

StorageGRID® 11.2

# **Administrator Guide**

August 2019 | 215-13573\_2019-08\_en-us doccomments@netapp.com



# **Contents**

Administering a StorageGRID system	9
Web browser requirements	9
Signing in to the Grid Manager	9
Signing out of the Grid Manager	13
Changing your password	13
Changing the browser session timeout	14
Viewing StorageGRID license information	15
Updating StorageGRID license information	15
Using the Grid Management API	16
Grid Management API versioning	19
Protecting against Cross-Site Request Forgery (CSRF)	20
Using the API if single sign-on is enabled	20
Controlling system access	26
Using identity federation	26
Configuring identity federation	26
Forcing synchronization with the identity source	31
Disabling identity federation	32
Managing admin groups	32
Creating admin groups	33
Admin group permissions	34
Modifying an admin group	38
Deleting an admin group	38
Managing local users	39
Creating a local user	39
Modifying a local user's account	40
Deleting a local user's account	40
Changing a local user's password	41
Using single sign-on (SSO) for StorageGRID	41
How single sign-on works	41
Requirements for using single sign-on	44
Configuring single sign-on	45
Creating and managing tenant accounts	60
Creating a tenant account	62
Creating a tenant account if StorageGRID is not using SSO	
Creating a tenant account if SSO is enabled	66
Changing the password for a tenant's local root user	67
Editing a tenant account	68
Deleting tenant accounts	
Managing platform services	71
What platform services are	71
Networking and ports for platform services	71

Per-site delivery of platform services messages	72
Troubleshooting platform services	74
Monitoring the StorageGRID system	77
Viewing the Dashboard	77
Using the Nodes page	79
Viewing common tabs for all node types	80
Viewing the Objects and ILM tabs for Storage Nodes	84
Viewing information about appliance Storage Nodes	86
Viewing site-level information	94
Resetting event counters	96
Viewing the Grid Topology tree	97
Understanding node icons	98
Information that you should monitor regularly	99
Key attributes to monitor	99
Monitoring storage capacity	100
Monitoring the recovery point objective through ILM	106
Monitoring object verification operations	108
Monitoring archival capacity	109
Monitoring the SSM service	110
Monitoring the Total Events alarm	112
About alarms and email notifications	112
Alarm notification types	112
Notification status and queues	113
Configuring notifications	114
Suppressing email notifications for a mailing list	118
Suppressing email notifications system wide	119
Selecting a preferred sender	120
Managing alarms	120
Alarm class types	121
Alarm triggering logic	123
Creating custom service or component alarms	127
Creating Global Custom alarms	129
Disabling alarms	131
Alarms and tables	131
Disabling Default alarms for services	
Disabling a Default alarm system wide	132
Disabling Global Custom alarms for services	133
Disabling Global Custom alarms system wide	134
Clearing triggered alarms	135
What AutoSupport is	135
Specifying the protocol for AutoSupport messages	136
Sending user-triggered AutoSupport messages	137
Disabling weekly AutoSupport messages	138
Disabling event-triggered AutoSupport messages	139
Troubleshooting AutoSupport messages	139

Using reports	. 141
Types of charts	141
Chart legend	. 142
Displaying charts	143
Generating charts	144
Types of text reports	145
Generating text reports	147
Exporting text reports	148
Managing objects through information lifecycle management	150
What an information lifecycle management policy is	150
What dual commit is	. 152
What an information lifecycle management rule is	. 152
What replication and erasure coding are	153
Configuring ILM rules	. 157
Configuring storage pools	. 158
Using Cloud Storage Pools	. 165
Configuring an Erasure Coding profile	177
Configuring regions (optional and S3 only)	179
Creating an ILM rule	181
Creating, simulating, and activating an ILM policy	192
Creating a proposed ILM policy	. 192
Simulating an ILM policy	196
Activating the ILM policy	203
Verifying an ILM policy with object metadata lookup	. 204
Working with ILM rules and ILM policies	207
Deleting an ILM rule	207
Editing an ILM rule	207
Cloning an ILM rule	208
Viewing the ILM policy activity queue	209
Example ILM rules and policies	. 210
Example 1: ILM rules and policy for object storage	210
Example 2: ILM rules and policy for EC object size filtering	214
Example 3: ILM rules and policy for better protection for image files	. 216
Managing S3 buckets and objects for compliance	220
How StorageGRID protects compliant data	220
Compliance workflow	221
Considerations for compliance	222
Enabling compliance	226
Example: Using compliant ILM rules in an ILM policy	227
S3 bucket for compliance example: bank-records	227
ILM rule 1 for compliance example: Erasure Coding profile with bucket	
matching	
ILM rule 2 for compliance example: Non-compliant rule	. 229
ILM rule 3 for compliance example: Default rule	
ILM policy for compliance example	232

Resolving consistency errors when updating a bucket's compliance	
configuration	234
Managing disk storage	236
What a Storage Node is	236
What the LDR service is	236
What the DDS service is	238
What the ADC service is	240
Managing Storage Nodes	240
What watermarks are	240
Storage Node configuration settings	243
Managing full Storage Nodes	246
Setting Grid Options	246
Configuring stored object encryption	246
Configuring stored object hashing	247
Configuring stored object compression	248
Enabling Prevent Client Modify	249
Enabling HTTP for client communications	250
What object segmentation is	251
Verifying object integrity	252
What background verification is	253
Configuring the background verification rate	253
What foreground verification is	255
Running foreground verification	255
How load balancing works	258
Managing archival storage	259
What an Archive Node is	259
What the ARC service is	259
About supported archive targets	260
Managing connections to archival storage	260
Configuring connection settings for S3 API	260
Migrating objects from Cloud Tiering - S3 to a Cloud Storage Pool	262
Modifying connection settings for S3 API	263
Configuring connections to Tivoli Storage Manager middleware	265
Managing Archive Nodes	267
Optimizing Archive Node for TSM middleware sessions	267
Managing an Archive Node when TSM server reaches capacity	268
Configuring Archive Node replication	269
Configuring retrieve settings	270
Configuring the archive store	271
Setting Custom alarms for the Archive Node	273
What an Admin Node is	274
Admin Node redundancy	275
Alarm acknowledgments	275
Email notifications and AutoSupport messages	276
Managing natworking	277

Viewing IP addresses	278
What link costs are	279
Updating link costs	280
Changing network transfer encryption	281
Configuring certificates	282
Configuring custom server certificates for the Grid Manager and the	
Tenant Manager	282
Restoring the default server certificates for the Grid Manager and the	
Tenant Manager	283
Configuring custom server certificates for storage API endpoints	283
Restoring the default server certificates for storage API endpoints	284
Copying the StorageGRID system's CA certificate	284
Configuring StorageGRID certificates for ONTAP clients using	
FabricPool	285
Configuring S3 API endpoint domain names	286
Configuring proxy settings	288
Configuring audit client access	290
Configuring audit clients for CIFS	
Configuring audit clients for Workgroup	
Configuring audit clients for Active Directory	293
Adding a user or group to a CIFS audit share	296
Removing a user or group from a CIFS audit share	298
Changing a CIFS audit share user or group name	299
Verifying CIFS audit integration	299
Configuring the audit client for NFS	300
Adding an NFS audit client to an audit share	301
Verifying NFS audit integration	303
Removing an NFS audit client from the audit share	
Changing the IP address of an NFS audit client	304
What data migration is	305
Confirming capacity of the StorageGRID system	305
Determining the ILM policy for migrated data	305
Impact of migration on operations	306
Scheduling data migration	306
Monitoring data migration	307
Creating custom notifications for migration alarms	307
Integrating Tivoli Storage Manager	309
Archive Node configuration and operation	
Configuration best practices	
Completing the Archive Node setup	
Installing a new TSM server	
Configuring the TSM server	
Copyright	
Trademark	316

How to send comments about documentation and receive update	
notifications	317

## Administering a StorageGRID system

This guide provides information and procedures you can use to administer and monitor the StorageGRID system on a day-to-day basis. This guide also includes information on how to configure a StorageGRID system to meet your deployment's unique operational requirements.

## Audience for this guide

This guide is for technical personnel who will be configuring, administering, and supporting a StorageGRID system after it has been installed.

This guide assumes a general understanding of the StorageGRID system. A fairly detailed level of computer literacy is assumed, including knowledge of Linux/UNIX command shells, networking, and server hardware setup and configuration.

## **About the Grid Manager**

This guide describes how to use the Grid Manager, which is the browser-based user interface that you use to perform most day-to-day activities in the StorageGRID system. You can use the Grid Manager to configure and monitor the system, create tenant accounts to allow S3 and Swift client applications to store and retrieve objects, configure ILM rules and policies, manage storage and networks, and more.

#### **Related information**

Grid primer

## Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	70
Microsoft Internet Explorer	11 (Native Mode)
Mozilla Firefox	63

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 sign-in page by entering the fully qualified domain name or IP address of an Admin Node into the address bar of a supported web browser.

## Before you begin

You know your login credentials.

- You have the fully qualified domain name or IP address of an Admin Node.
- You are using a supported web browser.
- Cookies are enabled in your web browser.
- You have specific access permissions.

### About this task

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 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 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.

### **Steps**

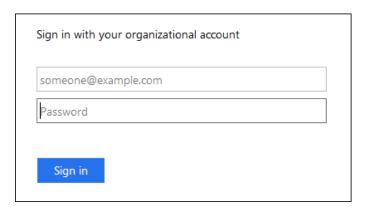
- 1. Launch a supported web browser.
- 2. In the browser's address bar, enter the fully qualified domain name or IP address of an Admin Node.
- 3. If you are prompted with a security alert, install the certificate using the browser's installation wizard.
- **4.** Sign in to the Grid Manager:
  - If single sign-on (SSO) is not being used for your StorageGRID system:
    - **a.** Enter your username and password for the Grid Manager.
    - b. Click Sign In.



- If SSO is enabled for your StorageGRID system and this is the first time you have accessed the URL on this browser:
  - a. Click Sign in. You can leave the Account ID field blank.



**b.** Enter your standard SSO credentials on your organization's SSO sign-in page. For example:

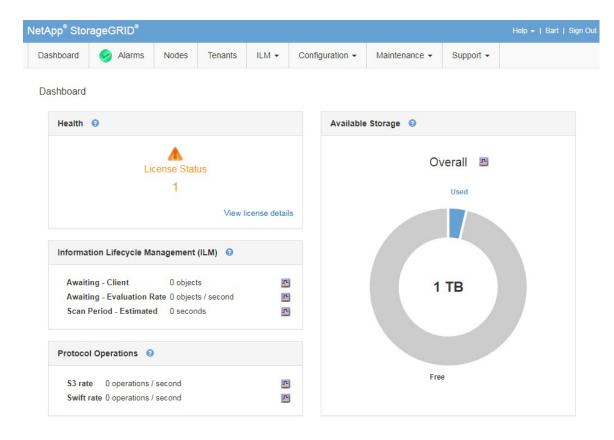


- If SSO is enabled for your StorageGRID system and you have previously accessed the Grid Manager or a tenant account:
  - **a.** Do either of the following:
    - Enter **0** (the account ID for the Grid Manager), and click **Sign in**.
    - Select **Grid Manager** if it appears in the list of recent accounts, and click **Sign in**.



**b.** 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.



## **5.** If you want to sign in to another Admin Node:

Option	Steps	
SSO not enabled	a. In the browser's address bar, enter the fully qualified domain name or IP address of the other Admin Node.	
	<b>b.</b> Enter your username and password for the Grid Manager.	
	c. Click Sign In.	
SSO enabled	In the browser's address bar, enter the fully qualified domain name or IP address of the other Admin Node.	
	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.	

## **Related concepts**

Configuring certificates on page 282

Admin group permissions on page 34

### Related tasks

Configuring single sign-on on page 45

### Related references

Web browser requirements on page 9

Using tenant accounts

## 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	You are signed out of the Admin Node.	
	The Grid Manager sign in page is displayed.	
	<b>Note:</b> If you signed into more than one Admin Node, you must sign out of each node.	
SSO enabled	You are signed out of all Admin Nodes you were accessing.	
	The StorageGRID Sign in page is displayed. <b>Grid Manager</b> is listed the default in the <b>Recent Accounts</b> drop-down, and the <b>Account ID</b> shows 0.	
	<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.	

#### Related tasks

Configuring single sign-on on page 45

### **Related information**

Using tenant accounts

## **Changing your password**

If you are a local user of the Grid Manager, you can change your own password.

### Before you begin

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 between 8 and 32 characters and is case-sensitive.

- **4.** Re-enter the new password.
- 5. Click Save.

## 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.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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.

**Note:** User session timeout can also be controlled by the following:

- A separate, non-configurable StorageGRID timer, 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 > 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.



Current Sender	ADMIN-DC1-ADM1	
Preferred Sender	ADMIN-DC1-ADM1	_
GUI Inactivity Timeout	900	
Notification Suppress All	Г	
		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 tasks

Signing out when SSO is enabled on page 43

#### Related information

Using tenant accounts

## 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.

#### Step

## 1. Select Maintenance > License.

The license information is displayed and includes the StorageGRID system ID, license serial number, licensed storage capacity of the grid, and the contents of the license text file. This information is read-only. 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.

## Before you begin

- You must have a new license file to apply to your StorageGRID system.
- You have specific access permissions.
- You must have the provisioning passphrase.

### **Steps**

- 1. Select Maintenance > License.
- 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. The Swagger user interface provides complete details and documentation for each API operation.

**Attention:** Any API operations you perform using the API Docs (Swagger) user interface are live operations. Be careful not to create, update, or delete configuration or other data by mistake.

To access the Swagger documentation for the Grid Management API:

- 1. Sign in to the Grid Manager.
- 2. Select **Help > API Docs** from the web application header.

#### **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, and return information about the health of the grid.
- 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/v2/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").

**Note:** 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.

• **compliance** – Operations to manage global compliance settings for the StorageGRID system.

- 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.
- 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.

Note: Metrics that include \_private\_ in their names are intended for internal use only. These metrics are subject to change between StorageGRID releases without notice.

- **ntp-servers** Operations to list or update external Network Time Protocol (NTP) servers.
- **objects** Operations on objects and object metadata.
- recovery Operations to list the grid nodes available for recovery.
- **recovery-package** Operations to download the Recovery Package.
- regions Operations to view and create regions.
- **server-certificates** Operations to view and update Grid Manager server certificates.
- **snmp** Operations on the current SNMP configuration.
- **users** Operations to view and manage Grid Administrator 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.



## **Issuing API requests**

**Attention:** Any API operations you perform using the API Docs (Swagger) user interface are live operations. Be careful not to create, update, or delete configuration or other data by mistake.

- 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.

## **Top-level resources**

The Grid Management API provides the following top-level resources:

• /grid: Access is restricted to Grid Administrator users.

- /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 Administrator users. These APIs are intended for internal use only and are not publicly documented. These APIs are also subject to change without notice.

#### Related concepts

Protecting against Cross-Site Request Forgery (CSRF) on page 20

#### Related tasks

Signing in to the API if single sign-on is enabled on page 21

#### Related information

Using tenant accounts Prometheus: Query basics

## **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
```

Changes in the Grid Management API that are backward incompatible bump the major version of the API. For example, an incompatible API change bumps the version from 2.1 to 3.0. Changes in the Grid Management API that are backward compatible bump the minor version instead. Backwardcompatible changes include the addition of new endpoints or new properties. For example, a compatible API change bumps the version from 3.0 to 3.1.

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.

**Note:** 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,
```

```
}
```

## 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.

**Note:** 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.

### **Steps**

- 1. Signing in to the API if single sign-on is enabled on page 21
- 2. Signing out of the API if single sign-on is enabled on page 24

## 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.

### Before you begin

- You know the SSO username and password for a federated user who belongs to a StorageGRID user group.
- If you want to access the Tenant Management API, you know the tenant account ID.

#### About this task

This is only a sample workflow that does not protect the password from being seen by other users.

This workflow might time out if you perform it too slowly. You might see the error: A valid SubjectConfirmation was not found on this Response.

If you have a URL-encoding issue, you might see the error: Unsupported SAML version.

#### **Steps**

1. 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'
```

**Note:** To access the Grid Management API, use 0 as TENANTACCOUNTID.

2. 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"
}
```

**3.** Save the SAMLRequest from the response for use in subsequent commands.

```
export SAMLREQUEST='fZHLbsIwEEV%2FJTuv7...sS1%2BfQ33cvfwA%3D'
```

**4.** 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=fZHRToMwFIZfhb...UJikvo77sXPw%3D
%3D&RelayState=12345&client-request-id=00000000-0000-0000-ee02-008000000de" >
```

**5.** Save the client request ID from the response.

```
export SAMLREQUESTID='00000000-0000-0000-ee02-0080000000de'
```

**6.** Send your credentials to the form action from the previous response.

