateway Node while it is disconnected.

### Steps

1. Attempt to bring any disconnected grid nodes back online or to recover them.

See the 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 check box for that node.



If your grid contains multiple disconnected nodes, the software requires you to decommission them at all the same time, which increases the potential for unexpected results.



Be very careful when selecting 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 cannot 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.

## Warning

The selected nodes are disconnected (health is Unknown or Administratively Down). If you continue and the node has the only copy of an object, the object will be lost when the node is removed.

The following grid nodes have been selected for decommissioning and will be permanently removed from the StorageGRID Webscale system.

DC1-S4

Do you want to continue?

Cancel

OK

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.

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.

The progress for each node is displayed while the decommission procedure is running. When all tasks are complete, the node selection list is redisplayed.

Name	Type	Progress	Stage
DC1-S4	Storage Node	<div style="width: 10%;"></div>	Prepare Task

Pause Resume

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.

#### Decommission Nodes

The previous decommission procedure completed successfully.

**i** Repair jobs for replicated and erasure-coded data have been started. These jobs restore object data that might have been on any disconnected Storage Nodes. To monitor the progress of these jobs and restart them as needed, see the Decommissioning section of the Recovery and Maintenance Guide.

Before decommissioning a grid node, review the health of all nodes. If possible, resolve any issues or alarms before proceeding.

Select the checkbox for each grid node you want to decommission. If decommission is not possible for a node, see the Recovery and Maintenance Guide to learn how to proceed.

#### Grid Nodes

Name	Site	Type	Has ADC	Health	Decommission Possible
DC1-ADM1	Data Center 1	Admin Node	-		No, primary Admin Node decommissioning is not supported.
<input type="checkbox"/> DC1-ADM2	Data Center 1	Admin Node	-		
<input type="checkbox"/> 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.

#### Passphrase

Provisioning  
Passphrase

- 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.



Do not perform this step until the nodes have shut down automatically.

- If you are decommissioning a Storage Node, monitor the status of the data repair jobs that are automatically started during the decommissioning process.
  - Select **Support > Tools > Grid Topology**.
  - Select **StorageGRID deployment** at the top of the Grid Topology tree.
  - On the Overview tab, locate the ILM Activity section.
  - 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 are not 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 do not 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.

e. Use the following commands to track or restart repairs:

- Use the `repair-data show-ec-repair-status` command to track repairs of erasure coded data.
- Use the `repair-data start-ec-node-repair` command with the `--repair-id` option to restart a failed repair.  
See the instructions for checking data repair jobs.

10. Continue to track the status of EC data repairs until all repair jobs have completed successfully.

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.

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.

[SG100 & SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

## Related information

[Grid node recovery procedures](#)

[Downloading the Recovery Package](#)

[Checking data repair jobs](#)

## Decommissioning connected grid nodes

You can decommission and permanently remove nodes that are connected to the grid.

### What you'll need

- You understand the requirements and considerations for decommissioning grid nodes.

[Considerations for decommissioning grid 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, 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.
- You have the provisioning passphrase.
- Grid nodes are connected.
- The **Decommission Possible** column for the node or nodes you want to decommission include a green checkmark.
- 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 the instructions for checking data repair jobs.



Do not remove a grid node's virtual machine or other resources until instructed to do so in this procedure.

### Steps

1. From the Decommission Nodes page, select the check box for each grid node you want to decommission.
2. Enter the provisioning passphrase.

The **Start Decommission** button is enabled.

3. Click **Start Decommission**.

A confirmation dialog box appears.

## Info

The following grid nodes have been selected for decommissioning and will be permanently removed from the StorageGRID Webscale system.

DC1-S5

Do you want to continue?

Cancel

OK

4. Review the list of selected nodes, and click **OK**.

The node decommission procedure starts, and the progress is displayed for each node. During the procedure, a new Recovery Package is generated to show the grid configuration change.

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.

The progress for each node is displayed while the decommission procedure is running. When all tasks are complete, the node selection list is redisplayed.

Name	Type	Progress	Stage
DC1-S5	Storage Node	<div style="width: 10%;"></div>	Prepare Task



Do not 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, 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.

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.

## Decommission Nodes

The previous decommission procedure completed successfully.

Before decommissioning a grid node, review the health of all nodes. If possible, resolve any issues or alarms before proceeding.

Select the checkbox for each grid node you want to decommission. If decommission is not possible for a node, see the Recovery and Maintenance Guide to learn how to proceed.

### Grid Nodes

	Name	Site	Type	Has ADC	Health	Decommission Possible
	DC1-ADM1	Data Center 1	Admin Node	-		No, primary Admin Node decommissioning is not supported.
<input checked="" type="checkbox"/>	DC1-ADM2	Data Center 1	Admin Node	-		
<input checked="" type="checkbox"/>	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.

### Passphrase

Provisioning  
Passphrase

## 7. 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.
- **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.

Complete these steps after you complete the node 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 use the appliance in another grid.

[SG100 & SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

## Related information

[Checking data repair jobs](#)

[Downloading the Recovery Package](#)

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

### Pausing and resuming the decommission process for Storage Nodes

If necessary, you can pause the decommission procedure for a Storage Node during certain stages. You must pause decommissioning on a Storage Node before you can start a second maintenance procedure. After the other procedure is finished, you can resume decommissioning.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.

#### Steps

1. Select **Maintenance > Maintenance Tasks > Decommission**.

The Decommission page appears.

2. Click **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. Click **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.

Name	Type	Progress	Stage
DC1-S5	Storage Node	<div style="width: 50%; background-color: #f4a460;"></div>	Evaluating ILM

4. After the other maintenance procedure is finished, click **Resume** to proceed with the decommission.

## Troubleshooting node decommissioning

If the node decommission procedure stops because of an error, you can take specific steps to troubleshoot the problem.

### What you'll need

You must be signed in to the Grid Manager using a supported browser.

### About this task

If you shut down the grid node being decommissioned, the task stops until the grid node is restarted. The grid node must be online.

### Steps

1. Select **Support > Tools > Grid Topology**.
2. In the Grid Topology tree, expand each Storage Node entry, and verify that the DDS and LDR services are both online.

To perform Storage Node decommissioning, the StorageGRID system's DDS services (hosted by Storage Nodes) must be online. This is a requirement of the ILM re-evaluation.

3. To view the active grid tasks, select **primary Admin Node > CMN > Grid Tasks > Overview**.
4. Check the status of the decommissioning grid task.
  - a. If the status of the decommissioning grid task indicates a problem with saving grid task bundles, select **primary Admin Node > CMN > Events > Overview**
  - b. Check the number of Available Audit Relays.

If the attribute Available Audit Relay is one or greater, the CMN service is connected to at least one ADC service. ADC services act as Audit Relays.


The CMN service must be connected to at least one ADC service and a majority (50 percent plus one) of the StorageGRID system's ADC services must be available in order for a grid task to move from one stage of decommissioning to another and finish.

- c. If the CMN service is not connected to enough ADC services, ensure that Storage Nodes are online, and check network connectivity between the primary Admin Node and Storage Nodes.

## Site decommissioning

You might need to remove a data center site from the StorageGRID system. To remove a site, you must decommission it.

The flowchart shows the high-level steps for decommissioning a site.



**Steps**

- Considerations for removing a site
- Gathering required materials

- [Step 1: Select Site](#)
- [Step 2: View Details](#)
- [Step 3: Revise ILM Policy](#)
- [Step 4: Remove ILM References](#)
- [Step 5: Resolve Node Conflicts \(and start decommission\)](#)
- [Step 6: Monitor Decommission](#)

### 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 cannot decommission a site that includes the primary Admin Node.
- You cannot decommission a site that includes an Archive Node.
- You cannot 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 cannot decommission a site if it contains a mixture of connected (✓) and disconnected (🔒 or 🚫) nodes.
- You cannot decommission a site if any node at any other site is disconnected (🔒 or 🚫).
- You cannot start the site decommission procedure if an ec-node-repair operation is in progress. See the following topic to track repairs of erasure-coded data.

### Checking data repair jobs

- While the site decommission procedure is running:
  - You cannot create ILM rules that refer to the site being decommissioned. You also cannot edit an existing ILM rule to refer to the site.
  - You cannot 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.
- You cannot 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 the active ILM policy. If it is, you must create and activate a new ILM policy with new ILM rules.
- No proposed ILM policy exists. If you have a proposed policy, you must delete it.

- No ILM rules refer to the site, even if those rules are not used in the active or proposed 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. The system-default All Storage Nodes storage pool 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 the instructions for managing objects with information lifecycle management.

## 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 policy, 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 add storage volumes or Storage Nodes to an existing site or add a new site before you perform this procedure. See the instructions for expanding a StorageGRID system.

- 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 cannot 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 perform an expansion procedure 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 cannot 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 the active ILM policy. In some cases, you might need to expand your StorageGRID system 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.

