



Learn about StorageGRID

StorageGRID 11.8

NetApp
May 17, 2024

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-118/primer/index.html> on May 17, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Learn about StorageGRID 1
 - What is StorageGRID?..... 1
 - Hybrid clouds with StorageGRID 3
 - StorageGRID architecture and network topology..... 4
 - Grid nodes and services 7
 - How StorageGRID manages data 19
 - Explore StorageGRID 31

Learn about StorageGRID


What is StorageGRID?

NetApp® StorageGRID® is a software-defined object storage suite that supports a wide range of use cases across public, private, and hybrid multicloud environments.

StorageGRID offers native support for the Amazon S3 API and delivers industry-leading innovations such as automated lifecycle management to store, secure, protect, and preserve unstructured data cost effectively over long periods.

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

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



Support for Archive Nodes is deprecated and will be removed in a future release. Moving objects from an Archive Node to an external archival storage system through the S3 API has been replaced by ILM Cloud Storage Pools, which offer more functionality.

StorageGRID benefits

Advantages of the StorageGRID system include the following:

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



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

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

Hybrid clouds with StorageGRID

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

Cloud Storage Pools

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

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



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

S3 platform services

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

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

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

ONTAP data tiering using FabricPool

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

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

Related information

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

StorageGRID architecture and network topology

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

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

For additional information about StorageGRID network topology, requirements, and grid communications, see

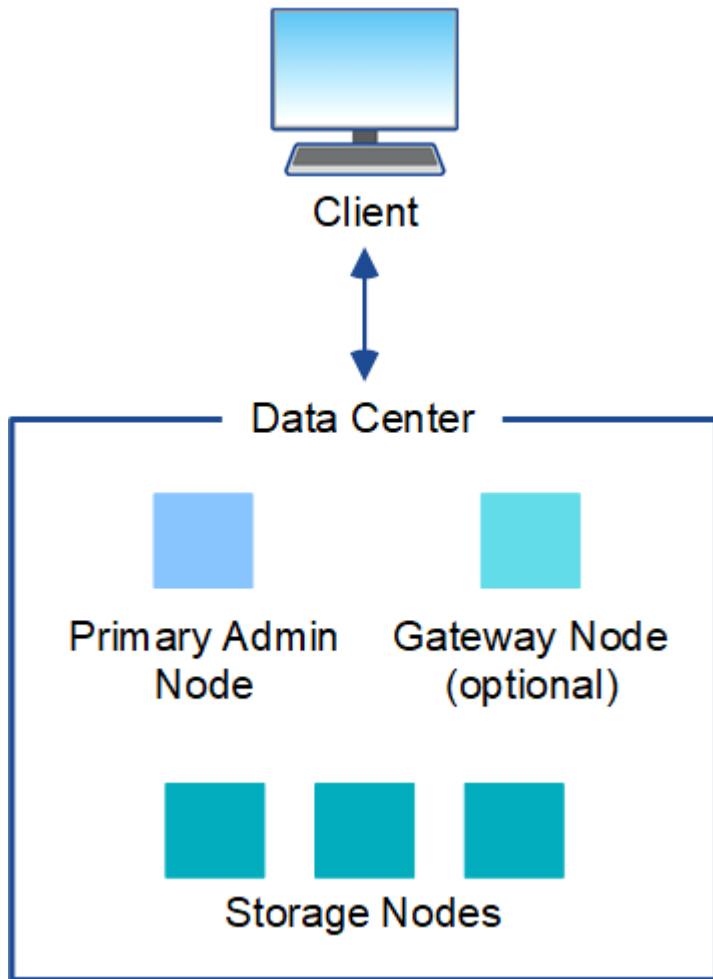
the [Networking guidelines](#).

Deployment topologies

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

Single site


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



Multiple sites

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

Different sites are often located in geographically different locations across different failure domains, such as an earthquake fault line or flood plain. Data sharing and disaster recovery are achieved by automated distribution of data to other sites.




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

Grid node redundancy

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

System architecture

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



S3 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, or in Cloud Storage Pools, which consist of external S3 buckets or Azure Blob storage containers.

Grid nodes and services

Grid nodes and services: Overview

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

Types of grid nodes

The StorageGRID system uses four types of grid nodes:

Admin Nodes

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

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

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

Storage Nodes

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

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

Gateway Nodes (optional)

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

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

Archive Nodes (deprecated)

Provide an optional interface through which object data can be archived to tape.

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

Hardware and software nodes

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

StorageGRID appliance nodes

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

See the following to learn about the available appliances:

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

Software-based nodes

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

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

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

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

StorageGRID services

The following is a complete list of StorageGRID services.