```
curl -X POST "https://$AD_FS_ADDRESS/adfs/ls/?SAMLRequest=
$SAMLREQUEST&RelayState=$TENANTACCOUNTID&client-request-id=
$SAMLREQUESTID" \
   --data "UserName=$SAMLUSER@$SAMLDOMAIN&Password=
$SAMLPASSWORD&AuthMethod=FormsAuthentication" --include
```

AD FS returns a 302 redirect, with additional information in the headers.

**Note:** 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...UJikvo77sXPw%3D
%3D&RelayState=12345&client-request-id=00000000-0000-0000-ee02-008000000de
Set-Cookie: MSISAuth=AAEAADAvsHpXk6ApV...pmP0aEiNtJvWY=; path=/adfs;
HttpOnly; Secure
Date: Tue, 06 Nov 2018 16:55:05 GMT
```

7. Save the MSISAuth cookie from the response.

```
export MSISAuth='AAEAADAvsHpXk6ApV...pmP0aEiNtJvWY='
```

8. 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
```

```
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=a3dpbnRlcnMtUHJpbWFyeS1BZG1pbi0xNzqmRmFsc2Umcnq4NnJDZmFKVX
\texttt{FxVWx3bk11MnFuUSUzZCUzZCYmJiYmXzE3MjAyZTA5LThmMDgtNDRkZC04Yzg5LTQ3NDUx}
YzA3ZjkzYw==; path=/adfs; HttpOnly; Secure
Set-Cookie: MSISAuthenticated=MTEvNy8yMDE4IDQ6MzI6NTkgUE0=; path=/
adfs; HttpOnly; Secure
Set-Cookie: MSISLoopDetectionCookie=MjAxOC0xMS0wNzoxNjozMjo10VpcMQ==;
path=/adfs; HttpOnly; Secure
Date: Wed, 07 Nov 2018 16:32:59 GMT
<form method="POST" name="hiddenform" action="https://</pre>
storagegrid.example.com:443/api/saml-response">
  <input type="hidden" name="SAMLResponse"</pre>
value="PHNhbWxwOlJlc3BvbnN...1scDpSZXNwb25zZT4=" /><input
type="hidden" name="RelayState" value="12345" />
```

**9.** Save the SAMLResponse from the hidden field:

```
export SAMLResponse='PHNhbWxwOlJlc3BvbnN...1scDpSZXNwb25zZT4='
```

**10.** 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"
}
```

11. 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://adfs.example.com/adfs/ls/?
SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3D",
    "responseTime": "2018-11-20T22:20:30.839Z",
    "status": "success"
}
```

**2.** Save the logout URL.

```
export LOGOUT_REQUEST='https://adfs.example.com/adfs/ls/?
SAMLRequest=fZDNboMwEIRfhZ...HcQ%3D%3D'
```

3. Send a request to the logout URL to trigger SLO and to redirect back to StorageGRID.

```
curl --include "$LOGOUT_REQUEST"
```

The 302 response is returned. The redirect location is not applicable to API-only logout.

```
HTTP/1.1 302 Found

Location: https://$STORAGEGRID_ADDRESS:443/api/saml-logout?