- All nodes at the site must have a Connection State of one of the following:
  - **Unknown** (🔵): The node is not connected to the grid for an unknown reason. For example, the network connection between nodes has been lost or the power is down.
  - **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 cannot be retrieved.

### Considerations for consistency controls when you remove a site

The consistency level for an S3 bucket or Swift container determines whether StorageGRID fully replicates object metadata to all nodes and sites before telling a client that object ingest was successful. The consistency level makes a trade-off 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 level 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.

#### Related information

[How site recovery is performed by technical support](#)

[Manage objects with ILM](#)

[Expand your grid](#)

#### Gathering 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.
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.

Item	Notes
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](#)

[Downloading the Recovery Package](#)

#### Step 1: Select Site

To determine if a site can be decommissioned, start by accessing the Decommission Site wizard.

#### What you'll need

- You must have obtained all required materials.
- You must have reviewed the considerations for removing a site.
- You must be signed in to the Grid Manager using a supported browser.
- You must have the Root Access permission, or the Maintenance and ILM permissions.

#### Steps

1. Select **Maintenance > Maintenance Tasks > Decommission**.

The Decommission page appears.

#### Decommission

Select **Decommission Nodes** to remove one or more nodes from a single site. Select **Decommission Site** to remove an entire data center site.


Learn important details about removing grid nodes and sites in the "Decommission procedure" section of the [recovery and maintenance instructions](#).



2. Select the **Decommission Site** button.

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



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
<input type="radio"/>	Raleigh	3.93 MB	
<input type="radio"/>	Sunnyvale	3.97 MB	
	Vancouver	3.90 MB	No. This site contains the primary Admin Node.

Next

3. 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.

4. 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 cannot be decommissioned, the most critical reason is shown.

Decommission Possible reason	Description	Next step
Green checkmark ()	You can decommission this site.	Go to <a href="#">the next step</a> .
No. This site contains the primary Admin Node.	You cannot decommission a site containing the primary Admin Node.	None. You cannot perform this procedure.

Decommission Possible reason	Description	Next step
No. This site contains one or more Archive Nodes.	You cannot decommission a site containing an Archive Node.	None. You cannot perform this procedure.
No. All nodes at this site are disconnected. Contact your NetApp account representative.	You cannot perform a connected site decommission unless every node in the site is connected (✓).	<p>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.</p> <p><b>IMPORTANT:</b> Never take online nodes offline so that you can remove a site. You will lose data.</p>

The example shows a StorageGRID system with three sites. The green checkmark (✓) for the Raleigh and Sunnyvale sites indicates that you can decommission those sites. However, you cannot 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.

### What you'll need

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 cannot 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.

### 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 and Swift 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  
Used Space: 3.93 MB

Free Space: 475.38 GB  
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 cannot decommission an Admin Node or a Gateway Node that is used in a HA

group. Before you start the decommission, you must edit HA groups to remove all nodes at the site. Or, you can remove the HA group if it only includes nodes from this site.

### Administer StorageGRID

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.

#### Related information

[Manage objects with ILM](#)

#### Step 3: Revise ILM Policy

From Step 3 (Revise ILM Policy) of the Decommission Site wizard, you can determine if the site is referred to by the active ILM policy.

#### What you'll need

You have a good understanding of how ILM works and you are familiar with creating storage pools, Erasure Coding profiles, ILM rules, and simulating and activating an ILM policy.

[Manage objects with ILM](#)

#### About this task

StorageGRID cannot decommission a site if that site is referred to by any ILM rule in the active ILM policy.

If your current ILM policy refers to the site you want to remove, you must activate a new ILM policy that meets certain requirements. Specifically, the new ILM policy:

- Cannot use a storage pool that refers to the site.
- Cannot use an Erasure Coding profile that refers to the site.
- Cannot use the default **All Storage Nodes** storage pool or the default **All Sites** site.
- Cannot use the stock **Make 2 Copies** rule.
- Must be designed to fully protect all object data.



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 are 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, the 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 the new policy.

### Steps

1. From Step 3 (Revise ILM Policy), determine if any ILM rules in the active ILM policy refer to the site you selected to remove.

## Decommission Site



If your current ILM policy refers to the site, you must activate a new policy before you can go to the next step.

The new ILM policy:

- Cannot use a storage pool that refers to the site.
- Cannot use an Erasure Coding profile that refers to the site.
- Cannot use the default **All Storage Nodes** storage pool or the default **All Sites** site.
- Cannot use the **Make 2 Copies** rule.
- Must be designed to fully protect all object data after one site is removed.

Contact technical support if you need assistance in designing the new policy.

If you are performing a connected site decommission, StorageGRID will begin to remove object data from the site as soon as you activate the new ILM policy. Moving or deleting all object copies might take weeks, but you can safely start a site decommission while object data still exists at the site.

### Rules Referring to Raleigh in the Active ILM Policy

The table lists the ILM rules in the active ILM policy that refer to the site.

- If no ILM rules are listed, the active ILM policy does not refer to the site. Select **Next** to go to Step 4 (Remove ILM References).
- If one or more ILM rules are listed, you must create and activate a new policy that does not use these rules.

Active Policy Name: [Data Protection for Three Sites](#)

The active ILM policy refers to Raleigh. Before you can remove this site, you must propose and activate a new policy.

Name	EC Profiles	Storage Pools
3 copies for S3 tenant	—	Raleigh storage pool
2 copy 2 sites for smaller objects	—	Raleigh storage pool
EC for larger objects	three site EC profile	All 3 Sites

Previous

Next

2. If no rules are listed, select **Next** to go to Step 4 (Remove ILM References)

### Step 4: Remove ILM References

3. If one or more ILM rules are listed in the table, select the link next to **Active Policy Name**.

The ILM Policies page appears in a new browser tab. Use this tab to update ILM. The Decommission Site page will remain open on the other tab.

- a. If necessary, select **ILM > Storage Pools** to create one or more storage pools that do not refer to the site.



For details, see the instructions for managing objects with information lifecycle management.

- b. If you plan to use erasure coding, select **ILM > Erasure Coding** to create one or more Erasure Coding profiles.

You must select storage pools that do not refer to the site.



Do not use the **All Storage Nodes** storage pool in the Erasure Coding profiles.

4. Select **ILM > Rules** and clone each of the rules listed in the table for Step 3 (Revise ILM Policy).



For details, see the instructions for managing objects with information lifecycle management.

- a. Use names that will make it easy to select these rules in a new policy.
- b. Update the placement instructions.

Remove any storage pools or Erasure Coding profiles that refer to the site and replace them with new storage pools or Erasure Coding profiles.



Do not use the **All Storage Nodes** storage pool in the new rules.

5. Select **ILM > Policies** and create a new policy that uses the new rules.



For details, see the instructions for managing objects with information lifecycle management.

- a. Select the active policy, and select **Clone**.
- b. Specify a policy name and a reason for change.
- c. Select rules for the cloned policy.
  - Unselect all rules listed for Step 3 (Revise ILM Policy) of the Decommission Site page.
  - Select a default rule that does not refer to the site.



Do not select the **Make 2 Copies** rule because that rule uses the **All Storage Nodes** storage pool, which is not allowed.

- Select the other replacement rules you created. These rules should not refer to the site.

## Select Rules for Policy

### Select Default Rule

This list shows the rules that do not use any filters. Select one rule to be the default rule for the policy. The default rule applies to any objects that do not match another rule in the policy and is always evaluated last. The default rule should retain objects forever.

	Rule Name
<input checked="" type="radio"/>	2 copies at Sunnyvale and Vancouver for smaller objects
<input type="radio"/>	2 copy 2 sites for smaller objects
<input type="radio"/>	Make 2 Copies

### Select Other Rules

The other rules in a policy are evaluated before the default rule and must use at least one filter. Each rule in this list uses at least one filter (tenant account, bucket name, or an advanced filter, such as object size).

	Rule Name	Tenant Account
<input type="checkbox"/>	3 copies for S3 tenant	S3 (61659555232085399385)
<input type="checkbox"/>	EC for larger objects	—
<input checked="" type="checkbox"/>	1-site EC for larger objects	—
<input checked="" type="checkbox"/>	2 copies for S3 tenant	S3 (61659555232085399385)

Cancel

Apply

d. Select **Apply**.

e. Drag and drop the rows to reorder the rules in the policy.

You cannot move the default rule.



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.

f. Save the proposed policy.

6. Ingest test objects, and simulate the proposed 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.

7. Activate the new policy.

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).

- Return to **Step 3 (Revise ILM Policy)** to ensure that no ILM rules in the new active policy refer to the site and the **Next** button is enabled.

#### Rules Referring to Raleigh in the Active ILM Policy

The table lists the ILM rules in the active ILM policy that refer to the site.

- If no ILM rules are listed, the active ILM policy does not refer to the site. Select **Next** to go to Step 4 (Remove ILM References).
- If one or more ILM rules are listed, you must create and activate a new policy that does not use these rules.