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

Service	Description	Location
DMV (Data Mover)	Moves data to cloud endpoints.	Storage Nodes
Dynamic IP (dynip)	Monitors the grid for dynamic IP changes and updates local configurations.	All nodes
Grafana	Used for metrics visualization in the Grid Manager.	Admin Nodes
High Availability	Manages high availability virtual IPs on nodes configured on the High Availability Groups page. This service is also known as the keepalived service.	Admin and Gateway Nodes
Identity (idnt)	Federates user identities from LDAP and Active Directory.	Storage Nodes that use the ADC service
Lambda Arbitrator	Manages S3 Select SelectObjectContent requests.	All nodes
Load Balancer (nginx-gw)	Provides load balancing of S3 and Swift traffic from clients to Storage Nodes. The Load Balancer service can be configured through the Load Balancer Endpoints configuration page. This service is also known as the nginx-gw service.	Admin and Gateway Nodes
LDR (Local Distribution Router)	Manages the storage and transfer of content within the grid.	Storage Nodes
MISCd Information Service Control Daemon	Provides an interface for querying and managing services on other nodes and for managing environmental configurations on the node such as querying the state of services running on other nodes.	All nodes
nginx	Acts as an authentication and secure communication mechanism for various grid services (such as Prometheus and Dynamic IP) to be able to talk to services on other nodes over HTTPS APIs.	All nodes
nginx-gw	Powers the Load Balancer service.	Admin and Gateway Nodes
NMS (Network Management System)	Powers the monitoring, reporting, and configuration options that are displayed through the Grid Manager.	Admin Nodes

Service	Description	Location
Persistence	Manages files on the root disk that need to persist across a reboot.	All nodes
Prometheus	Collects time series metrics from services on all nodes.	Admin Nodes
RSM (Replicated State Machine)	Ensures platform service requests are sent to their respective endpoints.	Storage Nodes that use the ADC service
SSM (Server Status Monitor)	Monitors hardware conditions and reports to the NMS service.	An instance is present on every grid node
Trace collector	Performs trace collection to gather information for use by technical support. The trace collector service uses open source Jaeger software.	Admin Nodes

What is an Admin Node?

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

Differences between primary and non-primary Admin Nodes

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

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

Capabilities	Primary Admin Node	Non-primary Admin Node
Includes the AMS service	Yes	Yes
Includes the CMN service	Yes	No
Includes the NMS service	Yes	Yes
Includes the Prometheus service	Yes	Yes
Includes the SSM service	Yes	Yes

Capabilities	Primary Admin Node	Non-primary Admin Node
Includes the Load Balancer and High Availability services	Yes	Yes
Supports the Management Application Program Interface (mgmt-api)	Yes	Yes
Can be used for all network-related maintenance tasks, for example IP address change and updating NTP servers	Yes	No
Can perform EC rebalance after Storage Node expansion	Yes	No
Can be used for the volume restoration procedure	Yes	Yes
Can collect log files and system data from one or more nodes	Yes	No
Sends alert notifications, AutoSupport packages, and SNMP traps and informs	Yes. Acts as the preferred sender .	Yes. Acts as a standby sender.

Preferred sender Admin Node

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

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

Multiple notifications might sent in these cases:

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



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

Primary services for Admin Nodes

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

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

What is a Storage Node?

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

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

Types of Storage Nodes

All Storage Nodes that were installed before StorageGRID 11.8 store both objects and the metadata for those objects. Starting in StorageGRID 11.8, you can choose the Storage Node type for new software-based storage nodes:

Object and metadata Storage Nodes

By default, all new Storage Nodes installed in StorageGRID 11.8 will store both objects and metadata.

Metadata-only Storage Nodes (software-based nodes only)

You can specify that a new software-based Storage Node be used to store only metadata. You can also add a metadata-only software-based Storage Node to your StorageGRID system during StorageGRID system expansion.



You can only select the Storage Node type when initially installing the software-based node or when you install the software-based node during StorageGRID system expansion. You can't change the type after the node installation is complete.

Installing a metadata-only node is typically not required. However, using a Storage Node exclusively for metadata can make sense if your grid stores a very large number of small objects. Installing dedicated metadata capacity provides a better balance between the space needed for a very large number of small objects and the space needed for the metadata for all those objects.

When installing a grid with software-based metadata-only nodes, the grid must also contain a minimum number of nodes for object storage:

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

Software-based Storage Nodes display a metadata-only indication for each metadata-only node on all pages that list the Storage Node type.