SAMLResponse=fVLLasMwEPwVo7ss%...%23rsa-sha256

Set-Cookie: MSISSignoutProtocol=U2FtbA==; expires=Tue, 20 Nov 2018

22:35:03 GMT; path=/adfs; HttpOnly; Secure
```

**4.** Delete the StorageGRID bearer token.

Deleting the StorageGRID bearer token works the same way as without SSO. If <code>cookie</code> "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 indicate the user is now signed out.

HTTP/1.1 204 No Content

## **Controlling system access**

You determine who can access the StorageGRID system by importing groups and users from an identity federation service or by setting up local groups and local users. You determine which tasks users can perform by assigning permissions to each group. Users must belong to a group to be granted access to the system. Optionally, you can enable single sign-on (SSO) if you want all StorageGRID users to be authenticated by an external identity provider.

#### **Steps**

- 1. Using identity federation on page 26
- 2. Managing admin groups on page 32
- **3.** Managing local users on page 39
- 4. Using single sign-on (SSO) for StorageGRID on page 41

## **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.

#### **Steps**

- 1. Configuring identity federation on page 26
- 2. Forcing synchronization with the identity source on page 31
- **3.** Disabling identity federation on page 32

## 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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### 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.

**Note:** Configuration of identity federation has been verified with Active Directory, OpenLDAP, and Oracle Directory Server. If you want to use another LDAP service, contact support.

**Note:** StorageGRID uses STARTTLS for securing LDAP communications. It does not support the LDAP over SSL (LDAPS) protocol. The default port used for communications with the LDAP server is 389, but you can use any port as long as your firewall is configured correctly.

Note: 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."

#### Steps

- **1.** Select Configuration > Identity Federation.
- 2. Select Enable Identity Federation.

The fields for configuring the LDAP server appear.

3. Select the type of LDAP service you want to configure from the LDAP Service Type drop-down

You can select Active Directory, OpenLDAP, or Other.

Note: If you select OpenLDAP, you must configure the OpenLDAP server. See "Guidelines for configuring an OpenLDAP server."

- **4.** If you selected **Other**, complete the fields in the **LDAP Attributes** section.
  - Unique User 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 16byte 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. Enter the required LDAP server and network connection information in the LDAP Server section:
  - **Hostname**: The server host name or IP address of the LDAP server.
  - **Port**: The port used to connect to the LDAP server. Enter 389.
  - Username: The full path of the distinguished name (DN) for the user that will connect to the LDAP server.

Note: 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 nsunique
- ° 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.

**Note:** The Unique Group 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.

**Note:** The Unique User Name values must be unique within the User Base DN they belong to.

- **6.** Select a security setting from the **Transport Layer Security** (**TLS**) drop-down list to specify if TLS is used to secure communications with the LDAP server:
  - **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 in the CA Certificate text box.
  - **Do not use TLS**: The network traffic between the StorageGRID system and the LDAP server will not be secured.
- 7. Optionally, click **Test Connection** to validate your connection settings for the LDAP server.

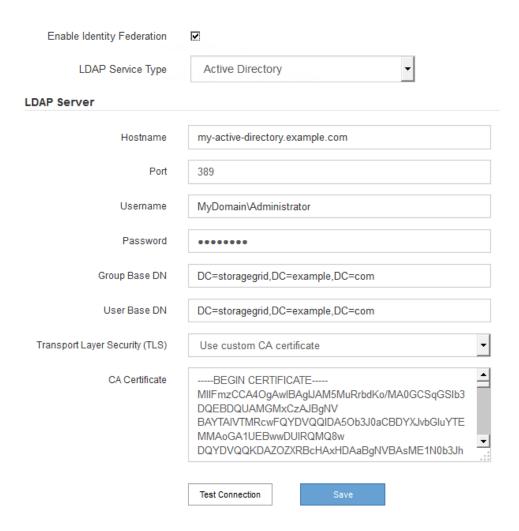
A green checkmark appears on the button if the connection is valid.



**8.** If the connection is valid, click **Save**.

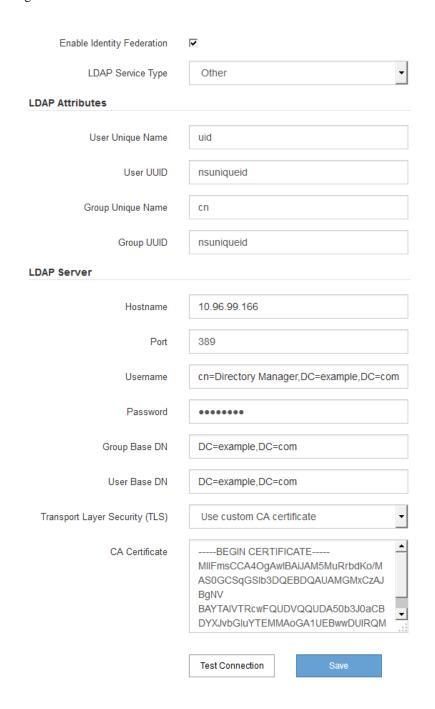
## Example

The following screenshot shows example configuration values for an LDAP server that uses Active Directory.



## Example

The following screenshot shows example configuration values for an LDAP server that uses Oracle Directory Server.



## **Related concepts**

Requirements for using single sign-on on page 44

## Related tasks

Creating a tenant account on page 62

## **Related information**

Using tenant accounts

## 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 member of 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.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- The identity source must be enabled.

#### Steps

## 1. Select Configuration > 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.

## 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 the StorageGRID and the identity source. However, any settings you have configured are retained, allowing you to easily reenable identity federation in the future.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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 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 > Identity Federation.
- 2. Uncheck the Enable Identity Federation check box.
- 3. Click Save.

## Related tasks

Disabling single sign-on on page 57

## 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.

## **Steps**

- 1. Creating admin groups on page 33
- 2. Admin group permissions on page 34
- 3. Modifying an admin group on page 38
- 4. Deleting an admin group on page 38

## Creating admin groups

Admin groups allow you to determine which admin users can access which StorageGRID features.

## Before you begin

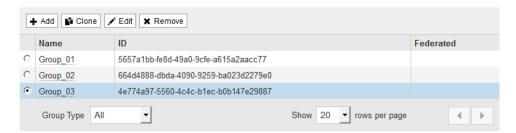
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- If you plan to import a federated group, you have configured identity federation and the federated group already exists in the configured identity source.

## **Steps**

### 1. Select Configuration > Admin Groups.

#### Admin Groups

Add and manage local and federated user groups, allowing member users to log into the Grid Administrator Interface. Set group permissions to control access to specific pages and features.



### 2. Click Add.

- 3. For 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 user interface. For example, "Maintenance Users" or "ILM administrators."
- Enter a unique name for the group.

For a local group, enter whatever unique name you want. For example, "ILM Admins." For a federated group, enter the group's name exactly as it appears in the configured identity source.

**6.** Select one or more management permissions.

You must assign at least one permission to each group; otherwise, users belonging to that group will not be able to sign in to StorageGRID.

#### Click 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 concepts

Admin group permissions on page 34

## Related tasks

Creating a local user on page 39

Modifying a local user's account on page 40

## **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
- Monitor alarms
- Change their own password (local users only)
- View certain information on the Configuration and Maintenance pages

The table shows the permissions you can assign when creating or editing an admin group. Any functionality not explicitly mentioned in the table requires the Root Access permission.

**Note:** You can use the Grid Management API to completely deactivate certain features. When a feature has been deactivated, the corresponding Management Permission no longer appears on the Groups page.

Management permission	Description
Root Access	Provides access to all grid administration features.
Grid Topology Page Configuration	Provides access to the Configuration tabs in Grid Topology. Also provides access to the <b>Reset event counts</b> links on the <b>Nodes &gt; Events</b> tabs to return event counts to zero.
Tenant Accounts	Provides access to the Tenant Accounts page from the <b>Tenants</b> option. Users who have this permission can add, edit, or remove tenant accounts. Users with this permission can also set the initial password for the tenant's local root user.
	Users who do not have this permission do not see the <b>Tenants</b> option in the menu.
	Note: 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.

Management permission	Description
Maintenance	Provides access to the following menu options:
	Maintenance > Maintenance Tasks
	<ul> <li>Expansion</li> </ul>
	<ul> <li>Decommission</li> </ul>
	· Recovery
	Maintenance > Network:
	Grid Network*
	DNS Servers*
	NTP Servers*
	• Maintenance > System:
	Apply Hotfix
	• License*
	Recovery Package
	Software Upgrade
	• Configuration > System Settings:
	<ul> <li>Domain Names*</li> </ul>
	<ul> <li>Server Certificates*</li> </ul>
	• Configuration > Monitoring:
	Audit*
	* Users who do not have the Maintenance permission can view, but not edit, the pages marked with an asterisk.
ILM	Provides access to the following menu options:
	• ILM > Rules
	• ILM > Policies
	• ILM > Erasure Coding
	• ILM > Regions
	Note: 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.
Acknowledge Alarms	Provides access to acknowledge and respond to alarms. All signed-in users can monitor alarms.
	If you want a user to monitor grid topology and acknowledge alarms only, you should assign this permission.

Management permission	Description
Other Grid Configuration	Provides access to the following grid configuration options:
	• Configuration > System Settings:
	<ul> <li>Grid Options</li> </ul>
	° Link Cost
	<ul> <li>Storage Options</li> </ul>
	<ul> <li>Display Options</li> </ul>
	• Configuration > Monitoring:
	∘ Global Alarms
	<ul> <li>Notifications</li> </ul>
	∘ Email Setup
	<ul> <li>AutoSupport</li> </ul>
	∘ Events
	• ILM:
	<ul> <li>Storage Pools</li> </ul>
	Storage Grades
	<b>Note:</b> Access to these items also requires the Grid Topology Page Configuration permission.
Change Tenant Root Password	Provides access to the <b>Change Root Password</b> button 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 <b>Change Root Password</b> button.
	<b>Note:</b> You must assign the Tenant Accounts permission to the group before you can assign this permission.
Metrics Query	Provides access to custom Prometheus metrics queries using the <b>Metrics</b> section of the Management API.
Object Metadata Lookup	Provides access to the <b>ILM &gt; Object Metadata Lookup</b> menu option.

## Related tasks

Deactivating features from the Grid Management API on page 37

### 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.

**Caution:** 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.

**Note:** 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 concepts

Using the Grid Management API on page 16

# Modifying an admin group

You can modify an admin group to update the display name or change the permissions associated with the group.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### **Steps**

- 1. Select Configuration > 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.** 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. Select a set of permissions.

See information about admin group permissions.

6. Click Save.

#### Related concepts

Admin group permissions on page 34

### 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.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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 > 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. Click Remove.
- 4. Click 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.

Note: If single sign-on (SSO) has been enabled, local users cannot sign in to StorageGRID.

### **Steps**

- 1. Creating a local user on page 39
- 2. Modifying a local user's account on page 40
- 3. Deleting a local user's account on page 40
- 4. Changing a local user's password on page 41

### 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.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

You can create local users only, 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 > 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 tasks

Creating admin groups on page 33

### 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.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

You can edit local users only. Federated user details are automatically synchronized with the external identity source.

### **Steps**

- 1. Select Configuration > 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.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### **Steps**

- 1. Select Configuration > Admin Users.
- **2.** Select the local user you want to delete.

**Note:** 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.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### 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 > 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.

#### **Steps**

- 1. How single sign-on works on page 41
- 2. Requirements for using single sign-on on page 44
- 3. Configuring single sign-on on page 45

### 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**

 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:



**Note:** 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. Go to step 4.

- 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:

someone@ex	ample.com	
Password		

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.

**Note:** 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.

### 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

**Note:** 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
- 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.

**Note:** 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 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 tasks

### 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.

### **Steps**

- 1. Confirming federated users can sign in on page 45
- 2. Using sandbox mode on page 46
- 3. Creating relying party trusts in AD FS on page 49
- 4. Testing relying party trusts on page 54
- 5. Enabling single sign-on on page 56
- **6.** Disabling single sign-on on page 57
- 7. Temporarily disabling and reenabling single sign-on for one Admin Node on page 57

### 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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You are using Active Directory as the federated identity source and AD FS as the identity provider. See "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.

Attention: 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 > 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	S3 tenant account	
Uses Own Identity Source		
Allow Platform Services	┍	
Storage Quota (optional)	GB ▼	
	Can	cel Save

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 concepts**

Requirements for using single sign-on on page 44 Admin group permissions on page 34

### Related tasks

Creating admin groups on page 33

### **Related information**

Using tenant accounts

### 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 reenable sandbox mode to configure or test new and existing relying party trusts. Reenabling sandbox mode temporarily disables SSO for StorageGRID users.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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.

**Note:** 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 > Single Sign-on.

The Single Sign-on page appears, with the **Disabled** option selected.



**Note:** 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."

#### 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).

### In the **Identity Provider** section:

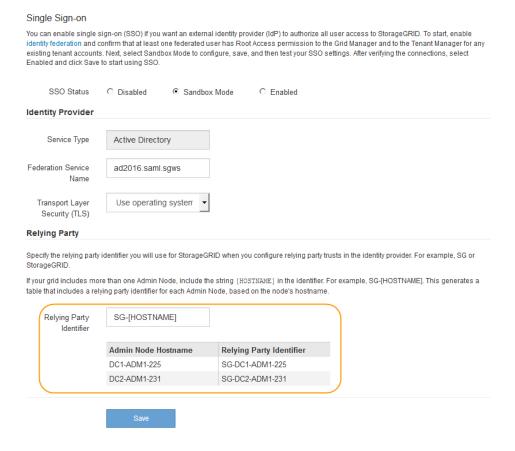
a. Enter the Federation Service name, exactly as it appears in AD FS.

Note: 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.

- 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.

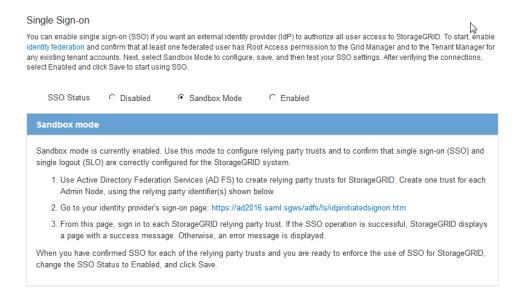


### 5. Click Save.

• A green check mark appears on the **Save** button for a few seconds.



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.



### Related concepts

Requirements for using single sign-on on page 44

### 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.

### Choices

- Creating a relying party trust using Windows PowerShell on page 49
- Creating a relying party trust by importing federation metadata on page 51
- Creating a relying party trust manually on page 52

#### Creating a relying party trust using Windows PowerShell

You can use Windows PowerShell to quickly create one or more relying party trusts.

### Before you begin

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.

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.

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**

- 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\_Identifer" -MetadataURL "https://Admin\_Node\_FQDN/api/saml-metadata"

- For Admin\_Node\_Identifier, enter the Relying Party Identifier for the Admin Node, exactly as it appears on the Single Sign-on page. For example, SG-DC1-ADM1.
- For Admin\_Node\_FQDN, enter the fully qualified domain name for the same Admin Node. (If necessary, you can use the node's IP address instead. However, if you enter an IP address here, be aware that you must update or recreate this relying party trust if that IP address ever changes.)
- 3. From Windows Server Manager, select **Tools > AD FS Management**.

The AD FS management tool appears.

4. Select AD FS > Relying Party Trusts.

The list of relying party trusts appears.

- **5.** Add an Access Control Policy to the newly created relying party trust:
  - a. Locate the relying party trust you just created.
  - b. Right-click the trust, and select Edit Access Control Policy.
  - c. Select an Access Control Policy.
  - d. 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.

#### Related tasks

Testing relying party trusts on page 54

### 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.

### Before you begin

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.

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.

- 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 FODN, 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.

**Note:** 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.

#### Related tasks

Testing relying party trusts on page 54

### Creating a relying party trust manually

If you choose not to import the data for the relying part trusts, you can enter the values manually.

#### Before you begin

• 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.

**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.

- 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.

Note: 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 dropdown 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.

#### Related tasks

Testing relying party trusts on page 54

### 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.

### Before you begin

• You must be signed in to the Grid Manager using a supported browser.

- You have specific access permissions.
- You have configured one or more relying party trusts in AD FS.

#### **Steps**

1. Select Configuration > 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.

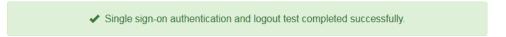


- 3. Click the link, or copy and paste the URL into a browser, to access your identity provider's signon page.
- To confirm you can use SSO to sign in to StorageGRID, select Sign in to one of the following sites, select the relying party identifier for your primary Admin Node, and click Sign in.



You are 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.



- If the SSO operation is unsuccessful, an error message appears. Fix the issue, clear the browser's cookies, and try again.
- 6. Repeat steps 4 and 5 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).