Active Policy Name: [Data Protection for Two Sites](#) 

No ILM rules in the active ILM policy refer to Raleigh.

Previous

Next



If any rules are listed, you must create and activate a new ILM policy before you can continue.

- If no rules are listed, 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 can remove the proposed policy if one exists and delete or edit any unused ILM rules that still refer to the site.

#### About this task

You are prevented from starting the site decommission procedure in these cases:

- A proposed ILM policy exists. If you have a proposed policy, you must delete it.
- Any ILM rule refers to the site, even if that rule is not used in any ILM policy. You must delete or edit all rules that refer to the site.

#### Steps

- If a proposed policy is listed, remove it.

## Decommission Site



Before you can decommission a site, you must ensure that no proposed ILM policy exists and that no ILM rules refer to the site, even if those rules are not currently used in an ILM policy.

**Proposed policy exists** ▲

You must delete the proposed policy before you can start the site decommission procedure.

Policy name: [Data Protection for Two Sites \(v2\)](#)  [Delete Proposed Policy](#)

**4 ILM rules** refer to Raleigh ▼

**1 Erasure Coding profile** will be deactivated ▼

**3 storage pools** will be deleted ▼

[Previous](#) [Next](#)

- a. Select **Delete Proposed Policy**.
  - b. Select **OK** in the confirmation dialog box.
2. Determine whether any unused ILM rules refer to the site.

## Decommission Site



Before you can decommission a site, you must ensure that no proposed ILM policy exists and that no ILM rules refer to the site, even if those rules are not currently used in an ILM policy.

No proposed policy exists

**4 ILM rules** refer to Data Center 3 ▲

This table lists the unused ILM rules that still refer to the site. For each rule listed, you must do one of the following:

- Edit the rule to remove the Erasure Coding profile or storage pool from the placement instructions.
- Delete the rule.

[Go to the ILM Rules page](#)

Name	EC Profiles	Storage Pools	Delete
Make 2 Copies	—	All Storage Nodes	
3 copies for S3 tenant	—	Raleigh storage pool	
2 copies 2 sites for smaller objects	—	Raleigh storage pool	
EC larger objects	three site EC profile	All 3 Sites	

**1 Erasure Coding profile** will be deactivated ▼

**3 storage pools** will be deleted ▼

Any ILM rules that are listed still refer to the site but are not used in any policy. In the example:

- The stock **Make 2 Copies** rule uses the system-default **All Storage Nodes** storage pool, which uses the All Sites site.
- The unused **3 copies for S3 tenant** rule refers to the **Raleigh** storage pool.
- The unused **2 copy 2 sites for smaller objects** rule refers to the **Raleigh** storage pool.
- The unused **EC larger objects** rules uses the Raleigh site in the **All 3 Sites** Erasure Coding profile.
- If no ILM rules are listed, select **Next** to go to **Step 5 (Resolve Node Conflicts)**.

### Step 5: Resolve Node Conflicts (and start decommission)



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 system-default All Storage Nodes storage pool is removed because it uses the All Sites site.

- If one or more ILM rules are listed, go to the next step.

3. 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)**.



For details, see the instructions for managing objects with information lifecycle management.

- To delete a rule, select the trash can icon  and select **OK**.



You must delete the stock **Make 2 Copies** rule before you can decommission a site.

4. Confirm that no proposed ILM policy exists, no unused ILM rules refer to the site, and the **Next** button is enabled.

#### Decommission Site



Before you can decommission a site, you must ensure that no proposed ILM policy exists and that no ILM rules refer to the site, even if those rules are not currently used in an ILM policy.

No proposed policy exists	
No ILM rules refer to Raleigh	
1 Erasure Coding profile will be deactivated	▼
3 storage pools will be deleted	▼

Previous **Next**

5. 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 system-default All Storage Nodes storage pool 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.

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.

- No node at the site you are removing can have an interface that belongs to a high availability (HA) group.

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 cannot create ILM rules that refer to the site being decommissioned. You also cannot edit an existing ILM rule to refer to the site.
  - You cannot 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 instructions for monitoring and troubleshooting StorageGRID and the grid node procedures. Contact technical support if you need assistance.

3. 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

4. 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.

5. 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.

### Related information

[Monitor & troubleshoot](#)

[Grid node procedures](#)

[Administer StorageGRID](#)

## 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



Name 	Type 	Progress 	Stage 
RAL-S1-101-196	Storage Node	<div style="width: 20px; height: 10px; background-color: #00a0e3; border: 1px solid #ccc;"></div>	Decommissioning Replicated and Erasure Coded Data
RAL-S2-101-197	Storage Node	<div style="width: 20px; height: 10px; background-color: #00a0e3; border: 1px solid #ccc;"></div>	Decommissioning Replicated and Erasure Coded Data
RAL-S3-101-198	Storage Node	<div style="width: 20px; height: 10px; background-color: #00a0e3; border: 1px solid #ccc;"></div>	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 <b>Note:</b> 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	<div style="width: 0%;"><div></div></div> 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 system-default All Storage Nodes storage pool 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
<input checked="" type="radio"/>	Sunnyvale	4.79 MB	
<input type="radio"/>	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.

#### Related information

[Downloading the Recovery Package](#)

### Network maintenance procedures

You can configure the list of subnets on the Grid Network or update IP addresses, DNS servers, or NTP servers for your StorageGRID system.

#### Choices

- [Updating subnets for the Grid Network](#)
- [Configuring IP addresses](#)

- [Configuring DNS servers](#)
- [Configuring NTP servers](#)
- [Restoring network connectivity for isolated nodes](#)

## Updating subnets for 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. When you add grid nodes or a new site in an expansion, you might need to update or add subnets to the Grid Network.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.
- You must 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 the new Grid subnet before you start the expansion procedure.

### Steps

1. Select **Maintenance > Network > Grid Network**.

#### Grid Network

Configure the subnets that are used on the Grid Network. These entries typically include the subnets for the Grid Network (eth0) for each site in your StorageGRID system as well as any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

#### Subnets

Subnet 1  **+**

#### Passphrase

Provisioning  
Passphrase

Save

2. In the Subnets list, click the plus sign to add a new subnet in CIDR notation.

For example, enter 10.96.104.0/22.

3. Enter the provisioning passphrase, and click **Save**.

The subnets you have specified are configured automatically for your StorageGRID system.

## Configuring IP addresses

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.



If you want to change the Grid Network IP address for all nodes in the grid, use the special procedure for grid-wide changes.

### Changing IP addresses for 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.



The IP change procedure can be a disruptive procedure. Parts of the grid might be unavailable until the new configuration is applied.

## 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 cannot 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

- You cannot change the Client network IP address outside the subnet of an HA group configured on the Client network interface.
- You cannot change the Client network IP address to the value of an existing virtual IP address assigned by an HA group configured on the Client network interface.
- You cannot change the Grid network IP address outside the subnet of an HA group configured on the Grid network interface.
- You cannot change the Grid network IP address to the value of an existing virtual IP address assigned by an HA group configured on the Grid network interface.

## Choices

- [Changing a node's network configuration](#)
- [Adding to or changing subnet lists on the Admin Network](#)
- [Adding to or changing subnet lists on the Grid Network](#)
- [Linux: Adding interfaces to an existing node](#)
- [Changing IP addresses for all nodes in the grid](#)

### Changing a node's 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.

#### What you'll need

You must 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.

**Appliances:** On StorageGRID appliances, if the Client or Admin Network was not configured in the StorageGRID Appliance Installer during the initial installation, the network cannot 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 in the installation and maintenance instructions for your appliance.

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 cannot be removed from the Grid Network.



IP address swaps are not 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 administering StorageGRID.



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` cannot be edited.

- a. To set a new value, enter it in the format shown for the current value.
- b. To leave the current value unchanged, press **Enter**.
- c. 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**.
- d. 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 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 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, are not violated.