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.

Service	Key function
Administrative Domain Controller (ADC)	<p>Maintains topology and grid-wide configuration.</p> <p>Details</p> <p>The Administrative Domain Controller (ADC) service authenticates grid nodes and their connections with each other. The ADC service is hosted on a minimum of three Storage Nodes at a site.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p>
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.

Service	Key function
Distributed Data Store (DDS)	<p data-bbox="475 153 899 191">Monitors object metadata storage.</p> <p data-bbox="475 222 565 254">Details</p> <div data-bbox="475 264 1489 569"> <p data-bbox="508 296 1456 401">Each Storage Node includes the Distributed Data Store (DDS) service. This service interfaces with the Cassandra database to perform background tasks on the object metadata stored in the StorageGRID system.</p> <p data-bbox="508 432 1456 537">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).</p> </div>
Identity (idnt)	Federates user identities from LDAP and Active Directory.

Service	Key function
Local Distribution Router (LDR)	<p data-bbox="475 153 1484 195">Processes object storage protocol requests and manages object data on disk.</p> <p data-bbox="475 222 565 254">Details</p> <div data-bbox="475 264 1484 1927"> <p data-bbox="508 296 1484 464">Each Storage Node includes the Local Distribution Router (LDR) service. This service handles content transport functions, including data storage, routing, and request handling. The LDR service does most of the StorageGRID system's hard work by handling data transfer loads and data traffic functions.</p> <p data-bbox="508 499 1065 531">The LDR service handles the following tasks:</p> <ul data-bbox="532 562 1341 898" style="list-style-type: none"> <li data-bbox="532 562 654 594">• Queries <li data-bbox="532 615 1138 646">• Information lifecycle management (ILM) activity <li data-bbox="532 667 743 699">• Object deletion <li data-bbox="532 720 800 751">• Object data storage <li data-bbox="532 772 1341 804">• Object data transfers from another LDR service (Storage Node) <li data-bbox="532 825 889 856">• Data storage management <li data-bbox="532 877 971 909">• Protocol interfaces (S3 and Swift) <p data-bbox="508 930 1406 961">The LDR service also maps each S3 and Swift object to its unique UUID.</p> <p data-bbox="508 993 686 1024">Object stores</p> <p data-bbox="548 1035 1484 1140">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.</p> <p data-bbox="548 1171 1484 1381">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.</p> <p data-bbox="548 1413 1484 1581">To ensure even space usage for replicated copies, object data for a given object is stored to one object store based on available storage space. When an object store fills to capacity, the remaining object stores continue to store objects until there is no more room on the Storage Node.</p> <p data-bbox="508 1612 776 1644">Metadata protection</p> <p data-bbox="548 1654 1409 1717">StorageGRID stores object metadata in a Cassandra database, which interfaces with the LDR service.</p> <p data-bbox="548 1749 1484 1896">To ensure redundancy and thus protection against loss, three copies of object metadata are maintained at each site. This replication is non-configurable and performed automatically. For details, see Manage object metadata storage.</p> </div>

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

What is a Gateway Node?

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

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

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

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

Primary services for Gateway Nodes

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

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

What is an Archive Node?

Support for Archive Nodes is deprecated and will be removed in a future release.



Support for Archive Nodes is deprecated and will be removed in a future release. 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 also deprecated. If you are currently using an Archive Node with this option, [migrate your objects to a Cloud Storage Pool](#) instead.

Additionally, you should remove Archive Nodes from the active ILM policies in StorageGRID 11.7 or earlier. Removing object data stored on Archive Nodes will simplify future upgrades. See [Working with ILM rules and ILM policies](#).

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.


How StorageGRID manages data

What is an object

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

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

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



What is object data?

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

What is object metadata?

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

Object metadata includes information such as the following:

- System metadata, including a unique ID for each object (UUID), the object name, the name of the S3 bucket or Swift container, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
- The current storage location of each object copy or erasure-coded fragment.
- Any user metadata associated with the object.

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

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


How is object data protected?