### Before you begin

- 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.

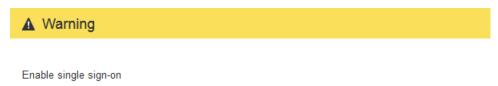
#### **Steps**

1. Select Configuration > 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.

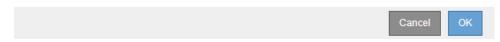


After you enable SSO, no local users—including the root user—will be able to sign in to the Grid Manager, the Tenant Manager, the Grid Management API, or the Tenant Management API.

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?



4. Review the warning, and click OK.

Single sign-on is now enabled.

**Attention:** 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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### **Steps**

1. Select Configuration > 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 reenable SSO for one Admin Node. To disable and then reenable SSO, you must access the node's command shell.

### Before you begin

- You have specific access permissions.
- You 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 reenable SSO on the Admin Node as soon as you sign out.

**Attention:** 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 > 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. Go to step *8*.

- 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 tasks

Configuring single sign-on on page 45

# Creating and managing tenant accounts

When you create a tenant account, you specify who can use your StorageGRID system to store and retrieve objects, and what functionality is available to them.

#### 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 (containers for Swift), 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.

**Note:** 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 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. 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.

### **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
- Using platform services (if enabled)
- Monitoring storage usage

Attention: 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

**Attention:** 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 Administrator permission to authenticate into the Swift REST API.

### **Steps**

- 1. Creating a tenant account on page 62
- 2. Changing the password for a tenant's local root user on page 67
- 3. Editing a tenant account on page 68
- 4. Deleting tenant accounts on page 70
- 5. Managing platform services on page 71

### Related tasks

Managing platform services on page 71

### **Related information**

Implementing S3 client applications Implementing Swift client applications Using tenant accounts

# Creating a tenant account

You must create at least one tenant account to control access to the storage in your StorageGRID system.

### Before you begin

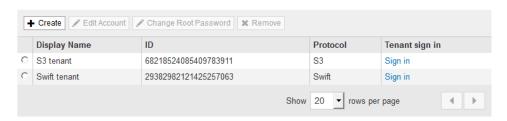
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### **Steps**

### 1. Select Tenants.

The Tenant Accounts page appears.

#### **Tenant Accounts**

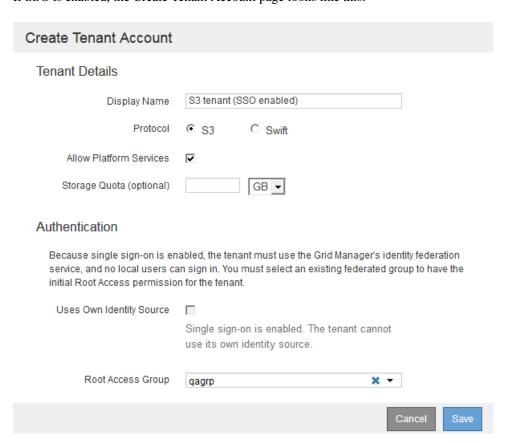


### 2. Click Create.

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.

• If SSO is enabled, the Create Tenant Account page looks like this.



#### Choices

- Creating a tenant account if StorageGRID is not using SSO on page 64
- Creating a tenant account if SSO is enabled on page 66

#### Related tasks

Configuring identity federation on page 26 Configuring single sign-on on page 45

### 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

### **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, uncheck the **Allow Platform Services** check box if 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.

**Note:** 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.

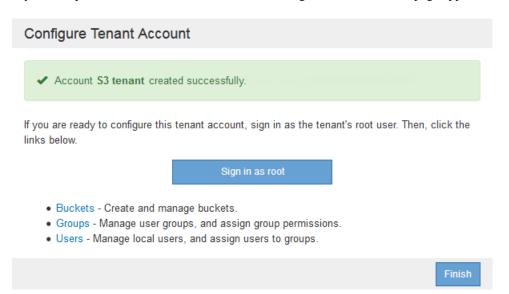
**Note:** You can monitor tenant storage usage from the Dashboard in the Tenant Manager or with the Tenant Management API. Note that a tenant's storage usage values might become out of date if nodes are isolated from other nodes in the grid. The totals will be updated when network connectivity is restored.

5. Determine if the tenant will use the identity source that was configured for the Grid Manager:

If the tenant will	St	eps
Manage its own groups and users		Select the Uses Own Identity Source check box (default).
		<b>Note:</b> 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>b.</b>	Specify a password for the tenant's local root user.

If the tenant will	St	eps
Use the groups and users configured for the Grid	a.	Uncheck the Uses Own Identity Source check box.
Manager	b.	Do either or both of the following:
		• Specify which existing federated group should have the initial Root Access permission for the tenant.
		<b>Note:</b> 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.

- **6.** Click **Save** to create the tenant account.
- 7. Decide whether to configure the tenant account now or later.
  - If you did not set a password for the local root user, the Tenant Accounts page appears, with a row for the new tenant.
    - If you are ready to configure the tenant and you belong to the Root Access federated group, click Sign In to immediately access the Tenant Manager.
    - Otherwise, provide the URL for the Sign in link to a user in the Root Access federated group. (The URL for a tenant is the fully qualified domain name or IP address of the Admin Node, followed by /?accountId=20-digit-account-id.)
  - If you set a password for the local root user, the Configure Tenant Account page appears.



- If you are ready to configure the tenant, go to step 8.
- Otherwise, click Finish. To access the tenant later, select Tenants from the menu and click the Sign in link for the account.
- 8. If you set a password for the local root user and you are ready to configure the tenant, click the Sign in as root button.

A green check mark appears on the button, indicating that you are now signed in to the tenant account as the root user.

9. Clink 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.

#### **10.** Click **Finish**.

The dialog closes. To access the Tenant Manager later, select **Tenants** from the menu, click the **Sign in** link, and sign in. Or, provide the URL for the **Sign in** link and the root user password to the tenant account's administrator. (The URL for a tenant is the fully qualified domain name or IP address of the Admin Node, followed by /?accountId=20-digit-account-id.)

#### Related tasks

Managing platform services on page 71

#### Related information

Using tenant accounts

### 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, uncheck the **Allow Platform Services** check box if 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.

**Note:** 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.

**Note:** You can monitor tenant 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.

Note: 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=20digit-account-id.)

Note: 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 tasks

Configuring single sign-on on page 45 Managing platform services on page 71

#### Related information

Using tenant accounts

# 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.

Note: 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.

### Before you begin

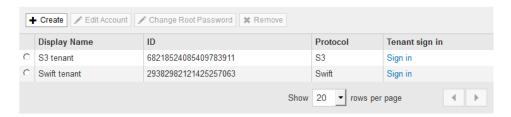
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### Steps

1. Select Tenants.

The Tenant Accounts page appears.

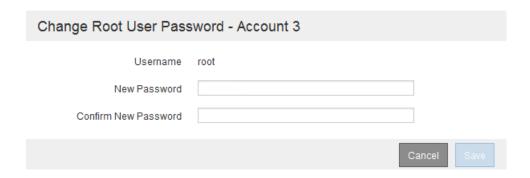
#### **Tenant Accounts**



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. You can then use your browser's find feature to search for a specific item in the currently displayed rows.

3. Select Change Root Password.



- 4. Enter the new password for the tenant account.
- 5. Click Save.

### Related concepts

Controlling system access on page 26

# 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.

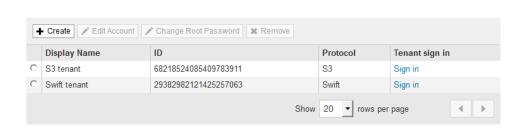
### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### Steps

1. Select Tenants.

The Tenant Accounts page appears.



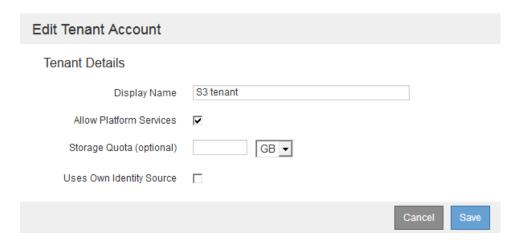
2. Select the tenant account you want to edit.

If your system includes more than 20 items, you can specify how many rows are shown on each page at one time. You can then use your browser's find feature to search for a specific item in the currently displayed rows.

#### 3. Select Edit Account.

Tenant Accounts

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.



- **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.

**Attention:** 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

**Note:** You can monitor tenant 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.

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.

### Note: 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. Click Save.

#### Related tasks

Managing platform services on page 71

# **Deleting tenant accounts**

You can delete a tenant account if you want to permanently remove the tenant's access to the system.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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. 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.

### **Related concepts**

Controlling system access on page 26

# Managing platform services

If you enable platform services for S3 tenant accounts, you must configure your grid such that tenants can access the external resources necessary to use these 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.
- Notifications: Per-bucket event notifications are used to send notifications about specific actions performed on objects to a specified external Simple Notification Service (SNS).
- 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.

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 must 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:

- NetApp recommends that you use 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 theStorageGRID 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.

### Related concepts

Viewing site-level information on page 94

#### Related information

Using tenant accounts

### 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.

**Note:** If a StorageGRID deployment is used as the destination for CloudMirror replication, replication messages are received by an API Gateway Node on port 8082. Ensure that this port is accessible through your enterprise network.

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.

#### Related tasks

Configuring proxy settings on page 288

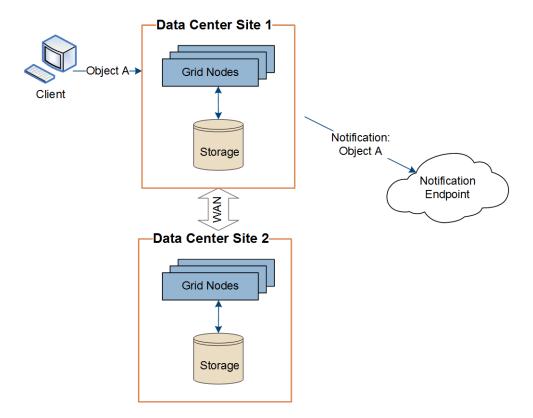
### Related information

Using tenant accounts

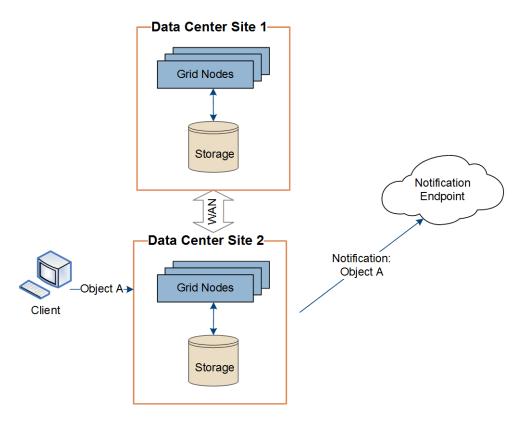
### Per-site delivery of platform services messages

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 an API 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.

Note: 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. 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. After resolving the issue, tenant users can test the endpoint. Clicking **Test** causes StorageGRID to validate that the 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.

### 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 > 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.

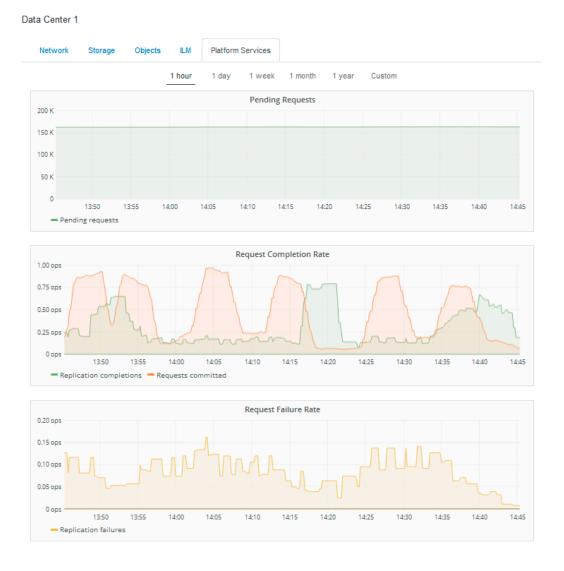
The only visible effect is that the incoming S3 requests will take longer to execute. If you start to detect significantly slower performance, you should reduce the ingest rate or use an endpoint with higher capacity. If the backlog of requests continues to grow, client S3 operations (such as PUT requests) 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.



# **Related concepts**

Viewing site-level information on page 94

# Monitoring the StorageGRID system

The Grid Manager enables you to monitor the daily activities of the StorageGRID system, including its health. Alarms and notifications help you evaluate and quickly resolve trouble spots that sometimes occur during normal operation.

You can use these areas of the Grid Manager to monitor the StorageGRID system to the level of detail you need:

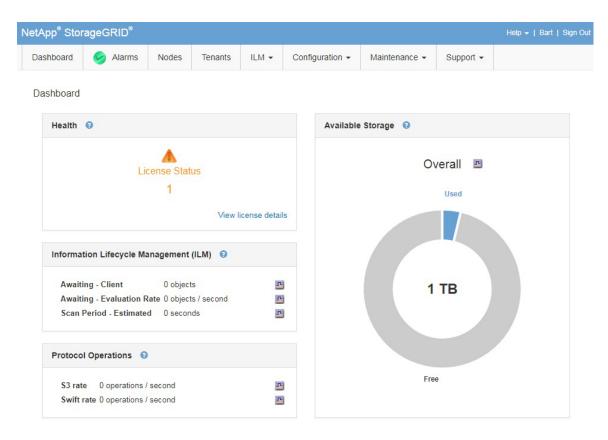
- Dashboard
- Nodes page
- Support > Grid Topology pages

### Related concepts

Viewing the Dashboard on page 77 Using the Nodes page on page 79

# 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 health and alerts, usage metrics, and operational trends and charts.



- **Health**: Provides an indication of the system's health by showing the number of disconnected grid nodes and the number of current alarms. Shows license status if there are license-related alerts.
- **Information Lifecycle Management (ILM)**: Displays current ILM operations and ILM queues for your StorageGRID system.
- **Protocol Operations**: Displays the number of protocol specific operations performed by your StorageGRID system. You can use this information to monitor your system's workloads and efficiencies.
- Available Storage: Shows the available and used storage capacity on the grid, not including archival media. With this information, you can compare the used storage with the available storage and, in a multi-site grid, determine which site is consuming more storage.

To view more detailed information about each panel on the Dashboard, click ②. Detailed descriptions are also listed in the following table.

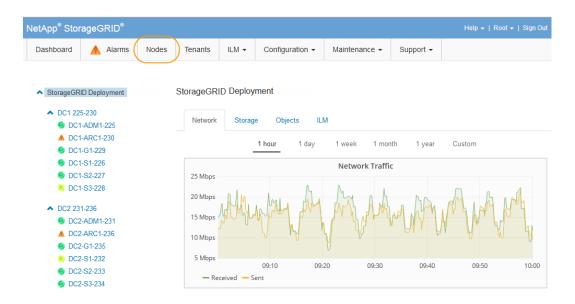
Panel	Description	View additional details	Learn more
Health	Provides an indication of the system's health by showing:  • The number of current alarms  • The number of disconnected grid nodes	<ul> <li>To see current alarms and alarm history, view custom alarms, and identify whether alarm notifications are sent, click View alarms details.</li> <li>To see details of all grid nodes, click View grid node details.</li> </ul>	<ul> <li>Managing alarms</li> <li>Using the Nodes page</li> <li>Grid primer</li> </ul>
Available Storage	Displays the available and used storage capacity in the entire grid, not including archival media.  The Overall chart presents grid-wide totals. If this is a multi-site grid, additional charts appear for each data center site.  With this information, you can compare the used storage with the available storage and in multi-site grids, determine which site is consuming more.	<ul> <li>To view the capacity, place your cursor over the chart's available and used capacity sections.</li> <li>To change the date range or select other data, click the chart icon in the upper right of the panel.         The default date range in charts on this page is one month.     </li> <li>To see available storage details, click Grid. Then, view the details for the entire grid, an entire site, or a single Storage Node.</li> </ul>	<ul> <li>Monitoring storage capacity on page 100</li> <li>Managing disk storage</li> </ul>
Information Lifecycle Management (ILM)	Displays current ILM operations and ILM queues for your system. With this information, you can monitor your system's workload.	<ul> <li>To see the existing ILM rules, click ILM &gt; Rules.</li> <li>To see the existing ILM policies, click ILM &gt; Policies.</li> </ul>	Managing objects through information lifecycle management

Panel	Description	View additional details	Learn more
Protocol Operations	Displays the number of protocol-specific operations (S3 or Swift) performed by your system.  With this information, you can monitor your system's workloads and efficiencies.  Protocol rates are averaged over the last two minutes.	<ul> <li>To change the date range or select other data, click the chart icon in the upper right of the panel.</li> <li>To manage tenant accounts for S3 or Swift, click Tenants.</li> </ul>	Creating and managing tenant accounts

# **Using 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, for each site in the grid, and for 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.



You can select the following tabs to view information for the entire grid:

- Network to view a graph showing network traffic
- Storage to view a graph showing storage used
- **Objects** to view a graph showing ingest and retrieve rates
- **ILM** to view information lifecycle management queue rates

The graphs on the Nodes page use the Grafana visualization tool and the Prometheus systems monitoring tool. Grafana displays time-series data in graph and chart formats, while Prometheus serves as the backend data source.

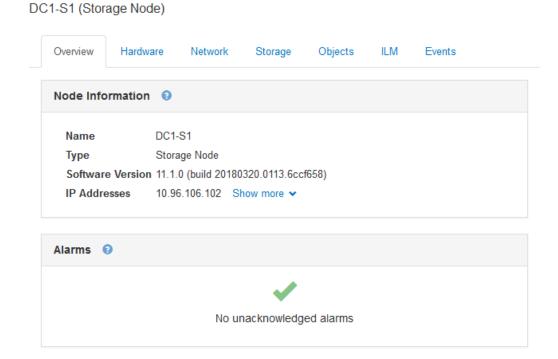
# Viewing common tabs for all node types

On the Nodes page, you can view the Overview, Hardware, Network, Storage, and Events tabs for all node types.

The StorageGRID system contains Admin Nodes, Storage Nodes, and, optionally, Archive Nodes and API Gateway Nodes.

### **Example Overview tab**

On the Overview tab, you can view general information about each node, such as the name, software version installed, and any unacknowledged alarms for the node. If any alarms are listed, you can click the link in the Service column to get details about each alarm.



### **Example Hardware tab**

The Hardware tab displays CPU utilization and memory information for each node.

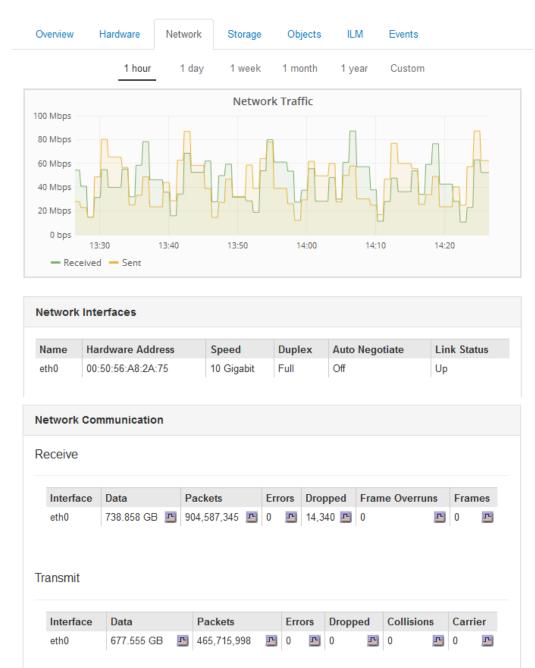
### DC1-S1-226 (Storage Node)

To display a different time interval, select one of the controls above the chart or graph. You can select to display the information available for intervals of 1 hour, 1 day, 1 week, 1 month, or 1 year. You can also set a custom interval, which allows you to specify date and time ranges.

### **Example Network tab**

The Network tab displays a graph showing network traffic for the node over time.

The Network Traffic graph indicates the total amount of network traffic across all of the network interfaces on the node. Communication to the node is "Received" and communication from the node is "Sent" or "Transmit."



DC1-S1-226 (Storage Node)

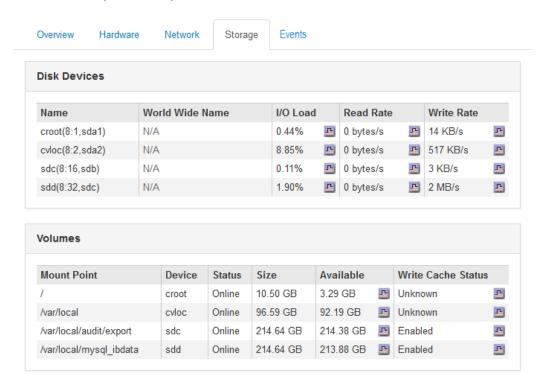
To display a different time interval, select one of the controls above the chart or graph. You can select to display the information available for intervals of 1 hour, 1 day, 1 week, 1 month, or 1 year. You can also set a custom interval, which allows you to specify date and time ranges.

### **Example Storage tab**

Although the Nodes page for each node type contains a Storage tab, the information on that tab differs depending on the node type:

 Admin Nodes, Archive Nodes, and Gateway Nodes each contain a list of disk devices and volumes on the node. Storage Nodes contain graphs showing data storage and metadata storage used over time, as well as a list of disk devices and volumes on the node.

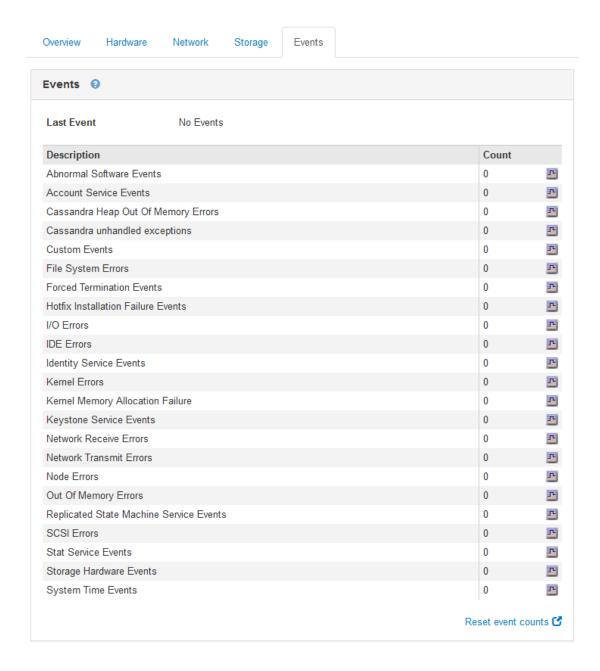
DC1-ADM1-225 (Admin Node)



### **Example Events tab**

The Events tab displays system events that help with troubleshooting. The Last Event provides an area of focus when an error occurs. After resolving issues, click Reset event counts to return the counts to zero. For details, see the information about troubleshooting.

Note: To reset event counts, you must be a user who belongs to a group that has the Grid Topology Page Configuration permission enabled.



### Related tasks

Resetting event counters on page 96

### **Related information**

Troubleshooting StorageGRID

# Viewing the Objects and ILM tabs for Storage Nodes

In addition to the Overview, Hardware, Network, Storage, and Events tabs, Storage Nodes also have Objects and ILM tabs.

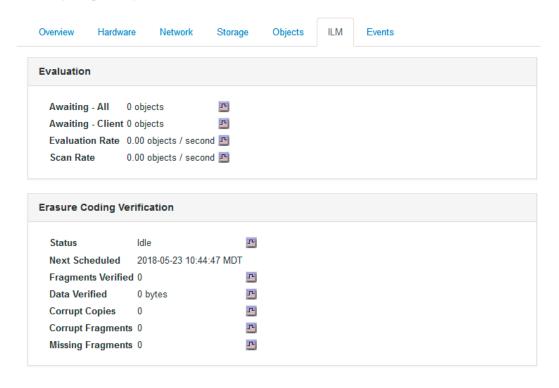
In the Objects tab, you can view information about object counts, queries, and verification.

### DC1-S1-226 (Storage Node)



In the ILM tab, you can view information about ILM operations.

### DC1-S1 (Storage Node)



### **Related concepts**

Managing objects through information lifecycle management on page 150

### **Related information**

Grid primer

## Viewing information about appliance Storage Nodes

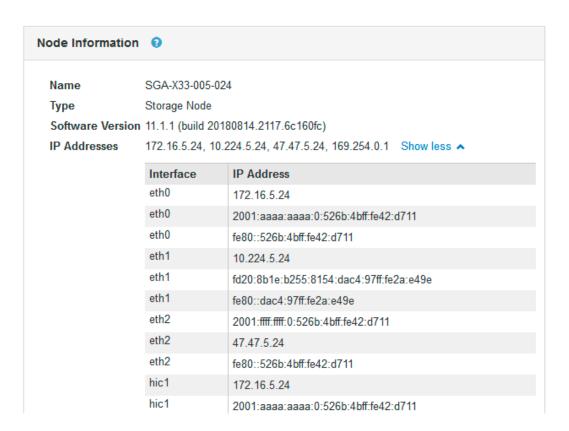