In this example, validation returned errors.

```
Validating new networking configuration... FAILED.
DK-10-224-5-20-G1: The admin subnet 172.18.0.0/16 overlaps the 172.18.0.0/21 grid network.
DK-10-224-5-22-S1: Duplicate Grid IP 172.16.5.18 (also in use by DK-10-224-5-21-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. Once 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**: Do not 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 applying or staging 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 “Adding interfaces to an 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.

#### Related information

[Linux: Adding interfaces to an existing node](#)

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

[Install Ubuntu or Debian](#)

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[Administer StorageGRID](#)

## Configuring IP addresses

### Adding to or changing subnet lists on the Admin Network

You can add, delete, or change the subnets in the Admin Network Subnet List of one or more nodes.

#### What you'll need

- You must have the `Passwords.txt` file.

You can add, delete, or change subnets to all nodes on the Admin Network Subnet List.

#### 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, you must 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 are not persistent. Host routes (/32 subnets) are applied automatically if the IP addresses provided for external services such as NTP or DNS are not 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 do not 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 do not 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 sample 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.

- Select **Maintenance > System > Recovery Package**.
- Enter the provisioning passphrase.

### Related information

[Configuring IP addresses](#)

### Adding to or changing subnet lists on the Grid Network

You can use the Change IP tool to add or change subnets on the Grid Network.

### What you'll need

- You have the `Passwords.txt` file.

### About this task

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.

### Steps

- 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 `#`.

- Start the Change IP tool by entering the following command: `change-ip`
- 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 sample 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, you must restart the node as soon as possible after staging to minimize disruptions.

- **cancel**: Do not 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 applying or staging 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 abort 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.

## Related information

[Configuring IP addresses](#)

### Linux: Adding interfaces to an existing node

If you want to add an interface to a Linux-based node that you did not install initially, you must use this procedure.

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 StorageGRID installation instructions for your Linux operating system.

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

[Install 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 the information about changing a node's network configuration.

[Changing a node's network configuration](#)

### Steps

1. Log in to the Linux server hosting the node that needs the new network assignment.
2. Edit the node configuration file at `/etc/storagegrid/nodes/node-name.conf`.



Do not specify any other network parameters, or a validation error will result.

- a. Add the new network target.

```
CLIENT_NETWORK_TARGET = bond0.3206
```

- b. Optional: Add a MAC address.

```
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`

### Related information

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

[Install Ubuntu or Debian](#)

[Changing a node's network configuration](#)

### Changing IP addresses for all nodes in the grid

If you need to change the Grid Network IP address for all nodes in the grid, you must follow this special procedure. You cannot do a grid-wide Grid Network IP change using the procedure to change individual nodes.

### What you'll need

- You must have the `Passwords.txt` file.

### About this task

To ensure that the grid starts up successfully, you must make all the changes at once.



This procedure applies to the Grid Network only. You cannot 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 instructions for changing a node's network configuration.

### 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 the instructions for administering StorageGRID.



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` cannot 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 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 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, are not violated.

In this example, validation returned errors.

```

Validating new networking configuration... FAILED.

DK-10-224-5-20-G1: The admin subnet 172.18.0.0/16 overlaps the 172.18.0.0/21 grid network.
DK-10-224-5-22-S1: Duplicate Grid IP 172.16.5.18 (also in use by DK-10-224-5-21-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. Once 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**. Do not 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 at once, 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. Once 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

[Administer StorageGRID](#)

[Changing a node's network configuration](#)

[Adding to or changing subnet lists on the Grid Network](#)

[Shutting down a grid node](#)

#### Configuring DNS servers

You can add, remove, and update domain name system (DNS) servers, so that you can use fully qualified domain name (FQDN) hostnames rather than IP addresses.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the IP addresses of the DNS servers to configure.

#### About this task

Specifying DNS server information allows you to use fully qualified domain name (FQDN) hostnames rather than IP addresses for email or SNMP notifications and AutoSupport. Specifying at least two DNS servers is recommended.



Provide between two to six IP addresses for DNS servers. In general, select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node.

### Modifying the DNS configuration for a single grid node

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

#### Steps

1. Select **Maintenance > Network > DNS Servers**.
2. In the Servers section, add update, or remove DNS server entries, as necessary.

The best practice is to specify at least two DNS servers per site. You can specify up to six DNS servers.

3. Click **Save**.

### Modifying the DNS configuration for a single grid node

Rather than configure the Domain Name System (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.

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`

## Configuring 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.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.
- You must have the IPv4 addresses of the NTP servers to configure.

### About this task

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.



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.

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

### [Support boundary to configure the Windows Time service for high-accuracy environments](#)

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.

### 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 4 NTP servers, and you can specify up to 6 servers.

3. In the **Provisioning Passphrase** text box, enter the provisioning passphrase for your StorageGRID

system and click **Save**.

The status of the procedure is displayed at the top of the page. 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, do not proceed. Contact technical support.

## Restoring network connectivity for isolated nodes

Under certain circumstances, such as site- or grid-wide IP address changes, one or more groups of nodes might not be able to contact the rest of the grid.

In the Grid Manager (**Support > Tools > Grid Topology**), if a node is gray, or if a node is blue with many of its services showing a status other than Running, you should check for node isolation.

The screenshot shows the Grid Manager interface. On the left is a tree view of the Grid Topology, showing a hierarchy from Grid1 down to Site1, then abrian-g1, SSM, Services, Events, Resources, Timing, CLB, and finally three server nodes: abrian-s1, abrian-s2, and abrian-s3. On the right is the 'Overview: SSM (abrian-g1) - Services' page. It includes tabs for Overview, Alarms, Reports, and Configuration. The Overview tab is active, showing the operating system as Linux 4.9.0-3-amd64. Below this is a table of services with columns for Service, Version, Status, Threads, Load, and Memory. The services table is as follows:

Service	Version	Status	Threads	Load	Memory
ADE Exporter Service	11.1.0-20171214.1441.c29e2f8	Running	11	0.011 %	7.87 MB
Connection Load Balancer (CLB)	11.1.0-20180120.0111.02137fe	Running	61	0.07 %	39.3 MB
Dynamic IP Service	11.1.0-20180123.1919.deeeba7.abrian	Not Running	0	0 %	0 B
Nginx Service	1.10.3-1+deb9u1	Running	5	0.002 %	20 MB
Node Exporter Service	0.13.0+ds-1+b2	Running	5	0 %	8.58 MB
Persistence Service	11.1.0-20180123.1919.deeeba7.abrian	Running	6	0.064 %	17.1 MB
Server Manager	11.1.0-20171214.1441.c29e2f8	Running	4	2.116 %	18.7 MB
Server Status Monitor (SSM)	11.1.0-20180120.0111.02137fe	Running	61	0.288 %	45.8 MB
System Logging	3.8.1-10	Running	3	0.006 %	8.27 MB
Time Synchronization	1:4.2.8p10+dfsg-3+deb9u1	Running	2	0.007 %	4.54 MB

Below the services table is a 'Packages' section with a table showing installed packages:

Package	Installed	Version
storage-grid-release	Installed	11.1.0-20180123.1919.deeeba7.abrian

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, the command line utility might have to be run on each isolated node.

## Steps

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 may 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.
- In the Grid Topology tree, 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-level and middleware procedures

Some maintenance procedures are specific to Linux or VMware deployments of StorageGRID, or are specific to other components of the StorageGRID solution.

### Linux: Migrating a grid node to a new host

You can migrate StorageGRID nodes from one Linux host to another to perform host maintenance (such as OS patching and reboot) without impacting the functionality or availability of your grid.

You migrate one or more nodes from one Linux host (the “source host”) to another Linux host (the “target host”). The target host must have previously been prepared for StorageGRID use.



You can use this procedure only if you planned your StorageGRID deployment to include migration support.

To migrate a grid node to a new host, both of the following conditions must be true:

- Shared storage is used for all per-node storage volumes
- Network interfaces have consistent names across hosts



In a production deployment, do not 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), do not install all instances on the same host.

For more information, see “Node migration requirements” in the StorageGRID installation instructions for your Linux operating system.

### Related information

[Deploying new Linux hosts](#)

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

[Install Ubuntu or Debian](#)

### Linux: Exporting the node from the source host

Shut down the grid node and export it from the source Linux host.

Run the following command on the source Linux host.

1. Obtain the status of all nodes currently running on the source host.

```
sudo storagegrid node status all
```

```
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
```

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
```

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. Take note of the import command suggested in the output of the `export` command.

You will run this command on the target host in the next step.

#### Linux: Importing the node on the target host

After exporting the node from the source host, you import and validate the node on the target Linux 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 command on the target Linux host.

1. Import the node on the target host.

```
sudo storagegrid node import /dev/mapper/sgws-dc1-s3-var-local
```

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
```

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.

### Related information

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

[Install Ubuntu or Debian](#)

### Linux: Starting the migrated node

After you validate the migrated node, you start the node by running a command on the target Linux host.

### Steps

1. Start the node on the new host.