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

Replication

When StorageGRID matches objects to an information lifecycle management (ILM) rule that is configured to create replicated copies, the system creates exact copies of object data and stores them on Storage Nodes, 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](#)
- [Use information lifecycle management](#)

The life of an object

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

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

- **Ingest:** The process of an S3 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 policies. During the copy management stage, StorageGRID protects object data from loss by creating and maintaining the specified number and type of object copies on Storage Nodes, 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](#)
- [Use information lifecycle management](#)

Ingest data flow

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

Data flow

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



1. The client application creates the object and sends it to the StorageGRID system through an HTTP PUT request.
2. The object is evaluated against the system's ILM policy.
3. The LDR service saves the object data as a replicated copy or as an erasure-coded copy. (The diagram shows a simplified version of storing a replicated copy to disk.)
4. The LDR service sends the object metadata to the metadata store.
5. The metadata store saves the object metadata to disk.
6. The metadata store propagates copies of object metadata to other Storage Nodes. These copies are also saved to disk.
7. The LDR service returns an HTTP 200 OK response to the client to acknowledge that the object has been ingested.

Copy management

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

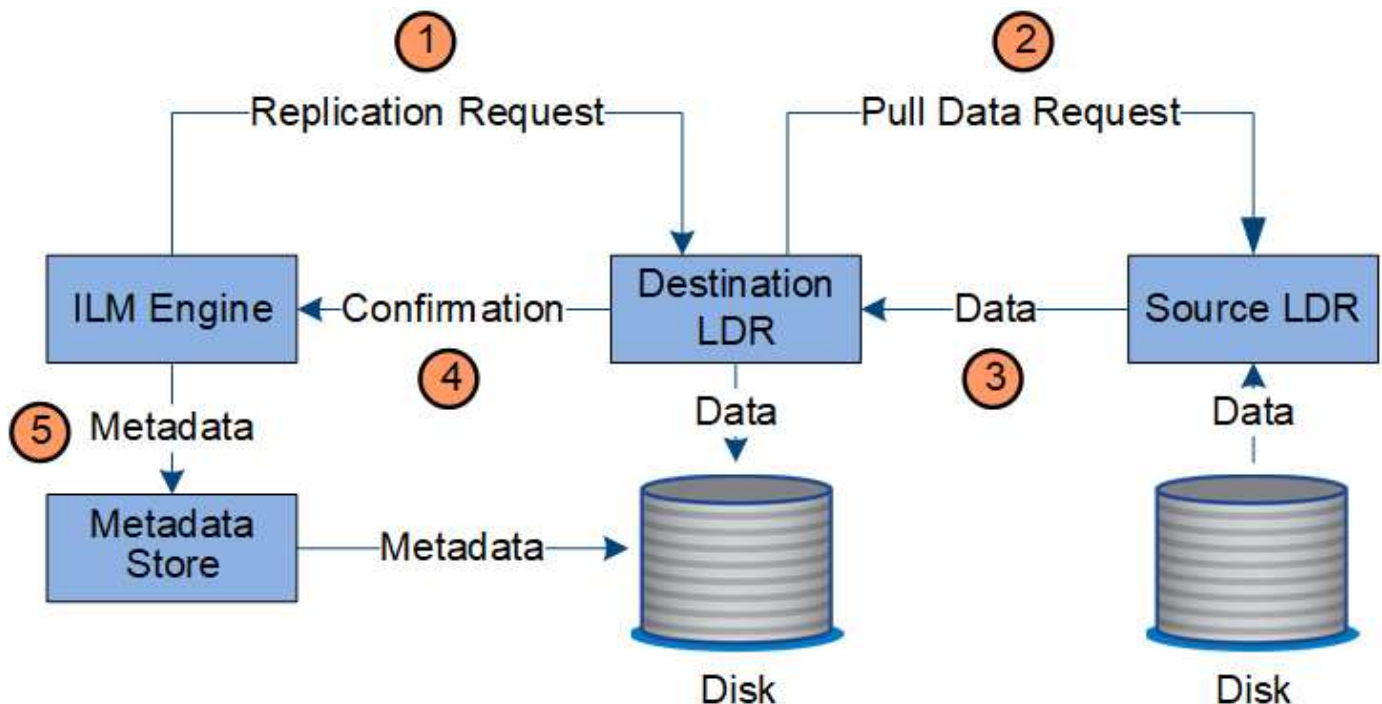
Different types or locations of object copies might be required at different times in the object's life. ILM rules are periodically evaluated to ensure that objects are placed as required.

Object data is managed by the LDR service.

Content protection: replication

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

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




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

Content protection: erasure coding

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

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




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

Content protection: Cloud Storage Pool

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

specified for the Cloud Storage Pool.

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



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


Retrieve data flow

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



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

Delete data flow

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

Deletion hierarchy

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

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

for each object.