The Nodes page lists information about service health and all computational, disk device, and network resources for each appliance Storage Node. You can also see memory, storage hardware, network resources, network interfaces, network addresses, and receipt and transmittal information.

### **Steps**

1. From the Nodes page, select an appliance Storage Node.

#### 2. Select Overview.

The Node Information table on the Overview tab displays the name of the node, the node type, the software version installed, and the IP addresses associated with the node. The Interface column contains the name of the interface. An *eth* interface is a Grid Network, Admin Network, or Client Network, *hic* is a bonded port, and *mtc* (not shown in the following screenshot) is for management IP addresses on the appliance.



- 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 select to display the information available for intervals of 1 hour, 1 day, 1 week, 1 month, or 1 year. 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.

**Note:** Some fields, such as Compute Controller BMC IP and Compute Hardware, appear only for appliances with that feature.

Appliance Model	SG6060	
Storage Controller Name	StorageGRID-SGA-X33-005-024	
Storage Controller A Management IP	10.224.5.234	
Storage Controller B Management IP	10.224.5.235	
Storage Controller WWID	600a098000de38d900000005b73bcc0	
Storage Appliance Chassis Serial Number	721827500166	
Storage Hardware	Nominal	r
Storage Controller Failed Drive Count	0	r
Storage Controller A	Nominal	r
Storage Controller B	Nominal	r
Storage Controller Power Supply A	Nominal	r
Storage Controller Power Supply B	Nominal	r
Storage Multipath Connectivity	Nominal	r
Overall Power Supply	Nominal	r
Compute Controller BMC IP	10.224.5.25	
Compute Controller Serial Number	QTFCR28250004	
Compute Hardware	Nominal	r
Compute Controller CPU Temperature	Nominal	r
Compute Controller Chassis Temperature	Nominal	r
Compute Controller Power Supply A	Nominal	r
Compute Controller Power Supply B	Nominal	r

Field	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.
Storage Controller WWID	The worldwide identifier of the storage controller shown in SANtricity software.
Storage Appliance Chassis Serial Number	The chassis serial number of the appliance.

Field	Description
Storage Hardware	The overall status of the storage controller hardware.  If the Storage Node is a StorageGRID appliance and it needs attention, then both the StorageGRID and SANtricity systems indicate that the storage hardware needs attention.
	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.
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 Multipath Connectivity	The multipath connectivity state.  For details about resolving performance or fault tolerance issues, refer to the E-Series documents.
Overall Power Supply	The status of all power supplies for the appliance.
Compute Controller BMC IP	The IP address of the baseboard management controller (BMC) port in the compute controller.
	You use this IP to connect to the BMC interface to monitor and diagnose the appliance hardware.
	This field is not displayed for appliance models that 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.
Compute Controller Power Supply A	The status of power supply A for the compute controller.
Compute Controller Power Supply B	The status of power supply B for the compute controller.

# c. Confirm that all statuses are "Nominal."

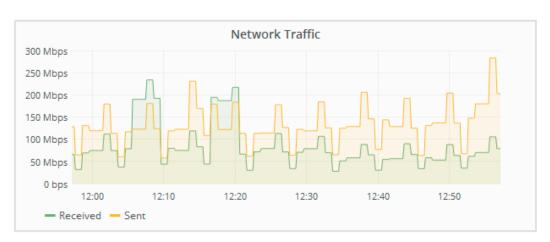
The statuses in this section correspond to the following alarm codes.

Field	Alarm code
Storage Controller Failed Drive Count	BADD
Compute Controller Chassis Temperature	BRDT
Compute Controller Hardware	CCNA
Compute Controller Power Supply A	CPSA
Compute Controller Power Supply B	CPSB
Compute Controller CPU Temperature	CPUT
Overall Power Supply	OPST
Storage Controller Power Supply A	PSAS
Storage Controller Power Supply B	PSBS
Storage Controller A	SCSA
Storage Controller B	SCSB
Storage Hardware	SOSS

For details about alarms in StorageGRID, see the information about troubleshooting.

**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.

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.

**Note:** 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

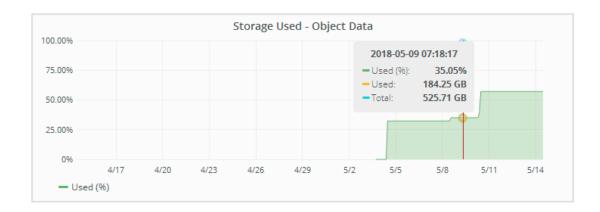
See the installation and maintenance instructions for your appliance for more information about configuring the 10/25-GbE ports.

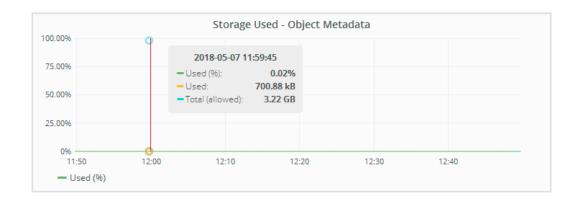
### b. Review the **Network Communication** section.

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 transmission metrics.

eceive													
Interface	Data	F	Packets		Err	ors	Dro	pped	Fran	ne Over	runs	Fra	mes
eth0	3.250 TB	5	,610,578,14	4 🝱	0	r	8,32	7 🖪	0		r	0	г
eth1	1.205 GB	9	,828,095	r	0	r	32,0	49 🝱	0		r	0	r
eth2	849.829 GB	9 1	86,349,407	r	0	r	10,2	69 🝱	0		r	0	ŗ
hic1	114.864 GB	3	03,443,393	r	0	r	0	r	0		r	0	r
hic2	2.315 TB	5	,351,180,95	6上	0	r	305	r	0		r	0	ŗ
hic3	1.690 TB	9 1	,793,580,23	0 🛂	0	r	0	-	0		r	0	r
hic4	194.283 GB	3	31,640,075	r	0	r	0	r	0		r	0	r
mtc1	1.205 GB	9	,828,096	r	0	r	0	-	0		r	0	r
mtc2	1.168 GB	9	,564,173	-	0	<u>.</u>	32,0	50 🝱	0		L	0	7
ansmit		9		T.	0					0-111-1			
ansmit Interface	Data		Packets			Err	ors	Drop	ped	Collisio	ons	Car	rier
ansmit Interface eth0	<b>Data</b> 5.759 TB	-	Packets 5,789,638,		r	Err 0	ors	Drop 0	ped	0	ons	Car 0	rier
Interface eth0 eth1	<b>Data</b> 5.759 TB 4.563 MB		Packets 5,789,638,41,520	626		<b>Err</b> 0 0	ors	<b>Drop</b> 0 0	ped	0	ons	<b>Car</b> 0 0	rier
Interface eth0 eth1 eth2	<b>Data</b> 5.759 TB 4.563 MB 855.404 GB		Packets 5,789,638, 41,520 139,975,19	626		0 0	ors	<b>Drop</b> 0 0 0	ped	0 0 0	ons F	Car 0 0	rier
Interface eth0 eth1 eth2 hic1	Data 5.759 TB 4.563 MB 855.404 GB 289.248 GB		Packets 5,789,638, 41,520 139,975,19 326,321,15	626 94 51		Errr 0 0 0 5	ors	<b>Drop</b> 0 0 0 0	ped	0 0 0	ons	Car 0 0 0	rier r r
Interface eth0 eth1 eth2 hic1 hic2	Data 5.759 TB 4.563 MB 855.404 GB 289.248 GB 1.636 TB		Packets 5,789,638, 41,520 139,975,19 326,321,15 2,640,416,	626 94 51 419		Errr 0 0 0 5 18	ors	Drop 0 0 0 0	ped F	0 0 0 0	ons	Car 0 0 0 5	rier r r
Interface eth0 eth1 eth2 hic1 hic2 hic3	Data 5.759 TB 4.563 MB 855.404 GB 289.248 GB 1.636 TB 3.219 TB		Packets 5,789,638, 41,520 139,975,19 326,321,15 2,640,416, 4,571,516,	626 94 51 419		Errr 0 0 0 5 18 33	ors R R	Drop 0 0 0 0 0	ped	0 0 0 0 0		Carro 0 0 0 5 18 33	rier
Interface eth0 eth1 eth2 hic1 hic2	Data 5.759 TB 4.563 MB 855.404 GB 289.248 GB 1.636 TB		Packets 5,789,638, 41,520 139,975,19 326,321,15 2,640,416,	626 94 51 419		Errr 0 0 0 5 18	ors	Drop 0 0 0 0	ped F	0 0 0 0	ons	Car 0 0 0 5	rier r r

**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.

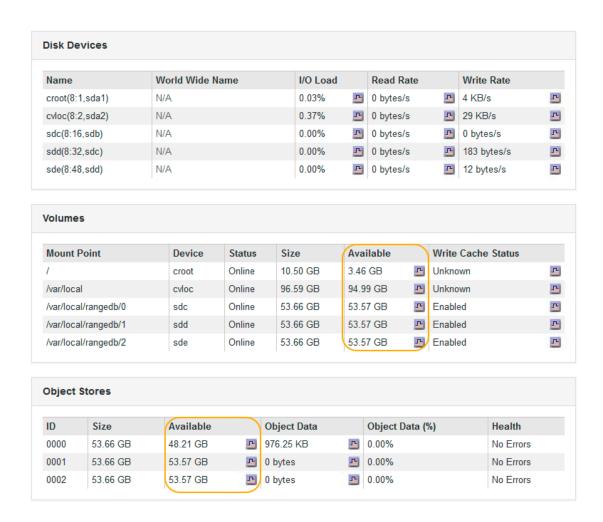




a. Scroll down to view the amounts of available storage for each volume and object store.

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.



### Related information

SG6000 appliance installation and maintenance

SG5700 appliance installation and maintenance

SG5600 appliance installation and maintenance

Troubleshooting StorageGRID

NetApp Documentation: SANtricity Storage Manager

# Viewing site-level information

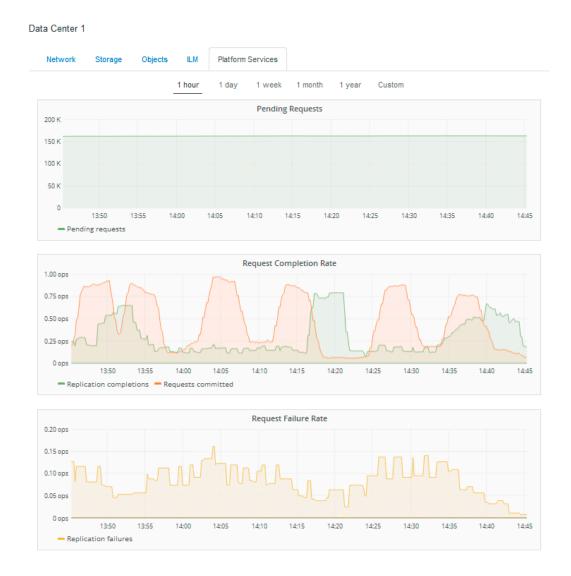
The site-level view provides information combined from all the nodes in one data center.

The nodes included in a data center are listed under the site name. When you click the site name, the site-level page appears.

Just as each individual node page provides graphical and tabular information for a single node, the site-level page provides similar information for the entire data center, including network activity, storage usage, ingest and retrieve rates, and information lifecycle management (ILM) queue rates.

DC3-S3

Additionally, the site-level page has a Platform Services tab. This tab provides information about services such as CloudMirror replication and search integration service. Graphs on this tab display metrics such as the number of pending requests, request completion rate, and request failure rate.



# Resetting event counters

On the Nodes page, the Events tab displays system events that help with troubleshooting. The Last Event provides an area of focus when an error occurs. After resolving issues, click **Reset event counts** to return the counts to zero.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### **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.

# Configuration: SSM (DC1-ADM1) - Events Updated: 2018-05-07 14:38:26 MDT

Description	Count	Rese
Abnormal Software Events	0	V
Account Service Events	0	
Cassandra Errors	0	V
Cassandra Heap Out Of Memory Errors	0	
Custom Events	0	5
File System Errors	0	
Forced Termination Events	0	
Grid Node Errors	0	
Hotfix Installation Failure Events	0	Г
I/O Errors	0	
IDE Errors	0	
Identity Service Events	0	
Kernel Errors	0	
Kernel Memory Allocation Failure	0	
Keystone Service Events	0	
Network Receive Errors	0	
Network Transmit Errors	0	
Out Of Memory Errors	0	
Replicated State Machine Service Events	0	
SCSI Errors	0	П
Stat Service Events	0	Г
Storage Hardware Events	0	
System Time Events	0	Г

- 4. Select the Reset check boxes for the specific counters you want to reset.
- 5. Click Apply Changes.

# **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 working with technical support.

To access the Grid Topology tree, select **Support > Grid Topology**.

# **Understanding node icons**

The Grid Manager uses the following icons to show the health of each grid node. The icons indicate whether nodes are connected to the grid and show if any alarms are active on the node.

Icon	Color	Node state	Alarm severity	Meaning
9	Green	Connected	Normal	The node is functioning normally. It is connected to the grid and there are no alarms.
	Yellow	Connected	Notice	The node is connected to the grid, but an unusual condition exists that does not affect normal operations.
*	Light Orange	Connected	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.
<b>A</b>	Dark Orange	Connected	Major	The node is connected to the grid, but an abnormal condition exists that currently affects operation. This requires prompt attention to prevent escalation.
8	Red	Connected	Critical	The node is connected to the grid, but an abnormal condition exists that has stopped normal operations. You should address the issue immediately.

Icon	Color	Node state	Alarm severity	Meaning
	Gray	Disconnected	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.
***	Blue	Disconnected	Unknown	The node is not connected to the grid, and the situation requires immediate attention. For example, the network connection between nodes has been lost or the power is down. This is the most severe condition.

# Information that you should monitor regularly

Monitor the StorageGRID system's key attributes regularly. Become familiar with system operations and spot trends before they turn into problems.

The following table lists the tasks to be performed on a regular basis.

Task	Frequency
Monitor system status. Note what has changed from the previous day.	Daily
Monitor the rate at which Storage Node capacity is being consumed.	Weekly
Check the capacity of the external archival storage system.	Weekly

The important attributes to monitor relate to:

- Object storage capacity
- Metadata storage capacity
- Object management related activities

When monitoring capacity, look at the absolute value and at the rate at which capacity is being consumed over time. Consumption rates can help you estimate when additional capacity might be required.

# Key attributes to monitor

To ensure that the StorageGRID system is able to store data as expected, you should monitor the amounts of available and used storage on a regular basis. The Dashboard provides a snapshot of used and free storage. The Nodes page contains graphs that enable you to monitor capacities over time.

Monitor the storage capacities shown in the following table.

Category	Component	Attributes
Storage capacity	<ul> <li>Dashboard: Available         Storage panel     </li> <li>Dashboard: Chart (Reports)         icon     </li> </ul>	<ul> <li>Dashboard Used and Free storage (Overall and for each Data Center)</li> <li>Dashboard Percentage Storage Capacity Used vs. Time (PSCU): The</li> </ul>
	• Nodes > Storage Node	percentage of installed storage capacity that has been consumed for the entire StorageGRID system or for each data center.
		Nodes page Storage Used - Object Data graph: hover over graph to see Used (%), Used, and Total storage. The Total value is the Total Usable Space (STAS) attribute.
Metadata storage capacity	Nodes > Storage Node	Hover over the Storage Used - Object Metadata graph to see the percentage of allowed space consumed by object metadata. This value is the Metadata Used Space (Percent) (CDLP) attribute.

**Note:** The attributes for storage capacity and metadata storage capacity do not include the capacity used for archived content.

# 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.

Storage capacity information is available at the grid, data center, and storage node levels of the StorageGRID system.

### **Steps**

- 1. Monitoring storage capacity for the entire grid on page 100
- 2. Monitoring storage capacity per Storage Node on page 102
- 3. Monitoring object metadata capacity per Storage Node on page 104

### Monitoring storage capacity for the entire grid

The Dashboard shows the available storage capacity for the entire grid and for individual data centers. You can use the Available Storage charts to quickly assess storage use for an entire StorageGRID system. In a multi-site grid, you can compare storage usage between sites (data centers).

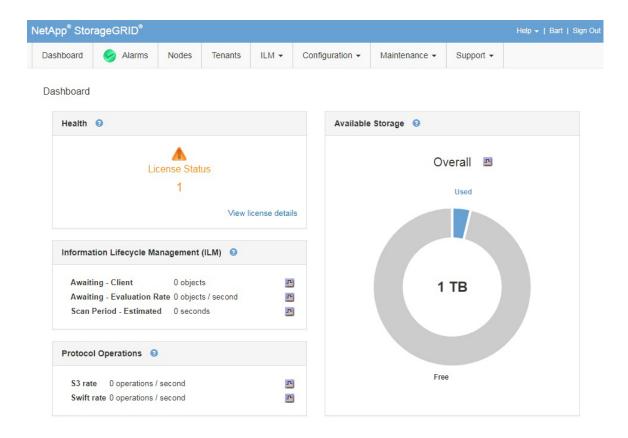
### Before you begin

You must be signed in to the Grid Manager using a supported browser.

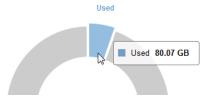
### Steps

- 1. Select Dashboard.
- **2.** In the **Available Storage** panel, note the overall summary of free and used storage capacity. For multi-site grids, review the chart for each data center.

Note: The summary does not include archival media.



3. Place your cursor over the chart's Free or Used capacity sections to see exactly how much space is free or used.

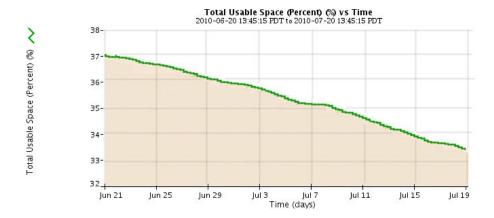


4. Click Chart (Reports) 1 to view a graph showing capacity usage over time for Overall storage or for individual data centers.

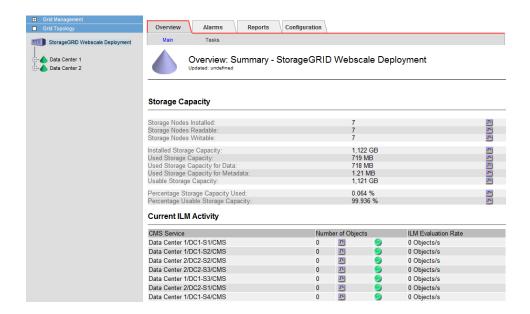
A graph showing Percentage Storage Capacity Used (%) vs. Time appears.

### Example

In the following example, usable storage space is being consumed at a rate of approximately 4% per month, which means that there are eight months left before this data center runs out of space.



- **5.** To monitor additional storage attributes:
  - Go to Nodes > storage node > Storage and view the graphs and tables.
  - Technical support could ask you to use this path: Select Support > Grid Topology. Then select StorageGRID Deployment > Overview > Main.



**6.** To maintain normal system operations, add Storage Nodes, add storage volumes, or archive object data before usable space is consumed.

### **Related information**

Expanding a StorageGRID system

### Monitoring storage capacity per Storage Node

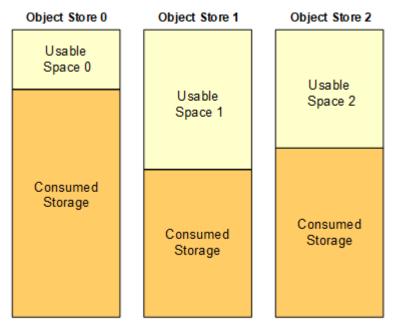
You must monitor the value of the Total Usable Space (STAS) attribute for each Storage Node to ensure that the node has enough space for new object data.

### Before you begin

You must be signed in to the Grid Manager using a supported browser.

### About this task

STAS is calculated by adding together the available space on all object stores within the Storage Node.



Total Usable Space = Usable Space 0 + Usable Space 1 + Usable Space 2

Note: If the value of STAS for a Storage Node falls below the value of the Storage Volume Hard Read-Only Watermark, the Storage Node becomes read-only and can no longer store new objects. You should add storage volumes or Storage Nodes before this happens.

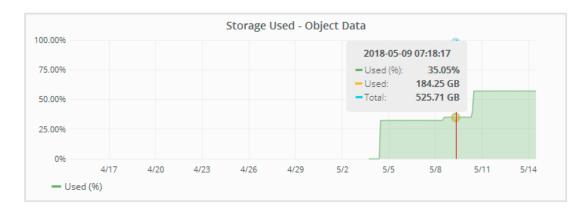
### **Steps**

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.

A pop-up displays Used (%), Used, and Total capacities.



3. Review the values in the tables below the graphs. To view graphs of the values, click Chart ... in the Available columns in the Volumes and Object Stores tables.

Name	World Wide Name	I/O Load		Read Rate		Write Rate	
croot(8:1,sda1)	N/A	0.03%	<u></u>	0 bytes/s	T.	4 KB/s	1
cvloc(8:2,sda2)	N/A	0.37%	T.	0 bytes/s	r	29 KB/s	1
sdc(8:16,sdb)	N/A	0.00%	r	0 bytes/s	r	0 bytes/s	1
sdd(8:32,sdc)	N/A	0.00%	-	0 bytes/s	T.	183 bytes/s	1
sde(8:48,sdd)	N/A	0.00%	J.	0 bytes/s	<b>J</b>	12 bytes/s	

Mount Point	Device	Status	Size	Available		Write Cache Status	
1	croot	Online	10.50 GB	3.46 GB	<u></u>	Unknown	9
var/local	cvloc	Online	96.59 GB	94.99 GB	J.	Unknown	9
/var/local/rangedb/0	sdc	Online	53.66 GB	53.57 GB	<u></u>	Enabled	9
/var/local/rangedb/1	sdd	Online	53.66 GB	53.57 GB	T	Enabled	1
/var/local/rangedb/2	sde	Online	53.66 GB	53.57 GB	T.	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	r	0.00%	No Errors

**4.** Monitor these values over time to estimate the rate at which usable storage space is being consumed.

Usable space is the actual amount of storage space available to store objects.

**5.** To maintain normal system operations, add Storage Nodes, add storage volumes, or archive object data before usable space is consumed.

### **Related concepts**

What watermarks are on page 240

### Related information

Expanding a StorageGRID system

### Monitoring object metadata capacity per Storage Node

StorageGRID reserves space on storage volume 0 of each Storage Node for the Cassandra database. This database stores three copies of all object metadata as well as certain configuration data. To ensure that adequate space remains available for essential Cassandra operations, you must monitor the metadata attributes for each Storage Node and add new Storage Nodes as required.

### Before you begin

• You must be signed in to the Grid Manager using a supported browser.

### About this task

The total space reserved on each Storage Node for metadata is known as Metadata Reserved Space, or CAWM. CAWM is subdivided into the space available for object metadata (the Metadata Allowed Space, or CEMS) and the space required for essential Cassandra operations, such as compaction and repair.

If object metadata uses more than 100% of the Metadata Allowed Space, Cassandra operations cannot run efficiently and errors will occur. For this reason, you must closely monitor how much object metadata space has been used.

The Metadata Used Space (Percent) attribute, or CDLP, measures how full the Metadata Allowed Space is. When the Metadata Used Space (Percent) reaches certain thresholds, the CDLP alarm is triggered, as follows:

- CDLP = 70% (minor alarm): You should add new Storage Nodes as soon as possible.
- CDLP = 90% (major alarm): You should add new Storage Nodes immediately.
- CDLP = 100% (critical alarm): You must add new Storage Nodes immediately and you must stop the ingest of new objects.

Note that when the Metadata Used Space (Percent) attribute reaches 70% (that is, when the Metadata Allowed Space becomes 70% full), the CDLP alarm is triggered as a minor alarm. You should add new Storage Nodes in an expansion procedure as soon as possible. When you add the new nodes, the system automatically rebalances object metadata across all Storage Nodes, and the alarms clear.

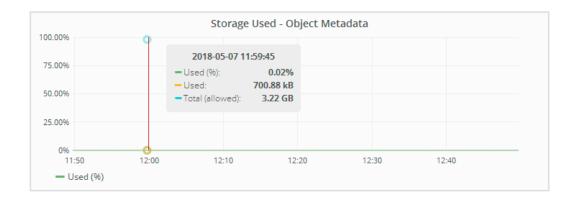
**Attention:** When the Metadata Used Space (Percent) attribute reaches 90%, the CDLP alarm is triggered as a major alarm, and a warning appears on the Dashboard. If this warning appears, you must add new Storage Nodes immediately. You must never allow object metadata to use more than 100% of the allowed space.