```
sudo storagegrid node start DC1-S3
Starting node DC1-S3
```

2. In the Grid Manager, verify that the status of the node is green with no alarms raised against it.



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, do not migrate any additional nodes so that you will not have more than one node out of service.

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.

### Archive Node maintenance for TSM middleware

Archive Nodes might be configured to target either tape through a TSM middleware server or the cloud through the S3 API. Once configured, an Archive Node's target cannot be changed.

If the server hosting the Archive Node fails, replace the server and follow the appropriate recovery procedure.

### Fault with archival storage devices

If you determine that there is a fault with the archival storage device that the Archive Node is accessing through Tivoli Storage Manager (TSM), take the Archive Node offline to limit the number of alarms displayed in the StorageGRID system. You can then use the administrative tools of the TSM server or the storage device,

or both, to further diagnose and resolve the problem.

## Taking the Target component offline

Before undertaking any maintenance of the TSM middleware server that might result in it becoming unavailable to the Archive Node, take the Target component offline to limit the number of alarms that are triggered if the TSM middleware server becomes unavailable.

### What you'll need

You must be signed in to the Grid Manager using a supported browser.

### Steps

1. Select **Support > Tools > Grid Topology**.
2. Select **Archive Node > ARC > Target > Configuration > Main**.
3. Change the value of Tivoli Storage Manager State to **Offline**, and click **Apply Changes**.
4. After maintenance is complete, change the value of Tivoli Storage Manager State to **Online**, and click **Apply Changes**.

### Tivoli Storage Manager administrative tools

The `dsmadm` tool is the administrative console for the TSM middleware server that is installed on the Archive Node. You can access the tool by typing `dsmadm` at the command line of the server. Log in to the administrative console using the same administrative user name and password that is configured for the ARC service.

The `tsmquery.rb` script was created to generate status information from `dsmadm` in a more readable form. You can run this script by entering the following command at the command line of the Archive Node:

```
/usr/local/arc/tsmquery.rb status
```

For more information about the TSM administrative console `dsmadm`, see the *Tivoli Storage Manager for Linux: Administrator's Reference*.

### Object permanently unavailable

When the Archive Node requests an object from the Tivoli Storage Manager (TSM) server and the retrieval fails, the Archive Node retries the request after an interval of 10 seconds. If the object is permanently unavailable (for example, because the object is corrupted on tape), the TSM API has no way to indicate this to the Archive Node, so the Archive Node continues to retry the request.

When this situation occurs, an alarm is triggered, and the value continues to increase. To see the alarm, select **Support > Tools > Grid Topology**. Then, select **Archive Node > ARC > Retrieve > Request Failures**.

If the object is permanently unavailable, you must identify the object and then manually cancel the Archive Node's request as described in the procedure, [Determining if objects are permanently unavailable](#).

A retrieval can also fail if the object is temporarily unavailable. In this case, subsequent retrieval requests should eventually succeed.

If the StorageGRID system is configured to use an ILM rule that creates a single object copy and that copy cannot be retrieved, the object is lost and cannot be recovered. However, you must still follow the procedure to determine if the object is permanently unavailable to "clean up" the StorageGRID system, to cancel the Archive Node's request, and to purge metadata for the lost object.

## Determining if objects are permanently unavailable

You can determine if objects are permanently unavailable by making a request using the TSM administrative console.

### What you'll need

- You must have specific access permissions.
- You must have the `Passwords.txt` file.
- You must know the IP address of an Admin Node.

### About this task

This example is provided for your information only; this procedure cannot help you identify all failure conditions that might result in unavailable objects or tape volumes. For information about TSM administration, see TSM Server documentation.

### 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.
2. Identify the object or objects that could not be retrieved by the Archive Node:
  - a. Go to the directory containing the audit log files: `cd /var/local/audit/export`

The active audit log file is named `audit.log`. 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`. After a day, the saved file is compressed and renamed, in the format `yyyy-mm-dd.txt.gz`, which preserves the original date.

- b. Search the relevant audit log file for messages indicating that an archived object could not be retrieved. For example, enter: `grep ARCE audit.log | less -n`

When an object cannot be retrieved from an Archive Node, the ARCE audit message (Archive Object Retrieve End) displays ARUN (archive middleware unavailable) or GERR (general error) in the result field. The following example line from the audit log shows that the ARCE message terminated with the result ARUN for CBID 498D8A1F681F05B3.

```
[AUDT: [CBID(UI64) :0x498D8A1F681F05B3] [VLID(UI64) :□20091127] [RSLT(FC32) :ARUN] [AVER(UI32) :7]
[ATIM(UI64) :1350613602969243] [ATYP(FC32) :ARCE] [ANID(UI32) :13959984] [AMID(FC32) :ARCI]
[ATID(UI64) :4560349751312520631]]
```

For more information see the instructions for understanding audit messages.

- c. Record the CBID of each object that had a request failure.

You might also want to record the following additional information used by the TSM to identify objects saved by the Archive Node:

- **File Space Name:** Equivalent to the Archive Node ID. To find the Archive Node ID, select **Support > Tools > Grid Topology**. Then, select **Archive Node > ARC > Target > Overview**.
- **High Level Name:** Equivalent to the volume ID assigned to the object by the Archive Node. The volume ID takes the form of a date (for example, 20091127), and is recorded as the VLID of the object in archive audit messages.
- **Low Level Name:** Equivalent to the CBID assigned to an object by the StorageGRID system.

d. Log out of the command shell: `exit`

3. Check the TSM server to see if the objects identified in step 2 are permanently unavailable:

a. Log in to the administrative console of the TSM server: `dsmadm`

Use the administrative user name and password that are configured for the ARC service. Enter the user name and password in the Grid Manager. (To see the user name, select **Support > Tools > Grid Topology**. Then, select **Archive Node > ARC > Target > Configuration**.)

b. Determine if the object is permanently unavailable.

For example, you might search the TSM activity log for a data integrity error for that object. The following example shows a search of the activity log for the past day for an object with CBID 498D8A1F681F05B3.

```
> query actlog begindate=-1 search=276C14E94082CC69
12/21/2008 05:39:15 ANR0548W Retrieve or restore
failed for session 9139359 for node DEV-ARC-20 (Bycast ARC)
processing file space /19130020 4 for file /20081002/
498D8A1F681F05B3 stored as Archive - data
integrity error detected. (SESSION: 9139359)
>
```

Depending on the nature of the error, the CBID might not be recorded in the TSM activity log. You might need to search the log for other TSM errors around the time of the request failure.

c. If an entire tape is permanently unavailable, identify the CBIDs for all objects stored on that volume:

```
query content TSM_Volume_Name
```

where `TSM_Volume_Name` is the TSM name for the unavailable tape. The following is an example of the output for this command:

```
> query content TSM-Volume-Name
Node Name      Type Filespace  FSID Client's Name for File Name
-----
DEV-ARC-20    Arch /19130020   216 /20081201/ C1D172940E6C7E12
DEV-ARC-20    Arch /19130020   216 /20081201/ F1D7FBC2B4B0779E
```

The Client's Name for File Name is the same as the Archive Node volume ID (or TSM "high level name") followed by the object's CBID (or TSM "low level name"). That is, the Client's Name for File Name takes the form /Archive Node volume ID /CBID. In the first line of the

example output, the Client's Name for File Name is /20081201/ C1D172940E6C7E12.

Recall also that the `Filespace` is the node ID of the Archive Node.

You will need the CBID of each object stored on the volume and the node ID of the Archive Node to cancel the retrieval request.

4. For each object that is permanently unavailable, cancel the retrieval request and issue a command to inform the StorageGRID system that the object copy was lost:



Use the ADE Console with caution. If the console is used improperly, it is possible to interrupt system operations and corrupt data. Enter commands carefully, and only use the commands documented in this procedure.

- a. If you are not already logged in to the Archive Node, log in as follows:

- 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. Access the ADE console of the ARC service: `telnet localhost 1409`

- c. Cancel the request for the object: `/proc/BRTR/cancel -c CBID`

where `CBID` is the identifier of the object that cannot be retrieved from the TSM.

If the only copies of the object are on tape, the “bulk retrieval” request is canceled with a message, “1 requests canceled”. If copies of the object exist elsewhere in the system, the object retrieval is processed by a different module so the response to the message is “0 requests canceled”.

- d. Issue a command to notify the StorageGRID system that an object copy has been lost and that an additional copy must be made: `/proc/CMSI/Object_Lost CBID node_ID`

where `CBID` is the identifier of the object that cannot be retrieved from the TSM server, and `node_ID` is the node ID of the Archive Node where the retrieval failed.

You must enter a separate command for each lost object copy: entering a range of CBIDs is not supported.

In most cases, the StorageGRID system immediately begins to make additional copies of object data to ensure that the system's ILM policy is followed.

However, if the ILM rule for the object specified that only one copy be made and that copy has now been lost, the object cannot be recovered. In this case running the `Object_Lost` command purges the lost object's metadata from the StorageGRID system.

When the `Object_Lost` command completes successfully, the following message is returned:

```
CLOC_LOST_ANS returned result 'SUCS'
```



The `/proc/CMSI/Object_Lost` command is only valid for lost objects that are stored on Archive Nodes.

- e. Exit the ADE Console: `exit`
  - f. Log out of the Archive Node: `exit`
5. Reset the value of Request Failures in the StorageGRID system:
- a. Go to **Archive Node > ARC > Retrieve > Configuration**, and select **Reset Request Failure Count**.
  - b. Click **Apply Changes**.

#### Related information

[Administer StorageGRID](#)

[Review audit logs](#)

#### VMware: Configuring a 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.

#### Grid node procedures

You might need to perform procedures on a specific grid node. While you can perform a few of these procedures from Grid Manager, most of the procedures require you to access Server Manager from the node's command line.

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.



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`

#### Choices

- [Viewing Server Manager status and version](#)
- [Viewing current status of all services](#)
- [Starting Server Manager and all services](#)

- [Restarting Server Manager and all services](#)
- [Stopping Server Manager and all services](#)
- [Viewing current status of a service](#)
- [Stopping a service](#)
- [Placing an appliance into maintenance mode](#)
- [Forcing a service to terminate](#)
- [Starting or restarting a service](#)
- [Removing port remaps](#)
- [Removing port remaps on bare metal hosts](#)
- [Rebooting a grid node](#)
- [Shutting down a grid node](#)
- [Powering down a host](#)
- [Powering off and on all nodes in the grid](#)
- [Using a DoNotStart file](#)
- [Troubleshooting Server Manager](#)

## Viewing 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.

### What you'll need

You must 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
```

This status is the equivalent of the status shown in the header of the local console display.