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



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

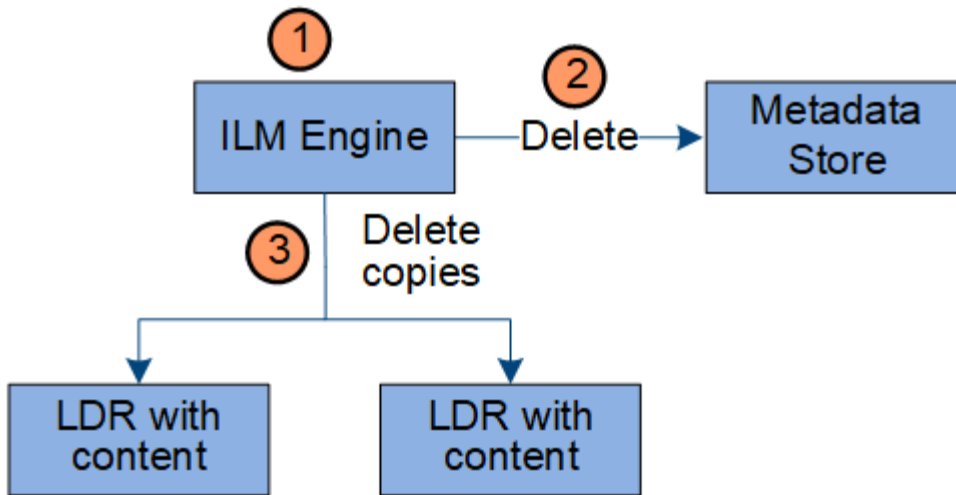
See [How objects are deleted](#) for more information.

Data flow for client deletes



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

Data flow for ILM deletes



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

Use information lifecycle management

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

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

ILM rules define:

- Which objects should be stored. A rule can apply to all objects, or you can specify filters to identify which objects a rule applies to. For example, a rule can apply only to objects associated with certain tenant accounts, specific S3 buckets or Swift containers, or specific metadata values.
- The storage type and location. Objects can be stored on Storage Nodes, 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.

Example ILM rule

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

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

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

How an ILM policy evaluates objects

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

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


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

Example ILM policy

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

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

Tenant A that are 1 MB or smaller).



Related information

- [Manage objects with ILM](#)

Explore StorageGRID

Explore the Grid Manager

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



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


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

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

Grid Manager dashboard

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

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



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

Search field

The **Search** field in the header bar allows you to quickly navigate to a specific page within Grid Manager. For example, you can enter **km** to access the Key management server (KMS) page. You can use **Search** to find entries in the sidebar of the Grid Manager and on the Configuration, Maintenance, and Support menus.

Help menu

The help menu  provides access to:

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

Alerts menu

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

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

- Review current alerts
- Review resolved alerts
- Configure silences to suppress alert notifications

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

Nodes page

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

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

Nodes

View the list and status of sites and grid nodes.

Q

Total node count: 14

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

Tenants page

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

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

Tenants

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

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

Displaying 2 results

<input type="checkbox"/>	Name	Logical space used	Quota utilization	Quota	Object count	Sign in/Copy URL
<input type="checkbox"/>	S3 Tenant	0 bytes	<div></div> 0%	100.00 GB	0	→ 📄
<input type="checkbox"/>	Swift Tenant	0 bytes	<div></div> 0%	100.00 GB	0	→ 📄

← Previous 1 Next →

ILM menu

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

From the ILM menu you can view and manage ILM:

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

Configuration menu

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

Network tasks

Network tasks include:

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

Security tasks

Security tasks include:

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

System tasks

System tasks include:

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

Monitoring tasks

Monitoring tasks include:

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

Access control tasks

Access control tasks include:

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

Maintenance menu

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

Tasks

Maintenance tasks include:

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

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

System

System maintenance tasks you can perform include:

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

Network

Network maintenance tasks you can perform include:

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

Support menu

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

Tools

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

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



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

Alarms (legacy)

From the [Alarms \(legacy\)](#) section of the Support menu, you can:

- Review current, historical, and global alarms
- Set up custom events
- Set up [email notifications for legacy alarms](#)



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

Other

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

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

Explore the Tenant Manager

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



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

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

Tenant Manager dashboard

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

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

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

Dashboard

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

Storage usage [?](#)

6.5 TB of 7.2 TB used

0.7 TB (10.1%) remaining



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

Total objects

8,418,886

objects

Tenant details [?](#)

Name: Tenant02

ID: 3341 1240 0546 8283 2208

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

Storage menu (S3)

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

My access keys

S3 tenant users can manage access keys as follows:

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



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

Buckets

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

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

system)

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

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

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

Platform services endpoints

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

Grid federation connections

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

Access Management menu

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

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.