```
The following Storage Nodes are using more than 90% of the space allowed for object metadata
                           Used Allowed
  Node
                 % Used
                 104.51% 6.73 GB
  DC1-S2-227
                                       6.44 GB
                104.36% 6.72 GB
104.20% 6.71 GB
  DC1-S3-228
                                       6 44 GB
  DC2-S2-233
                                      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
```

Attention: If the first storage volume is smaller than the Metadata Reserved Space, the CDLP calculation might be inaccurate.

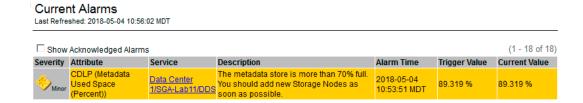
### **Steps**

- 1. Select Nodes > Storage Node > Storage.
- 2. View the Storage Used Object Metadata graph.



- **3.** If the **Used %** (CDLP Metadata Used Space (Percent)) value is 70% or higher, expand the StorageGRID system by adding Storage Nodes.
- **4.** To view alarm details, select the **Alarms** tab from the **Data Store Overview** page, or select **Alarms** from the top menu bar.

### **Example**



CDLP (Metadata Used Space (Percent)) alarms of all severities (minor, major, and critical) are displayed on this page. In the example, all the alarms listed are critical.

**5.** If a major or critical CDLP alarm appears on the Current Alarms page, add Storage Nodes immediately.

When you add the new nodes, the system automatically rebalances object metadata across all Storage Nodes, and the alarms clear.

### Related concepts

What watermarks are on page 240

### Related information

Expanding a StorageGRID system

# Monitoring the recovery point objective through ILM

You can track ILM evaluation attributes to determine the recovery point objective (RPO) of the StorageGRID system as defined by the ILM policy. The RPO defines the maximum tolerable period in which data might be lost because of a site failure, a Storage Node failure, or both.

### Before you begin

You must be signed in to the Grid Manager using a supported browser.

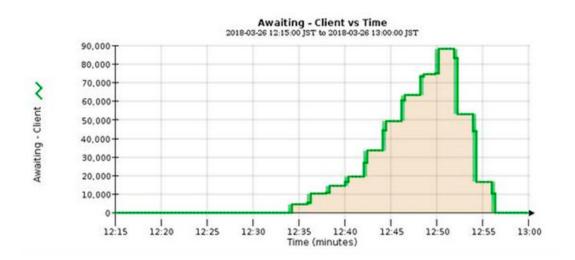
### About this task

The StorageGRID system manages objects by applying the defined ILM policy. The ILM policy and associated ILM rules determine how many copies are made, how those copies are made, the appropriate placement, and the length of time each copy is retained.

ILM rules are processed asynchronously after specific operations such as ingest or delete. ILM processing classifies objects into four categories and operates simultaneously on all four categories to ensure fairness and prioritization:

- Repair of replicated copies with only one copy remaining.
- Awaiting Background: excludes repair of replicated copies with only one copy remaining.
- Awaiting Client: includes new ingests and metadata updates; excludes deletions.
- Deletions.

Ingest or other activity can exceed the rate at which the system can process ILM. When this scenario occurs, the system will begin to queue objects whose ILM can no longer be fulfilled in near real time. In the example shown, the chart of the Awaiting—Client indicates that the number of objects awaiting ILM evaluation temporarily increases in an unsustainable manner, then eventually decreases. Such a trend indicates that ILM was temporarily not fulfilled in near real time.



### **Steps**

- 1. Select Support > Grid Topology.
- 2. Select deployment > Overview > Main.
- 3. In the ILM Activity section, review the key attributes for ILM evaluations:

### Awaiting - All

The total number of objects awaiting ILM evaluation.

### **Awaiting - Client**

The total number of objects awaiting ILM evaluation from client operations (for example, ingest).

#### Scan Rate

The rate at which objects in the grid are scanned and queued for ILM.

### **Scan Period - Estimated**

The estimated time to complete a full ILM scan of all objects.

**Note:** A full scan does not guarantee that ILM has been applied to all objects.

### **Awaiting - Evaluation Rate**

The current rate at which objects are evaluated against the ILM policy in the grid.

### **Repairs Attempted**

The total number of object repair operations for replicated data that have been attempted. This count increments each time an LDR tries to repair a high-risk object.

**Note:** The same object repair might increment again if replication failed after the repair.

### Related information

Grid primer

# Monitoring object verification operations

The StorageGRID system can verify the integrity of object data on Storage Nodes using background and foreground verification.

#### Before you begin

You must be signed in to the Grid Manager using a supported browser.

#### About this task

There are two verification processes: background verification and foreground verification. They work together to ensure data integrity. Background verification runs automatically, continuously checking the correctness of object data. Foreground verification can triggered by a user to more quickly verify the existence (although not the correctness) of object data.

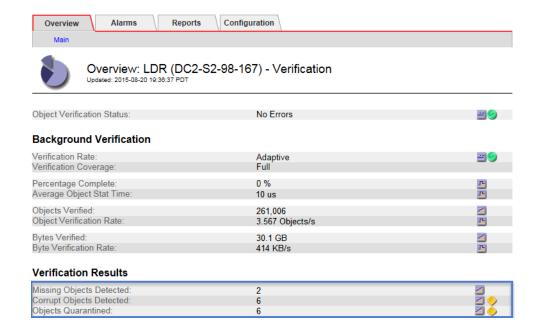
Background verification automatically and continuously checks all Storage Nodes to determine if there are corrupt copies of replicated and erasure-coded object data. If problems are found, the StorageGRID system automatically attempts to replace the corrupt object data from copies stored elsewhere in the system. Background verification does not run on Archive Nodes.

The foreground verification process 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.

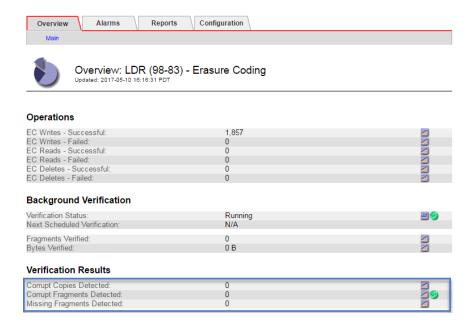
The **LDR** > **Verification** and **LDR** > **Erasure Coding** pages enable you to review results from background and foreground verifications, such as corrupt or missing objects detected. You should investigate any instances of corrupt or missing object data immediately, to determine the root cause.

### Steps

- 1. Select Support > Grid Topology.
- **2.** To check the verification results:
  - To check replicated object data verification, select Storage Node > LDR > Verification >
     Overview > Main.



To check erasure-coded fragment verification, select Storage Node > LDR > Erasure Coding > Overview > Main.



## 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 is still able to send object data to the archival destination, which may be an indication that an expansion of archival media is required.

### Before you begin

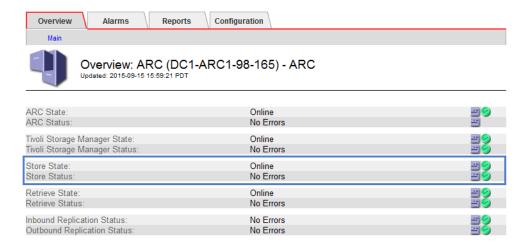
You must be signed in to the Grid Manager using a supported browser.

#### 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. As well, the triggering of the Store Failures (ARVF) alarm may be an indication that the targeted archival storage system has reach capacity and can no longer accept object data.

### **Steps**

- 1. Select Support > 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.



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.

## Monitoring the SSM service

The Server Status Monitor (SSM) service is present on all grid nodes and monitors the node's status, services, and resources. You can look at the SSM entry for each grid node to see the status of services on that node.

#### About this task

The SSM service monitors the condition of the server and related hardware, polls the server and hardware drivers for information, and displays the processed data. The information monitored includes:

- Service status
- Computation resources (restarts, runtime, uptime, load)
- Memory information (installed, available)
- CPU information (type, mode, speed)
- Volumes (status, available space)
- Network (addresses, interfaces, resources)
- · NTP synchronization

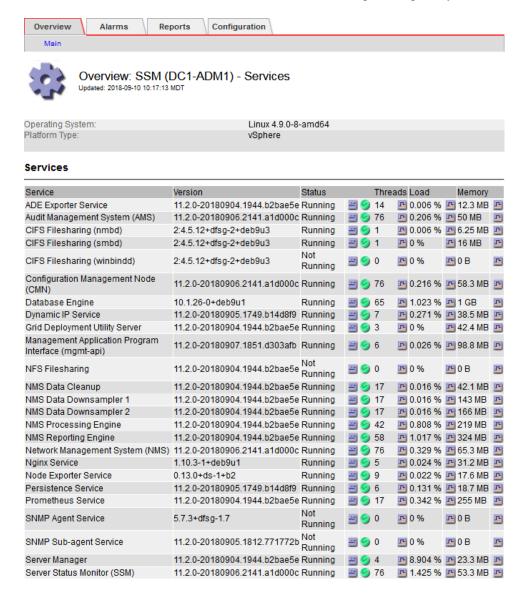
### **Steps**

- 1. Select Support > Grid Topology.
- 2. Select the grid node, and select SSM.

The state and status of the node's SSM service is shown.

3. Select SSM > Services.

The status of each service on the node is shown, as in this example for a primary Admin Node.



**4.** To monitor the events for the node, select **SSM** > **Events**.

**Note:** This information is the same as displayed on the **Nodes > Events** tab.

5. To monitor the node's resources, select SSM > Resources.

As required, you can use the Configuration tab to reset network error counters.

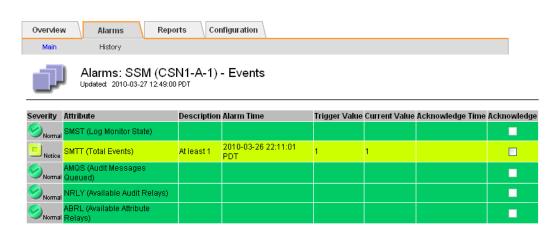
**6.** To monitor the clock settings for the node, select **SSM > Timing**.

The timing attributes report on the state of the grid node's local clock and the state of neighboring grid node clocks. In addition, these attributes report on NTP synchronization.

## Monitoring the Total Events alarm

When the Total Events alarm is raised, monitor the situation and take appropriate action.

Category	Code	Service	Notes
Total events	SMTT	SSM	The total number of logged error or fault events (Total Events SMTT) includes errors such as network errors. Unless these errors have been cleared (that is, the count has been reset to 0), total events alarms can be triggered.  Note: This alarm is safe to ignore only if the events that triggered the alarm have been investigated.



For more information about alarms, see the troubleshooting information.

### **Related information**

Troubleshooting StorageGRID

## About alarms and email notifications

An email notification is a message automatically sent by the StorageGRID system to notify recipients of a newly triggered alarm or service state change. You can configure email notifications and set mailing lists to receive these email notifications for any particular alarm severity or state change.

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.

## **Alarm notification types**

The StorageGRID systems sends out two types of alarm notifications: severity level and service state.

### Severity level notifications

Severity level notifications are sent at the alarm level and are associated with attributes. A mailing list receives all notifications related to alarm for the selected severity: Notice, Minor, Major, and Critical. A notification is sent when an event triggers an alarm for the selected alarm level. A notification is

also sent when the alarm leaves the alarm level — either by being resolved or by entering a different alarm severity level.

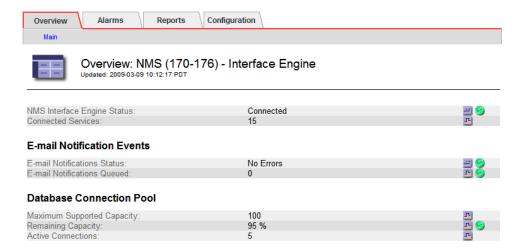
#### Service state notifications

Service State notifications are sent at the services level and are associated with services; for example, the LDR service or NMS service. A mailing list receives all notifications related to changes in the selected state: Unknown, or Administratively Down. A notification is sent when a service enters the selected Service State and when it leaves the selected Service State.

## Notification status and queues

You can view the current status of the NMS service's ability to send notifications to the mail server and the size of its notifications queue through the Interface Engine page.

To access the Interface Engine page, select Support > Grid Topology. Then, select site > Admin Node > NMS > Interface Engine.



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.

For a StorageGRID system configured with multiple Admin Nodes (and thus multiple NMS services), if the "standby" sender detects a Server Connection Error with the preferred sender, it will begin sending notifications to the mail server. The standby sender will continue to send notifications until it detects that the preferred sender is no longer in an error state and is again successfully sending notifications to the mail server. Notifications in the preferred sender's queue are not copied to the standby sender. Note that in a situation where the preferred sender and the standby sender are islanded from each other, duplicate messages can be sent.

### Related tasks

Selecting a preferred sender on page 120

## **Configuring notifications**

By default, notifications are not sent. You must configure the StorageGRID to send notifications when alarms are raised.

## Steps

- 1. Configuring email server settings on page 114
- 2. Creating email templates on page 115
- 3. Creating mailing lists on page 116
- 4. Configuring global email notifications on page 117

## Configuring email server settings

The E-mail Server page allows you to configure SMTP mail server settings that enable the sending of alarm notifications and AutoSupport messages. The StorageGRID system only sends email; it cannot receive email.

### Before you begin

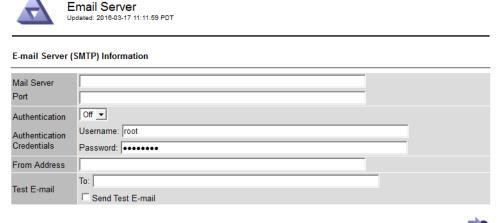
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### About this task

Only the SMTP protocol is supported for the sending email.

### **Steps**

- 1. Select Configuration > Email Setup.
- 2. From the Email menu, select Server.



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 host name rather than an IP address if you have previously configured DNS settings on the Admin Node.
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 alarm notification or AutoSupport 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 MINS (NMS Notification Status) Minor alarm.

### Related tasks

Creating mailing lists on page 116

### Creating email templates

Create an email template to customize the header, footer, and subject line of a notification. You can use email templates to send unique notifications that contain the same body text to different mailing lists.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

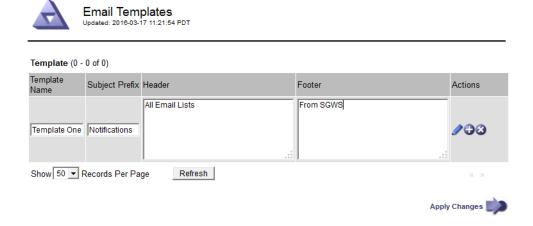
#### About this task

Different mailing lists might require different contact information. Templates do not include the body text of the email message.

### **Steps**

1. Select Configuration > Email Setup.

- 2. From the Email menu, select **Templates**.
- 3. Click **Edit** (or **Insert** if this is not the first template).



**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**

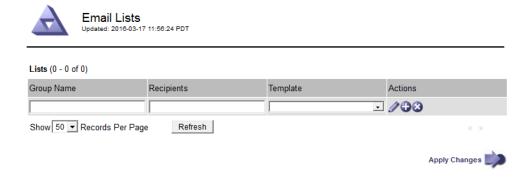
You can create mailing lists for notifications. A mailing list enables you to send one email message to multiple email addresses. These mailing lists are used to send notifications when an alarm is triggered or when a service state changes. You must create a mailing list before you can send notifications. To send a notification to a single recipient, create a mailing list with one email address.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

- 1. Select Configuration > Email Setup.
- 2. From the Email menu, select Lists.

3. Click **Edit** (or **Insert** if this is not the first mailing list).



**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.
Recipients	Single email address, a previously configured mailing list, or a commadelineated list of email addresses and mailing lists to which notifications will be sent.
Template	Optionally, select an email template to add a unique header, footer, and subject line to notifications sent to all recipients of this mailing list.

## 5. Click Apply Changes.

A new mailing list is created.

## Related tasks

Creating email templates on page 115

## Configuring global email notifications

In order to receive global email notifications, 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.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You must have configured an email list.

- 1. Select Configuration > Notifications.
- 2. Click **Edit** (or **Insert** if this is not the first notification).

- 3. Under **E-mail List**, add a mailing list.
- **4.** Select one or more alarm severity levels and service states:

Notification Type	Category	Description
Notice	Severity Level	An unusual condition exists that does not affect normal operation.
Minor	Severity Level	An abnormal condition exists that could affect operation in the future.
Major	Severity Level	An abnormal condition exists that is currently affecting operation.
Critical	Severity Level	An abnormal condition exists that has stopped normal operation.
Unknown	Service State	An unknown condition exists that has stopped normal service operation.
Administratively Down	Service State	A condition whereby a service has been purposefully stopped.

### 5. 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 tasks

Creating mailing lists on page 116

## Suppressing email notifications for a mailing list

You can suppress notifications for a mailing list system-wide when you do not want a mailing list to receive notifications, for example while performing maintenance procedures.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

## Steps

- 1. Select Configuration > Notifications.
- 2. Click **Edit** next to the mailing list for which you want to suppress notifications.
- **3.** 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.
- 4. Click Apply Changes.

Notifications are suppressed for the selected mailing lists.

## Suppressing email notifications system wide

You can block the StorageGRID system's ability to send notifications when an alarm is triggered.

### Before you begin

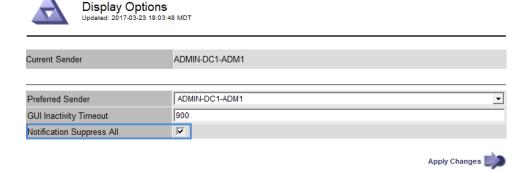
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### About this task

Attention: Suppressing email notifications system wide also suppresses event-triggered AutoSupport emails, even when the Enabled check box is selected on the Event-Triggered AutoSupport page (Support > AutoSupport > Event-triggered).

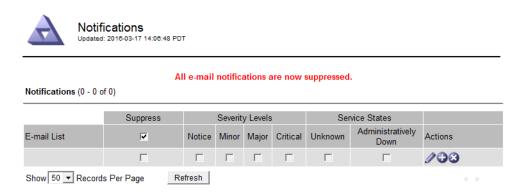
### **Steps**

- 1. Select Configuration > Display Options.
- From the Display Options menu, select **Options**.
- Select Notification Suppress All.



## 4. Click Apply Changes.

The Notifications page (**Configuration > Notifications**) displays the following message:





### Related concepts

Troubleshooting AutoSupport messages on page 139

## Selecting a preferred sender

Each site in a StorageGRID deployment can include one or more Admin Nodes. If a deployment includes multiple Admin Nodes, you must configure one Admin Node as the preferred sender of notifications and AutoSupport messages. You can select any Admin Node as the preferred sender, and you can change which Admin Node is selected at any time.

### Before you begin

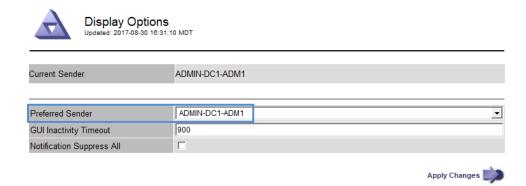
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

The **Display Options** page lists the Admin Node that is currently sending notifications. In most cases, this Admin Node is the same as the Preferred Sender; however, if an Admin Node is islanded from the rest of the system, it is unable to use the Preferred Sender and automatically updates to become the Current Sender. In the case of islanded Admin Nodes, multiple Admin Nodes will attempt to send notifications and AutoSupport message and thus it is possible that multiple copies of notifications will be received.