3. View the current version of Server Manager running on a grid node: **service servermanager version**

The current version is listed. For example:

```
11.1.0-20180425.1905.39c9493
```

4. Log out of the command shell: **exit**

### Viewing current status of all services

You can view the current status of all services running on a grid node at any time.

#### What you'll need

You must 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`

## Starting Server Manager and all services

You might need to start Server Manager, which also starts all services on the grid node.

### What you'll need

You must 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`

### Restarting Server Manager and all services

You might need to restart server manager and all services running on a grid node.

#### What you'll need

You must 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`

### Stopping 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.

#### What you'll need

You must have the `Passwords.txt` file.

#### About this task

The only scenario that requires you to stop Server Manager while keeping the operating system running is when you need to integrate Server Manager to other services. If there is a requirement to stop the Server Manager for servicing of the hardware or reconfiguration of the server, the entire server should be halted.

#### 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`

### Viewing current status of a service

You can view the current status of a services running on a grid node at any time.

#### What you'll need

You must 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`

### Stopping a 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.

#### What you'll need

You must 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:
  - 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 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

[Forcing a service to terminate](#)

### Placing an appliance into maintenance mode

You must place the appliance into maintenance mode before performing specific maintenance procedures.

#### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission. For details, see the instructions for administering StorageGRID.

#### About this task

Placing a StorageGRID appliance into maintenance mode might make the appliance unavailable for remote access.



The password and host key for a StorageGRID appliance in maintenance mode remain the same as they were when the appliance was in service.

### Steps

1. From the Grid Manager, select **Nodes**.

- From the tree view of the Nodes page, select the appliance Storage Node.
- Select **Tasks**.

The screenshot shows a navigation bar with tabs: Overview, Hardware, Network, Storage, Objects, ILM, Events, and Tasks. The 'Tasks' tab is selected and highlighted. Below the navigation bar, there are two sections: 'Reboot' and 'Maintenance Mode'. The 'Reboot' section has a description 'Shuts down and restarts the node.' and a blue button labeled 'Reboot'. The 'Maintenance Mode' section has a description 'Places the appliance's compute controller into maintenance mode.' and a blue button labeled 'Maintenance Mode'.

- Select **Maintenance Mode**.

A confirmation dialog box appears.

The dialog box has a yellow header with a warning icon and the text 'Enter Maintenance Mode on SGA-106-15'. Below the header, there is a paragraph: 'You must place the appliance's compute controller into maintenance mode to perform certain maintenance procedures on the appliance.' followed by an attention note: 'Attention: All StorageGRID services on this node will be shut down. Wait a few minutes for the node to reboot into maintenance mode.' Below this, it says 'If you are ready to start, enter the provisioning passphrase and click OK.' There is a text input field labeled 'Provisioning Passphrase'. At the bottom right, there are two buttons: 'Cancel' and 'OK'.

- Enter the provisioning passphrase, and select **OK**.

A progress bar and a series of messages, including "Request Sent", "Stopping StorageGRID", and "Rebooting", indicate that the appliance is completing the steps for entering maintenance mode.


## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

**Attention:** Your request has been sent, but the appliance might take 10-15 minutes to enter maintenance mode. Do not perform maintenance procedures until this tab indicates maintenance mode is ready, or data could become corrupted.

 Request Sent

When the appliance is in maintenance mode, a confirmation message lists the URLs you can use to access the StorageGRID Appliance Installer.

## Reboot

Shuts down and restarts the node.

Reboot

## Maintenance Mode

This node is currently in maintenance mode. Navigate to one of the URLs listed below and perform any necessary maintenance procedures.

- <https://172.16.2.106:8443>
- <https://10.224.2.106:8443>
- <https://47.47.2.106:8443>
- <https://169.254.0.1:8443>

When you are done with any required maintenance procedures, you must exit maintenance mode by clicking Reboot Controller from the StorageGRID Appliance Installer.

6. To access the StorageGRID Appliance Installer, browse to any of the URLs displayed.

If possible, use the URL containing the IP address of the appliance's Admin Network port.




Accessing <https://169.254.0.1:8443> requires a direct connection to the local management port.


7. From the StorageGRID Appliance Installer, confirm that the appliance is in maintenance mode.

⚠ This node is in maintenance mode. Perform any required maintenance procedures. If you want to exit maintenance mode manually to resume normal operation, go to **Advanced > Reboot Controller** to **reboot** the controller.

8. Perform any required maintenance tasks.
9. After completing maintenance tasks, exit maintenance mode and resume normal node operation. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.



It can take up to 20 minutes for the appliance to reboot and rejoin the grid. To confirm that the reboot is complete and that the node has rejoined the grid, go back to the Grid Manager. The **Nodes** tab should display a normal status ✓ for the appliance node, indicating that no alerts are active and the node is connected to the grid.



## Forcing a service to terminate

If you need to stop a service immediately, you can use the `force-stop` command.

### What you'll need

You must 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. 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`

### Starting or restarting a service

You might need to start a service that has been stopped, or you might need to stop and restart a service.

#### What you'll need

You must 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`

## Removing 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.



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.



If the port remap provides client access, the client should be reconfigured to use a different port configured as an load balancer endpoint if possible, 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 Docker container 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.

Follow the instructions for rebooting a grid node.

4. Repeat these steps on each Admin Node and Gateway Node that has conflicting remapped ports.

### Related information

[Administer StorageGRID](#)

[Rebooting a grid node](#)

[Removing port remaps on bare metal hosts](#)

### Removing 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. 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
<strong>PORT_REMAP = client/tcp/8082/443</strong>
<strong>PORT_REMAP_INBOUND = client/tcp/8082/443</strong>
```

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 listing 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.

## Rebooting a grid node

You can reboot a grid node from the Grid Manager or from the node's command shell.

### About this task

When you reboot a grid node, the node shuts down and restarts. All services are restarted automatically.

If you plan to reboot Storage Nodes, note the following:

- If an ILM rule specifies an ingest behavior of Dual commit or the rule specifies Balanced and it is not possible to immediately create all required copies, StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site and evaluates ILM later. If you want to reboot two or more Storage Nodes on a given site, you might not be able to access these objects for the duration of the reboot.
- 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.

### Related information

[Administer StorageGRID](#)

### Choices

- [Rebooting a grid node from the Grid Manager](#)
- [Rebooting a grid node from the command shell](#)

### Rebooting a grid node from the Grid Manager

Rebooting a grid node from the Grid Manager issues the `reboot` command on the target node.

### What you'll need

- You must be signed in to the Grid Manager using a supported browser.
- You must have the Maintenance or Root Access permission.
- You must have the provisioning passphrase.

### Steps

1. Select **Nodes**.
2. Select the grid node you want to reboot.
3. Select the **Tasks** tab.

## DC3-S3 (Storage Node)

Overview

Hardware

Network

Storage

Objects

ILM

Events

Tasks

### Reboot

Reboot shuts down and restarts the node.

Reboot

#### 4. Click **Reboot**.

A confirmation dialog box appears.

#### Reboot Node DC3-S3

Reboot shuts down and restarts a node, based on where the node is installed:

- 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 are ready to reboot this node, enter the provisioning passphrase and click OK.

Provisioning Passphrase

Cancel

OK



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 click **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 icon (Administratively Down) appears on the left side of the Nodes page. When all services have started again, the icon changes back to its original color.

### Rebooting a grid node from the 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 into the grid node and run the Server Manager reboot command from the command shell.

### What you'll need

- You must 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`

## Shutting down a grid node

You can shut down a grid node from the node's command shell.

### What you'll need

- You must 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.
- Do not 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 Storage Nodes, note the following:
  - If an ILM rule specifies an ingest behavior of Dual commit or the rule specifies Balanced and it is not possible to immediately create all required copies, StorageGRID immediately commits any newly ingested objects to two Storage Nodes on the same site and evaluates ILM later. If you want to shut down two or more Storage Nodes on a given site, you might not be able to access these objects for the duration of the shutdown.
  - To ensure you can access all objects when a Storage Node is shut down, stop ingesting objects at a site for approximately one hour before shutting down the node.

## 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. Log out of the command shell: `exit`

After being shut down you can power off the grid node.

## Powering down a host

### Related information

[Administer StorageGRID](#)

## Powering down a 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:

- For the SG100 or SG1000 services appliance

- i. Turn off the power to the appliance.
- ii. Wait for the blue power LED to turn off.

- For the SG6000 appliance

- i. 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.

- ii. Turn off the power to the appliance, and wait for the blue power LED to turn off.

- For the SG5700 appliance

- i. 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.

- ii. Turn off the power to the appliance, and wait for all LED and seven-segment display activity to stop.

7. Log out of the command shell: `exit`

### Related information

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

### Powering off and on all nodes in the 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.

## Stopping services and shutting 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, Docker containers, and StorageGRID appliances.

### About this task

If possible, you should stop services on the grid nodes in this order:

- Stop services on Gateway Nodes first.
- Stop services on the primary Admin Node last.

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, do not 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.

## Linux: Migrating a grid node to a new host

### 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, Archive 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:
  - For the SG100 or SG1000 services appliance
    - i. Turn off the power to the appliance.
    - ii. Wait for the blue power LED to turn off.
  - For the SG6000 appliance
    - i. 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.
    - ii. Turn off the power to the appliance, and wait for the blue power LED to turn off.
  - For the SG5700 appliance
    - i. 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.
    - ii. 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.

#### **Related information**

[SG100 & SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

#### **Starting up the grid nodes**

Follow this sequence to start up the grid nodes after a complete shutdown.



If the entire grid has been shut down for more than 15 days, you must contact technical support before starting up any grid nodes. Do not attempt the recovery procedures that rebuild Cassandra data. Doing so might result in data loss.

### About this task

If possible, you should 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 Archive Nodes and Storage Nodes.

You can power on these nodes in any order.

3. Power on the hosts for all Gateway Nodes.
4. Sign into the Grid Manager.
5. Click **Nodes**, and monitor the status of the grid nodes. Verify that all nodes return to “green” status.



## Using 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 are not 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 cannot run if the NTP service is not running.

### Adding a DoNotStart file for a service

You can prevent an individual service from starting by adding a DoNotStart file to that service's directory on a grid node.

### What you'll need

You must 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`

## Removing a DoNotStart file for a service

When you remove a DoNotStart file that is preventing a service from starting, you must start that service.

### What you'll need

You must 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`

## Troubleshooting Server Manager

Technical support might direct you to troubleshooting tasks to determine the source of Server Manager-related problems.

### Accessing the Server Manager log file

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.

### What you'll need

You must 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`

## Related information

[Removing a DoNotStart file for a service](#)

## Appliance node cloning

You can clone an appliance node in StorageGRID to use an appliance of newer design or increased capabilities. Cloning transfers all information on the existing node to the new appliance, provides a hardware-upgrade process that is easy to perform, and provides an alternative to decommissioning and expansion for replacing appliances.

## How appliance node cloning works

Appliance node cloning lets you easily replace an existing appliance node (source) in your grid with a compatible appliance (target) that is part of the same logical StorageGRID site. The process transfers all data to the new appliance, placing it in service to replace the old appliance node and leaving the old appliance in a pre-install state.

### Why clone an appliance node?

You can clone an appliance node if you need to:

- Replace appliances that are reaching end-of-life.
- Upgrade existing nodes 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 from DDP-8 to DDP-16, or to RAID-6.
- Efficiently implement node encryption to allow the use of external key management servers (KMS).

### Which StorageGRID network is used?

Cloning transfers data from the source node directly to the target appliance over any of the three StorageGRID networks. The Grid Network is typically used, but you can also use the Admin Network or the Client Network if the source appliance is connected to these networks. Choose the network to use for cloning traffic that provides the best data-transfer performance without degrading StorageGRID network performance or data availability.

When you install the replacement appliance, you must specify temporary IP addresses for StorageGRID connection and data transfer. Since the replacement appliance will be part of the same networks as the appliance node it replaces, you must specify temporary IP addresses for each of these networks on the replacement appliance.

### Target appliance compatibility

Replacement appliances must be the same type as the source node they are replacing and both must be part of the same logical StorageGRID site.

- A replacement services appliance can be different than the Admin Node or Gateway Node it is replacing.
  - You can clone an SG100 source node appliance to an SG1000 services target appliance to give the Admin Node or Gateway Node greater capability.
  - You can clone an SG1000 source node appliance to an SG100 services target appliance to redeploy the SG1000 for a more demanding application.

For example, if an SG1000 source node appliance is being used as an Admin Node and you want to use it as a dedicated load-balancing node.

- Replacing an SG1000 source node appliance with an SG100 services target appliance reduces the maximum speed of the network ports from 100-GbE to 25-GbE.
- The SG100 and SG1000 appliances have different network connectors. Changing the appliance type might require replacing the cables or SFP modules.

- A replacement storage appliance must have equal or greater capacity than the Storage Node it is replacing.
  - If the target storage appliance has the same number of drives as the source node, the drives in the target appliance must have the same capacity (in TB) or larger.
  - If the number of standard drives installed in a target storage appliance is less than the number of drives in the source node, due to installation of solid-state drives (SSDs), the overall storage capacity of the standard drives in the target appliance (in TB) must meet or exceed the total functional drive capacity of all drives in the source Storage Node.

For example, when cloning an SG5660 source Storage Node appliance with 60 drives to an SG6060 target appliance with 58 standard drives, larger drives should be installed in the SG6060 target appliance before cloning to maintain storage capacity. (The two drive slots containing SSDs in the target appliance are not included in the total appliance-storage capacity.)

However, if a 60-drive SG5660 source node appliance is configured with SANtricity Dynamic Disk Pools DDP-8, configuring a 58-drive same-size-drive SG6060 target appliance with DDP-16 might make the SG6060 appliance a valid clone target due to its improved storage efficiency.

You can view information about the current RAID mode of the source appliance node on the **Nodes** page in Grid Manager. Select the **Storage** tab for the appliance.

#### What information is not cloned?

The following appliance configurations do not transfer to the replacement appliance during cloning. You must configure them during initial set up of the replacement appliance.

- BMC interface
- Network links
- Node encryption status
- SANtricity System Manager (for Storage Nodes)
- RAID mode (for Storage Nodes)

#### What issues prevent cloning?

If any of the following issues are encountered while cloning, the cloning process halts and an error message is generated:

- Wrong network configuration
- Lack of connectivity between the source and target appliances
- Source and target appliance incompatibility
- For Storage Nodes, a replacement appliance of insufficient capacity

You must resolve each issue for cloning to continue.

#### Considerations and requirements for appliance node cloning

Before cloning an appliance node, you must understand the considerations and requirements.

## Hardware requirements for the replacement appliance

Ensure that the replacement appliance meets the following criteria:

- The source node (appliance being replaced) and the target (new) appliance must be the same type of appliance:
  - You can only clone an Admin Node appliance or a Gateway Node appliance to a new services appliance.
  - You can only clone a Storage Node appliance to a new storage appliance.
- For Admin Node or Gateway Node appliances, the source node appliance and the target appliance do not need to be the same type of appliance; however, changing the appliance type might require replacing the cables or SFP modules.

For example, you can replace a SG1000 node appliance with a SG100 or replace a SG100 appliance with a SG1000 appliance.