#### **Steps**

- 1. Select Configuration > 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.



4. Click Apply Changes.

The Admin Node is set as the preferred sender of notifications.

# Managing alarms

Customizing alarms lets you customize your StorageGRID system based on your unique monitoring requirements.

You can configure customized alarms either globally (Global Custom alarms) or for individual services (Custom alarms). You can create customized alarms with alarm levels that override Default alarms, and you can create alarms for attributes that do not have a Default alarm. Alarm

customization is restricted to accounts with the Grid Topology Page Configuration and Other Grid Configuration permissions.

**Attention:** Using the Default alarm settings is recommended. 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 until it prevents a critical operation from completing. If you do need to change an alarm setting, you should discuss your proposed changes with technical support.

For a list of alarm codes, see the troubleshooting information.

### Related concepts

Controlling system access on page 26

#### Related information

Troubleshooting StorageGRID

## Alarm class types

Alarms are separated into three mutually exclusive alarm classes.

- **Default**: Standard alarm configurations. Default alarms are set during installation.
- **Global Custom**: Custom alarms that are set at a global level and that apply to all services of a given type in the StorageGRID system. Global Custom alarms are configured after installation to override default settings.
- **Custom:** Custom alarms that are set on individual services or components. Custom alarms are configured after installation to override default settings.

### **Default alarms**

Default alarms are configured on a global basis and cannot be modified. However, Default alarms can be disabled or overridden by Custom alarms and Global Custom alarms.

Default alarms can be disabled both globally and at the services level. If a Default alarm is disabled globally, the **Enabled** check box appears with an adjacent asterisk at the services level on the Configuration page. The asterisk indicates that the Default alarm has been disabled through the **Configuration > Global Alarms** page even though the **Enabled** check box is selected.

You can view the default alarms for a particular service or component. Select **Support > Grid** Topology. Then select service or component > Configuration > Alarms.

### Related tasks

Disabling Default alarms for services on page 131 Disabling a Default alarm system wide on page 132

### Viewing all Default alarms

You can view all Default alarms that are standard alarm configurations set as part of the installation.

### Before you begin

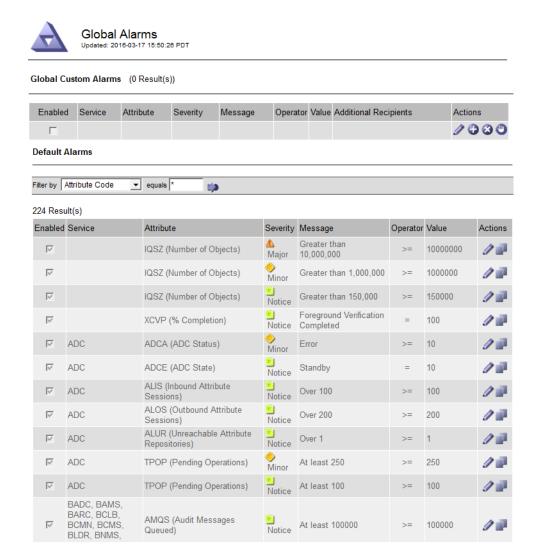
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### **Steps**

1. Select Configuration > Global Alarms.

- 2. For Filtered by select Attribute Code or Attribute Name.
- 3. For equals, enter the wildcard symbol: \*
- **4.** Click the arrow or press **Enter**.

All Default alarms are listed.



#### **Global Custom alarms**

Global Custom alarms monitor the status of conditions system-wide. By creating a Global Custom alarm, you can override a Default alarm system-wide. You can also create a new Global Custom alarm that will monitor status system-wide. This can be useful for monitoring any customized conditions of your StorageGRID system.

You can create Global Custom alarms, and disable Global Custom alarms system wide or for individual services.

#### Related tasks

Creating Global Custom alarms on page 129

Disabling Global Custom alarms for services on page 133

Disabling Global Custom alarms system wide on page 134

#### **Custom alarms**

Custom alarms can be created to override a Default alarm or Global Custom alarm at the service or component level. You can also create new Custom alarms based on the service's unique requirements.

You can configure Custom alarms by going to each service's Configuration > Alarms page in the Grid Topology tree.

### Related tasks

Creating custom service or component alarms on page 127

## Alarm triggering logic

Each alarm class is organized into a hierarchy of five severity levels from Normal to Critical. An alarm is triggered when a threshold value is reached that evaluates to true against a combination of alarm class and alarm severity level. Note that a severity level of Normal does not trigger an alarm.

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.

Attribute values are evaluated against the list of enabled alarms defined for that attribute in the Alarms table on the Alarms page for a specific service or component (for example, LDR > Storage > **Alarms > Main**). 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. Custom alarms with alarm severities from Critical down to Notice.
- 2. Global Custom alarms with alarm severities from Critical down to Notice.
- 3. 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 Custom alarm for an attribute, the NMS service only evaluates the attribute value against Custom alarms. Global Custom alarms and Default alarms are not evaluated. Thus, an enabled Global Custom alarm for an attribute can meet the criteria needed to trigger an alarm, but it will not be triggered because a Custom alarm (that does not meet the specified criteria) for the same attribute is enabled. No alarm is triggered and no notification is sent.

### Related concepts

What an Admin Node is on page 274

#### Alarm triggering examples

You can use these example to understand how Custom alarms, Global Custom alarms, and Default alarms are triggered.

### Example 1

For the following example, an attribute has a Global Custom alarm and a Default alarm defined and enabled as shown in the following table.

	Thresho	Threshold Values	
	Global Custom alarm (enabled)	Default alarm (enabled)	
Notice	>= 1500	>= 1000	

	Threshold Values	
	Global Custom alarm (enabled)	Default alarm (enabled)
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.

Note, however, that if there are Custom alarms for an attribute, these alarms are still evaluated as Custom alarms have a higher priority than Global Custom alarms.

#### Example 2

For the following example an attribute has a Custom alarm, a Global Custom alarm, and a Default alarm defined and enabled as shown in the following table.

	Threshold Values		
	Custom alarm (enabled)	Global Custom alarm (enabled)	Default alarm (enabled)
Notice	>= 500	>= 1500	>=1000
Minor	>= 750	>= 15,000	>=10,000
Major	>=1,000	>= 150,000	>= 250,000

If the attribute is evaluated when its value is 1000, a Major alarm is triggered and an email notification is sent to the designated personnel. The Custom alarm takes precedence over both the Global Custom alarm and Default alarm. A value of 1000 reaches the threshold value of the Major severity level for the Custom alarm. As a result, the attribute value triggers a Major level alarm.

Within the same scenario, if the Custom alarm is then disabled and the attribute value reevaluated at 1000, the alarm level is changed to Normal. The attribute value is evaluated against the threshold values of the Global Custom alarm, the next alarm class that is defined and enabled. A value of 1000 does not reach any threshold level for this alarm class. As a result, the attribute value is evaluated to be Normal and no notification is sent. The Notice level alarm from the previous evaluation is cleared.

### Example 3

For the following example, an attribute has a Custom alarm, Global Custom alarm, and Default alarm defined and enabled/disabled as shown below in the following table.

	Threshold Values		
	Custom alarm (disabled)	Global Custom alarm (enabled)	Default alarm (enabled)
Notice	>= 500	>= 1500	>=1000
Minor	>= 750	>= 15,000	>=10,000

	Threshold Values		
	Custom alarm (disabled)	Global Custom alarm (enabled)	Default alarm (enabled)
Major	>=1,000	>= 150,000	>= 250,000

If the attribute is evaluated when its value is 10,000, a Notice alarm is triggered and an email notification is sent to the designated personnel.

The Custom alarm is defined, but disabled; therefore, the attribute value is evaluated against the next alarm class. The Global Custom alarm is defined, enabled, and it takes precedence over the Default alarm. The attribute value is evaluated against the threshold values set for the Global Custom alarm class. A value of 10,000 reaches the Notice severity level for this alarm class. As a result, the attribute value triggers a Notice level alarm.

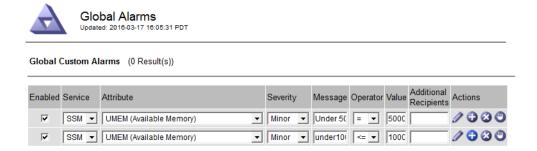
If the Global Custom alarm is then disabled and the attribute value reevaluated at 10,000, a Minor level alarm is triggered. The attribute value is evaluated against the threshold values for the Default alarm class, the only alarm class in that is both defined and enabled.

A value of 10,000 reaches the threshold value for a Minor level alarm. As a result, the Notice level alarm from the previous evaluation is cleared and the alarm level changes to Minor. An email notification is sent to the designated personnel.

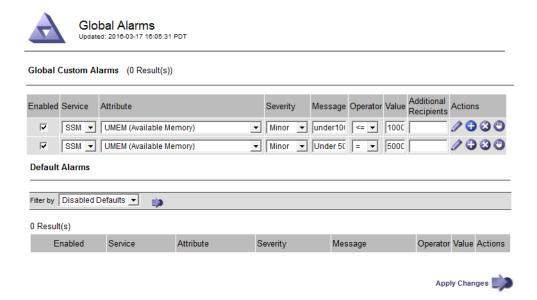
### Alarms of same severity

If two Global Custom or 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).



If the order is reversed, when UMEM drops to 100MB, the first alarm (<=100000000) is triggered, but not the one below it (= 50000000).



#### Alarm class overrides

To override a class of alarms, disable all alarms within that class. If all alarms within a class for an attribute are disabled, the NMS service interprets the class as having no alarms configured for the attribute and evaluates the next lower class for enabled alarms.

For example, if an alarm is triggered at the Global Custom alarm class level, it means that there are no enabled alarms at the Custom alarms class level for that attribute.

For example, to override a Default alarm, add a Global Custom alarm or Custom alarm for that attribute. This override is achieved because the NMS service does not evaluate lower priority alarm classes once an alarm setting is detected within a class. If this override is performed after an alarm has already been triggered, the override will not take effect until the alarm is triggered again.

### Severity changes

If an alarm's severity changes, the severity is propagated up the network hierarchy as needed. If there is a notification configured, a notification is sent. The notification is sent only at the time the alarm enters or leaves the new severity level.

### **Notifications**

A notification reports the occurrence of an alarm or the change of state for a service. It is an email communication to designated personnel that the system requires attention.

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.

## Creating custom service or component alarms

Customizing alarm settings enables you to create a customized methodology for monitoring the StorageGRID system. You can create alarms on individual services or components in addition to creating global alarms.

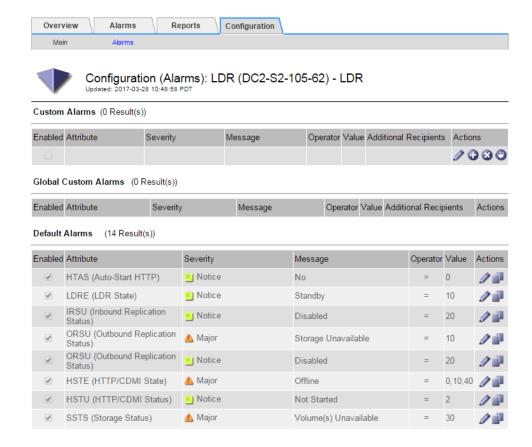
#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

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

- 1. Select Support > Grid Topology.
- Select a service or component in the Grid Topology tree.
- 3. Click Configuration > Alarms.



- **4.** Add a new row to the Custom alarms table:
  - Click **Edit** (if this is the first entry) or **Insert** to add a new alarm.

- Copy an alarm from the Default alarms or Global Custom alarms tables. Click **Copy** next to the alarm you want to modify.
- **5.** Make any necessary changes to the Custom alarm settings:

Heading	Description
Enabled	Select or clear to enable or disable the alarm.
Attribute	Select the name and code of the attribute being monitored from the list of all attributes applicable to the selected service or component.
	To display information about the attribute, click <b>Info</b> next to the attribute's name.
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	Operators for testing the current attribute value against the Value threshold:
	• = equal to
	• > greater than
	• < less than
	• >= greater than or equal to
	• <= less than or equal to
	• ≠ not equal to
Value	The alarm's threshold value used to test against the attribute's actual value using the operator.
	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/or ranges.
Additional Recipients	A supplementary list of email addresses to be notified when the alarm is triggered, in addition to the mailing list's configuration on the <b>Configuration</b> > <b>Notifications</b> page. Lists are comma delineated.
	<b>Note:</b> Mailing lists require SMTP server setup in order to operate. Before adding mailing lists, confirm that SMTP is configured.
	Notifications for Custom alarms can override notifications from Global Custom or Default alarms.
Actions	Control buttons to:
	Z Edit a row
	Insert a row
	Delete a row
	Drag-and-drop a row up or down
	Copy a row

## 6. Click Apply Changes.

## Creating Global Custom alarms

You can configure Global Custom alarms when you require a unique alarm that is the same for every service of the same type. Customizing alarm settings enables you to create a customized methodology for monitoring the StorageGRID system.

### Before you begin

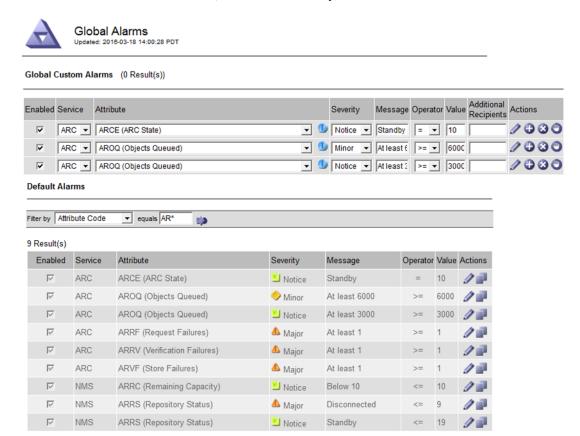
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

Global 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.

#### Steps

- 1. Select Configuration > 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**





To modify a Default alarm, search for the Default alarm.

- a. Under Filter by, select either Attribute Code or Attribute Name.
- 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**.
- **d.** 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 clear to enable or disable the alarm.
Attribute	Select the name and code of the attribute being monitored from the list of all attributes applicable to the selected service or component.
	To display information about the attribute, click <b>Info</b> next to the attribute's name.
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	Operators for testing the current attribute value against the Value threshold:
	• equals
	• > greater than
	• < less than
	• >= greater than or equal to
	• <= less than or equal to
	• ≠ not equal to
Value	The alarm's threshold value used to test against the attribute's actual value using the operator.
	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/or ranges.
Additional Recipients	A supplementary list of email addresses to be notified when the alarm is triggered. This is in addition to the mailing list's configuration on the <b>Configuration &gt; Notifications &gt; Main</b> page. Lists are comma delineated.
	<b>Note:</b> Mailing lists require SMTP server setup in order to operate. Before adding mailing lists, confirm that SMTP is configured.
	Notifications for Custom alarms can override notifications from Global Custom or Default alarms.

Heading	Description
Actions	Control buttons to:
	Edit a row
	Insert a row
	Opelete a row
	Drag-and-drop a row up or down
	Copy a row

4. Click Apply Changes.

# **Disabling alarms**

Alarms are enabled by default, but you can disable alarms that are not required.

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.

**Attention:** There are consequences to disabling alarms and extreme care should be taken. Disabling an alarm can result in no alarm being triggered. Because alarms are evaluated by alarm class and then severity level within the class, disabling an alarm at a higher class does not necessarily result in a lower class alarm being evaluated. All alarms for a specific attribute must be disabled before a lower alarm class will be evaluated.

#### Related tasks

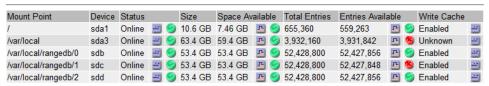
Clearing triggered alarms on page 135

### Alarms and tables

Alarm attributes displayed in tables can be disabled at the service, component, or system level. Alarms cannot be disabled for individual rows in a table.

For example, in the following figure, there are two critical Entries Available (VMFI) alarms. 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



## Disabling Default alarms for services

To temporarily stop alarms for a specific service, you can disable Default alarms for that service.

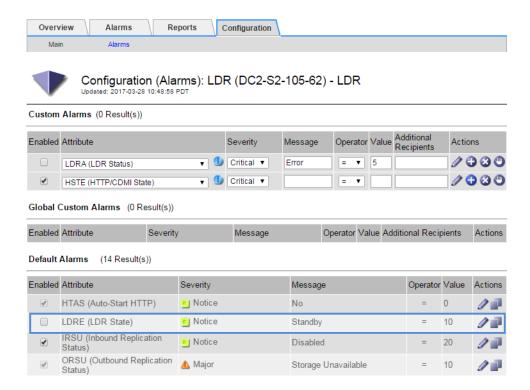
### Before you begin

• You must be signed in to the Grid Manager using a supported browser.

• You have specific access permissions.

### **Steps**

- 1. Select Support > Grid Topology.
- 2. Select a service or component in the Grid Topology tree.
- 3. Click Configuration > Alarms.
- 4. In the **Default Alarms** table, click **Edit** next to the alarm you want to disable.
- 5. Clear the **Enabled** check box for the alarm.



6. Click Apply Changes.

The Default alarm is disabled for the service or component.

## Disabling a Default alarm system wide

You can temporarily disable a Default alarm system wide.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- · You have specific access permissions.

- 1. Click Configuration > Global Alarms.
- 2. Search for the Default alarm to disable.
  - a. In the Default Alarms section, select Filter by > Attribute Code or Attribute Name.

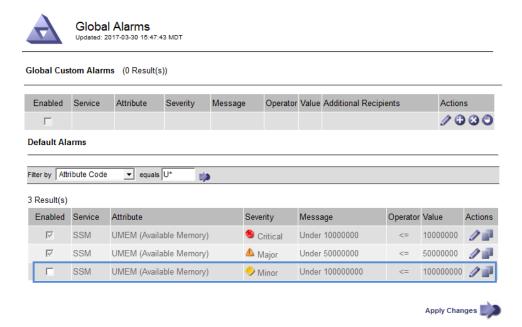
b. Type a search string.

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**.

Note: Selecting Disabled Defaults displays a list of all currently disabled Default alarms.

- 3. In the Default Alarms table, click the Edit icon next to the alarm you want to disable.
- Clear the **Enabled** check box.



5. Click Apply Changes.

The Default alarm is disabled system wide.

## Disabling Global Custom alarms for services

To disable a global alarm for a service, create another enabled global alarm for the attribute. You must create another enabled global alarm, because if all alarms within a class for an attribute are disabled, the NMS service interprets the class as having no alarms configured for the attribute and evaluates the next lower class for the enabled alarm.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### About this task

Instead of creating a Global Custom alarm and disabling it for selected services, reconfigure the alarms such that you create individual local Custom alarms for the services that require the alarm. If you want to ensure that all these Custom alarms have the same configuration, you can create a Global Custom alarm, disable it, and then enable it for selected services as a Custom alarm.