- For Storage Node appliances, the source node appliance and the target appliance do not need to be the same type of appliance; however, the target appliance must have the same or greater storage capacity as the source appliance.

For example, you can replace a SG5600 node appliance with a SG5700 or a SG6000 appliance.

Contact your StorageGRID sales representative for help choosing compatible replacement appliances to clone specific appliance nodes in your StorageGRID installation.

### Preparing to clone an appliance node

You must have the following information before you clone an appliance node:

- Obtain a temporary IP address for the Grid Network from your network administrator for use with the target appliance during initial installation. If the source node belongs to an Admin Network or Client Network, obtain temporary IP addresses for these networks.

Temporary IP addresses are normally on the same subnet as the source node appliance being cloned and are not needed after cloning completes. The source and target appliances must both connect to the primary Admin Node of your StorageGRID to establish a cloning connection.

- Determine which network to use for cloning data-transfer traffic that provides the best data-transfer performance without degrading StorageGRID network performance or data availability.



Using the 1-GbE Admin Network for clone data transfer results in slower cloning.

- Determine if node encryption using a key management server (KMS) will be used on the target appliance, so that you can enable node encryption during initial target appliance installation before cloning. You can check if node encryption is enabled on the source appliance node as described in appliance installation.

The source node and target appliance can have different node-encryption settings. Data decryption and encryption is performed automatically during data transfer and when the target node restarts and joins the grid.

- [SG100 & SG1000 services appliances](#)
- [SG5600 storage appliances](#)

- [SG5700 storage appliances](#)
- [SG6000 storage appliances](#)
- Determine if the RAID mode on the target appliance should be changed from its default setting, so you can specify this information during initial target appliance installation before cloning. You can view information about the current RAID mode of the source appliance node on the **Nodes** page in Grid Manager. Select the **Storage** tab for the appliance.

The source node and target appliance can have different RAID settings.

- Plan for sufficient time to complete the node cloning process. Several days might be required to transfer data from an operational Storage Node to a target appliance. Schedule cloning at a time that minimizes the impact to your business.
- You should only clone one appliance node at a time. Cloning can prevent you from performing other StorageGRID maintenance functions at the same time.
- After you have cloned an appliance node, you can use the source appliance that was returned to a pre-install state as the target to clone another compatible node appliance.

## Appliance node cloning procedure

The cloning process might take several days to transfer data between the source node (appliance being replaced) and the target (new) appliance.

### What you'll need

- You have installed the compatible target appliance into a cabinet or rack, connected all cables, and applied power.
- You have verified that the StorageGRID Appliance Installer version on the replacement appliance matches the software version of your StorageGRID system, upgrading the StorageGRID Appliance Installer firmware, if necessary.
- You have configured the target appliance, including configuring StorageGRID connections, SANtricity System Manager (storage appliances only), and the BMC interface.
  - When configuring StorageGRID connections, use the temporary IP addresses.
  - When configuring network links, use the final link configuration.



Leave the StorageGRID Appliance Installer open after you complete initial target appliance configuration. You will return to the target appliance's installer page after you start the node cloning process.

- You have optionally enabled node encryption for the target appliance.
- You have optionally set the RAID mode for the target appliance (storage appliances only).
- [Considerations and requirements for appliance node cloning](#)

[SG100 & SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

You should clone only one appliance node at a time to maintain StorageGRID network performance and data availability.

## Steps

1. Place the source node you are cloning into maintenance mode.

### Placing an appliance into maintenance mode

2. From the StorageGRID Appliance Installer on the source node, in the Installation section of the Home page, select **Enable Cloning**.

The screenshot shows the NetApp StorageGRID Appliance Installer interface. At the top, there is a blue header with the text "NetApp® StorageGRID® Appliance Installer" and a "Help" link on the right. Below the header is a navigation bar with tabs: "Home", "Configure Networking", "Configure Hardware", "Monitor Installation", and "Advanced". The "Home" tab is selected.

Under the "Home" tab, there is a yellow warning box that reads: "⚠ This node is in maintenance mode. Perform any required maintenance procedures. If you want to exit maintenance mode manually to resume normal operation, go to Advanced > Reboot Controller to **reboot** the controller."

Below the warning box, there is a section titled "This Node". It contains a "Node type" dropdown menu set to "Storage" and a "Node name" text input field containing "hrmny2-1-254-sn". There are "Cancel" and "Save" buttons below the input fields.

Next is the "Primary Admin Node connection" section. It has an "Enable Admin Node discovery" checkbox which is unchecked. The "Primary Admin Node IP" text input field contains "172.16.0.62". The "Connection state" is "Connection to 172.16.0.62 ready." There are "Cancel" and "Save" buttons below.

The "Installation" section is at the bottom. It shows the "Current state" as "Maintenance mode. Reboot the node to resume normal operation." There are two buttons: "Start Expansion" and "Enable Cloning". The "Enable Cloning" button is highlighted with a yellow rectangular box.

The Primary Admin Node connection section is replaced with the Clone target node connection section.

Home

**⚠** This node is in maintenance mode. Perform any required maintenance procedures. If you want to exit maintenance mode manually to resume normal operation, go to Advanced > Reboot Controller to **reboot** the controller.

This Node

Node type: Storage ▾  
Node name: hrmny2-1-254-sn  
[Cancel] [Save]

**Clone target node connection**  
Clone target node IP: 0.0.0.0  
Connection state: No connection information available.  
[Cancel] [Save]

Installation

Current state: Waiting for configuration and validation of clone target.  
[Start Cloning] [Disable Cloning]

3. For **Clone target node IP**, enter the temporary IP address assigned to the target node for the network to use for clone data-transfer traffic, and then select **Save**.

Typically, you enter the IP address for the Grid Network, but if you need to use a different network for clone data-transfer traffic, enter the IP address of the target node on that network.

**i** Using the 1-GbE Admin Network for clone data transfer results in slower cloning.

After the target appliance is configured and validated, in the Installation section, **Start Cloning** is enabled on the source node.

Home

Configure Networking ▾

Configure Hardware ▾

Monitor Installation

Advanced ▾

Home

⚠ This node is in maintenance mode. Perform any required maintenance procedures. If you want to exit maintenance mode manually to resume normal operation, go to Advanced > Reboot Controller to **reboot** the controller.

ℹ The cloning process is ready to be started. Select **Start Cloning** when you are ready. To terminate cloning before it completes and return this node to service, trigger a reboot.

**This Node**

Node type

Storage ▾

Node name

hmnny2-1-254-sn

Cancel

Save

**Clone target node connection**

Clone target node IP

10.224.1.253

Connection state

Connection to 10.224.1.253 ready.

Cancel

Save

**Installation**

Current state

Ready to start cloning all data from this node to the clone target node using the Admin Network connection.  
 ⚠ Attention: the Admin Network typically has less bandwidth than the Grid or Client Networks. Use the Grid or Client IP of the target node for faster cloning.

Start Cloning

Disable Cloning

If issues exist that prevent cloning, **Start Cloning** is not enabled and issues that you must resolve are listed as the **Connection state**. These issues are listed on the StorageGRID Appliance Installer Home page of both the source node and the target appliance. Only one issue displays at a time and the state automatically updates as conditions change. Resolve all cloning issues to enable **Start Cloning**.

When **Start Cloning** is enabled, the **Current state** indicates the StorageGRID network that was selected for cloning traffic, along with information about using that network connection.

[Considerations and requirements for appliance node cloning](#)

4. Select **Start Cloning** on the source node.
5. Monitor the cloning progress using the StorageGRID Appliance Installer on either the source or target node.

The StorageGRID Appliance Installer on both the source and target nodes indicates the same status.

NetApp® StorageGRID® Appliance Installer Help

Home | Configure Networking ▾ | Configure Hardware ▾ | Monitor Installation | Advanced ▾

Monitor Cloning

1. Establish clone peering relationship		Complete
2. Clone another node from this node		Running
Step	Progress	Status
Send data to clone target node	<div style="width: 10%;"></div>	Sending data, 0% complete, 8.99 GB transferred
3. Activate cloned node and leave this one offline		Pending

The Monitor Cloning page provides detailed progress for each stage of the cloning process:

- **Establish clone peering relationship** shows the progress of cloning set up and configuration.
  - **Clone another node from this node** shows the progress of data transfer. (This part of the cloning process can take several days to complete.)
  - **Activate cloned node and leave this one offline** shows the progress of transferring control to the target node and placing the source node in a pre-install state, after data transfer is complete.
6. If you need to terminate the cloning process and return the source node to service before cloning is complete, on the source node go to the StorageGRID Appliance Installer Home page and select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.

If the cloning process is terminated:

- The source node exits maintenance mode and rejoins StorageGRID.
- The target node remains in the pre-install state.  
To restart cloning the source node, start the cloning process again from step 1.

When cloning successfully completes:

- The source and target nodes swap IP addresses:
  - The target node now uses the IP addresses originally assigned to the source node for the Grid, Admin, and Client Networks.
  - The source node now uses the temporary IP address initially assigned to the target node.
- The target node exits maintenance mode and joins StorageGRID, replacing the source node.
- The source appliance is in a pre-installed state, as if you had prepared it for reinstallation.

[Preparing an appliance for reinstallation \(platform replacement only\)](#)

 If the appliance does not rejoin the grid, go to the StorageGRID Appliance Installer Home page for the source node, select **Advanced > Reboot Controller**, and then select **Reboot into Maintenance Mode**. After the source node reboots in maintenance mode, repeat the node cloning procedure.

User data remains on the source appliance as a recovery option if an unexpected issue occurs with the target node. After the target node has successfully rejoined StorageGRID, user data on the source appliance is

outdated and is no longer needed. If desired, ask StorageGRID Support to clear the source appliance to destroy this data.

You can:

- Use the source appliance as a target for additional cloning operations: no additional configuration is required. This appliance already has the temporary IP address assigned that were originally specified for the first clone target.
- Install and set up the source appliance as a new appliance node.
- Discard the source appliance if it is no longer of use with StorageGRID.

# Other versions of NetApp StorageGRID documentation

You can find documentation for other releases of NetApp StorageGRID software here:

- [StorageGRID 11.9 documentation](#)
- [StorageGRID 11.8 documentation](#)
- [StorageGRID 11.7 documentation](#)
- [StorageGRID 11.6 documentation](#)
- [StorageGRID 11.4 Documentation Center](#)
- [StorageGRID 11.3 Documentation Center](#)
- [StorageGRID 11.2 Documentation Center](#)

# 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.

[Notice for StorageGRID 11.5](#)

## Copyright information

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

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

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

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

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

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

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

## Trademark information

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