If you want to create a Global Custom alarm and disable it for selected services, you must create a local Custom alarm for that service that will never be triggered. A local Custom alarm that is never triggered overrides all Global Custom alarms for that service.

Note: Alarms cannot be disabled for individual rows in a table.

#### **Steps**

- 1. Select Configuration > Global Alarms
- 2. In the Global Custom alarm table, click next to the alarm you want to disable.

  The alarm is copied to the Custom Alarms table.
- 3. Clear the **Enabled** check box for the alarm.
- 4. Click Apply Changes.

#### Related tasks

Creating custom service or component alarms on page 127

## Disabling Global Custom alarms system wide

You can disable a Global Custom alarm for the entire system.

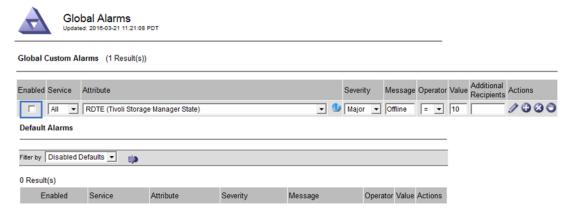
### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

Note: Alarms cannot be disabled for individual rows in a table.

- 1. Select Configuration > Global Alarms.
- 2. In the Global Custom Alarms table, click **Edit** next to the alarm you want to disable.
- 3. Clear the **Enabled** check box.





### 4. Click Apply Changes.

The Global Custom alarm is disabled system wide.

## Clearing triggered alarms

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).

### Before you begin

You must have the root/admin password as listed in the Passwords.txt file.

## **Steps**

- **1.** Disable the alarm.
- **2.** From the service laptop, 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 concepts**

Disabling alarms on page 131

# What AutoSupport is

AutoSupport enables technical support to proactively monitor the health of your StorageGRID system.

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

- Weekly: Automatically send AutoSupport messages once per week (default setting: Enabled)
- User-triggered: Manually send AutoSupport messages at any time
- **Event-triggered**: Automatically send AutoSupport messages when significant system events occur (default setting: Enabled)



You can choose from three protocols for sending AutoSupport messages: HTTPS, HTTP, and SMTP.

By analyzing AutoSupport information, technical support can help you determine the health and status of your StorageGRID system and troubleshoot any problems that might occur. Technical support can also monitor the storage needs of the system and help you determine if you need to add new nodes or sites.

### Related tasks

Specifying the protocol for AutoSupport messages on page 136

### **Related information**

NetApp Support

## Specifying the protocol for AutoSupport messages

You can select from three protocols for sending AutoSupport messages.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- If you will choose the HTTPS or HTTP protocol for sending AutoSupport messages, you must have provided internet access to the primary Admin Node.
- If you will choose the SMTP protocol for sending AutoSupport messages, you must have configured an email server.

#### About this task

The following protocol choices are available:

- HTTPS: This is the default setting for new installations.
- HTTP: Based on security issues, this protocol is not recommended.
- SMTP: This is the default setting for upgrades from StorageGRID 11.1 to 11.2.

**Note:** The HTTPS protocol uses port 443, and the HTTP protocol uses port 80.

### Steps

1. Select Support > AutoSupport.

The Weekly AutoSupport page appears.







Send

- 2. In the Protocol drop-down, choose the desired protocol.
- 3. Click Apply Changes.

All Weekly, User-Triggered, and Event-Triggered messages are sent using the selected protocol.

#### Related tasks

Configuring email server settings on page 114

## Sending user-triggered AutoSupport messages

You can manually trigger an AutoSupport message at any time.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### **Steps**

- 1. Select Support > AutoSupport.
- From the AutoSupport menu, select **User-triggered**.
- Click Send.



The StorageGRID system attempts to send an AutoSupport message to technical support. If the attempt is successful, the Last Attempt attribute updates to Successful. If there is a problem, the Last Attempt attribute updates to Failed. The StorageGRID system does not try again.

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."

The AutoSupport message includes the following information:

- StorageGRID software version
- · Operating system version
- · System-level and location-level attribute information
- · All alarms raised in the last seven days
- · 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

## Disabling weekly AutoSupport messages

By default, the StorageGRID system is configured to send an AutoSupport message to NetApp Support once a week.

### Before you begin

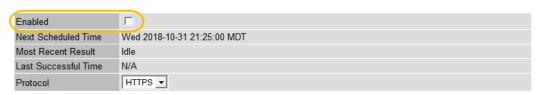
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

### About this task

To determine when the weekly AutoSupport message is sent, see the **Next Scheduled Time** attribute on the **AutoSupport > Weekly** page. You can disable the automatic sending of an AutoSupport message at any time.

- 1. Select Support > AutoSupport.
- 2. From the AutoSupport menu, select Weekly.
- 3. Clear the **Enabled** check box.







## Disabling event-triggered AutoSupport messages

By default, the StorageGRID system is configured to send an AutoSupport message to NetApp Support when a significant system event occurs.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

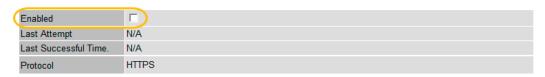
#### About this task

You can disable the automatic sending of an AutoSupport message at any time.

### **Steps**

- 1. Select Support > AutoSupport.
- **2.** From the AutoSupport menu, select **Event-triggered**.
- 3. Clear the **Enabled** check box.







4. Click Apply Changes.

## **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.

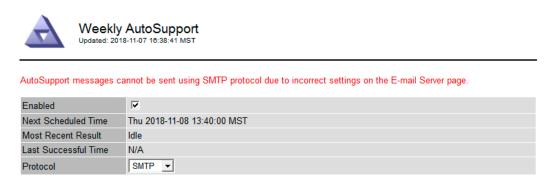
**Attention:** Suppressing email notifications system wide (**Configuration > Display Options**, Notifications Suppress All check box) suppresses event-triggered AutoSupport emails, even when the Enabled check box is selected on the Event-Triggered AutoSupport page.

### 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, an email server configuration error could appear.
- 2. Does not attempt to send the message again.
- 3. Logs the error in nms.log.

### 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."

Check that the StorageGRID system's email server is correctly configured and that your email server is running.

#### Related tasks

Configuring email server settings on page 114
Suppressing email notifications system wide on page 119

# **Using reports**

You can use reports to monitor the state of the StorageGRID system and troubleshoot problems. The types of reports available in the Grid Manager include pie charts (on the Dashboard only), graphs, and text reports.

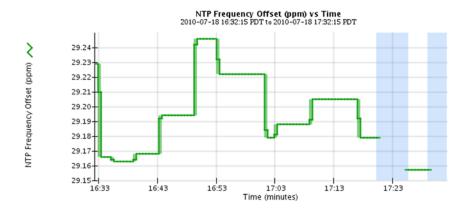
## Types of charts

In addition to the summary pie charts shown on the Dashboard, you can access more detailed graphs that present the data with the attribute value (vertical axis) over a specified time span (horizontal axis).

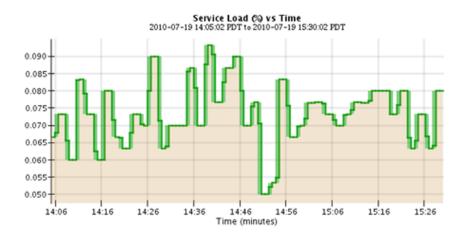
The Dashboard provides access to pie charts summarizing available storage as well as graphs for ILM and protocol operations.

In addition, graphs are available from the Nodes page and the Grid Topology tree. There are three types of graphs:

Line graph: Used to plot the values of an attribute that has 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 graph: Used to plot volumetric 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.



• State graph: State graphs are used to plot 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.

Contine

2004-07-09 16:40:23 to 2004-07-09 17:17:11

Online

Standby

Offline
16:41 16:45 16:49 16:53 16:57 17:01 17:05 17:09 17:13 17:17

Time - Minutes

# **Chart legend**

The lines and colors used to draw charts have specific meaning.

Sample	Meaning	
	Reported attribute values are plotted using dark green lines.	
\$1pt	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.	
THIT	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**

The Nodes page contains the charts you should access regularly to monitor attributes such as storage capacity and throughput. In some cases, especially when working with technical support, you might use the Grid Topology tree to access additional charts.

#### Before you begin

You must be signed in to the Grid Manager using a supported browser.

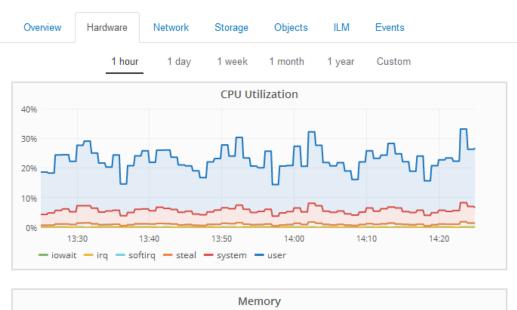
### About this task

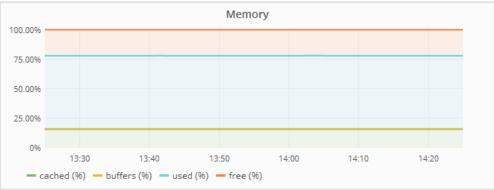
Note: You cannot create charts for all attributes; for example, text attributes such as Node ID, version number, and build number.

### **Steps**

- 1. Select Nodes > grid or grid node.
- Select the appropriate tab, or select **Chart** to the right of an attribute to display a chart. For some attributes, the chart appears when you select the tab. For other attributes, you select the **Chart** icon to display the chart.

## DC1-S1-226 (Storage Node)





- 3. To display additional attributes and charts, select **Support > Grid Topology**.
- 4. Select grid node > component or service > Overview > Main.



#### **Computational Resources**

Service Restarts:	1	
Service Runtime:	6 days	
Service Uptime:	6 days	
Service CPU Seconds:	10666 s	
Service Load:	0.266 %	<u>r</u>

#### Memory

Installed Memory: 8.38 GB	with the same of t		
Available Memory: 2.9 GB	nstalled Memory:	8.38 GB	- <del>8</del>
	Available Memory:	2.9 GB	□ •

#### **Processors**

Processor Number	Vendor	Туре	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

5. Click **Chart** 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.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions. For details, see information about controlling system
  access with administration user accounts and groups.

- 1. Select Support > 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** checkbox, or to round values to a maximum of three decimal places (for example, for attributes reported as percentages), deselect the **Raw** Data checkbox.
- **6.** Select the time period to report on from the **Quick Query** drop-down list.
  - Select the Custom Query option to select a specific time range.
  - The chart appears after a few moments. Allow several minutes for tabulation of long time ranges.
- 7. If you selected Custom Query, customize the time period for the chart by entering the **Start Date** and End Date.

Use the format YYYY/MM/DD HH: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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions. For details, see information about controlling system access with administration user accounts and groups.

#### 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.

Grey 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 > 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), 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 report appears after a few moments. Allow several minutes for tabulation of long time ranges.
- 7. If you selected Custom Query, you need to customize the time period to report on by entering the Start Date and End Date.

Use the format YYYY/MM/DD HH:MM:SS in local time. Leading zeros are required to match the format. For example, 2017/4/6 7:30:00 fails validation. The correct format is: 2017/04/06 07:30:00.

#### 8. Click Update.

A text report is generated after a few moments. Allow several minutes for tabulation of long time ranges. Depending on the length of time set for the query, either a raw text report or aggregate text report is displayed.

9. If you want to print the report, right-click and select Print, and modify any necessary printer settings and click Print.

## **Exporting text reports**

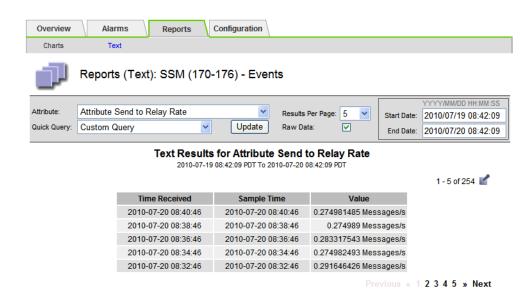
Exported text reports open a new browser tab, which enables you to select and copy the data.

#### 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**

- 1. Select Support > Grid Topology.
- **2.** Create a text report.
- 3. Click Export



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)
Ouerv 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.

# Managing objects through information lifecycle management

You manage objects by configuring information lifecycle management (ILM) rules and policies, which determine how the StorageGRID system creates and distributes copies of object data and how it manages those copies over time.

Designing and implementing an ILM policy requires careful planning and an understanding of ILM. You must define the logic for how objects are to be filtered, copied, and distributed throughout the system, taking into account the topology of your StorageGRID system, object protection requirements, and available storage types.

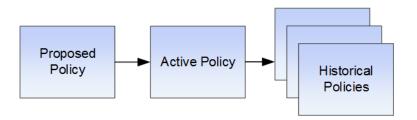
# What an information lifecycle management policy is

An information lifecycle management (ILM) policy is a set of prioritized ILM rules that determines how the StorageGRID system manages object data over time.

## 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. Activating the proposed policy causes the previously active policy to become a historical policy. Historical ILM policies are those ILM policies that are no longer active. Historical ILM policies cannot be deleted.

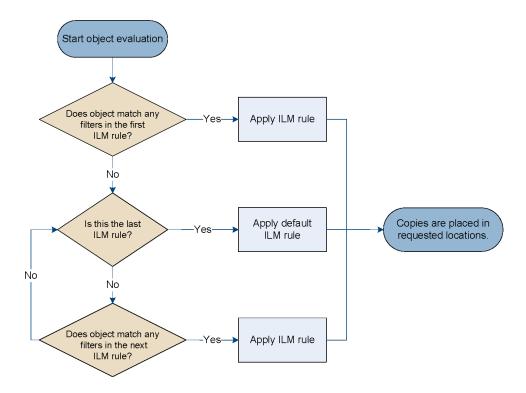


#### How ILM policies evaluate objects

When an object is ingested into the StorageGRID system, its metadata is evaluated against the first ILM rule in the active ILM policy. If the object metadata matches the filters in the first rule, the content placement instructions for that rule distribute object data to the specified storage locations. If the metadata does not match the filters in the first rule, the object is evaluated against each subsequent ILM rule in the active ILM policy, until a match is made.

One ILM rule must be set as the default ILM rule for the policy. If none of the other ILM rules in the policy matches the object metadata, the placement instructions specified by the default rule are applied. When the StorageGRID system is first installed, the stock ILM rule, Make 2 Copies, is the default ILM rule in the active ILM policy.

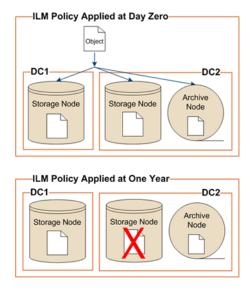
This figure illustrates how each object is evaluated by the ILM rules in the active ILM policy.



## **ILM** policy example

The following figure illustrates an ILM policy that includes ILM rules that specify the following:

- 1. When an object is ingested, one copy is placed on disk (Storage Node) at Data Center 1 (DC1), one copy is placed on disk at Data Center 2 (DC2), and one copy is placed on archival media (Archive Node) at DC2.
- 2. At the end of one year, the copy on disk at DC2 is deleted.



#### What dual commit is

Dual commit is the system's default functionality when an object is ingested. Dual commit is designed to prevent the loss of object data if an object's initial storage location fails before the object can be evaluated against the active ILM policy.

As soon as an object is ingested, a second copy of that object is created and distributed to a different Storage Node within the same site. When the object is matched by an ILM rule in the active policy, StorageGRID determines if the initial, dual-commit copies satisfy the placement instructions in the rule. If they do not, the object is added to the ILM evaluation queue. When the object is re-evaluated, new object copies might need to be made in different locations, and the initial dual-commit copies might need to be deleted.

If the request to create the dual-commit copies fails (for example, a network issue prevents the second initial copy from being made), the StorageGRID system does not retry and ingest fails.

Dual commit is enabled by default. If ILM rules are configured to store only one instance of replicated object data, you can specify a single-commit ingest operation to avoid unnecessarily creating and then deleting copies generated by the dual-commit ingest. For configuration information, see the instructions for implementing S3 or Swift client applications.

**Note:** You cannot specify a single-commit ingest operation (that is, you cannot use REDUCED\_REDUNDANCY) for S3 buckets that have compliance enabled. This is to ensure that compliance requirements are satisfied (two copies of each object exist) before the object is evaluated against the active ILM policy.

#### Related tasks

Managing S3 buckets and objects for compliance on page 220

#### **Related information**

Implementing S3 client applications
Implementing Swift client applications

# What an information lifecycle management rule is

An information lifecycle management (ILM) rule determines how the StorageGRID system stores object data over time. You configure ILM rules and then add them to an ILM policy.

#### ILM rules determine:

- Where an object's data is stored and the type of storage used (storage grades and storage pools)
- The number and type of copies made (replicated or erasure coded)
- How specific objects are managed (object filtering)
- How the object's data is managed over time, where it is stored, and how it is protected from loss (placement instructions)

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 load balanced across all Storage Nodes.

## What replication and erasure coding are

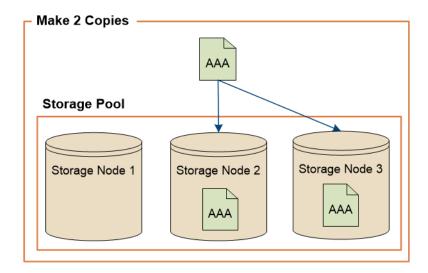
StorageGRID provides two methods for protecting object data from loss: replication and erasure coding.

## What replication is

Replication is one of two methods used by StorageGRID to store object data. When StorageGRID matches objects to an ILM rule that is configured to create replicated copies, 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.

**Attention:** In general, do not configure ILM rules to create only one replicated copy. If the only replicated copy is lost or corrupt, data will be lost. In addition, you might lose access to the object during maintenance or recovery operations. Use ILM to create a single replicated copy only if you are willing to risk the loss of object data when a failure occurs.

Note: 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.

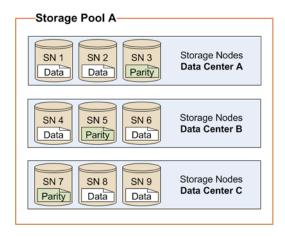
#### Related concepts

Using multiple storage pools for cross-site replication on page 161

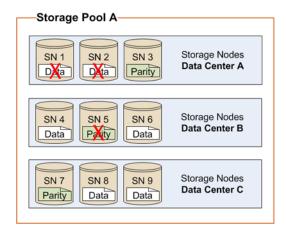
## What erasure coding is

Erasure 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.

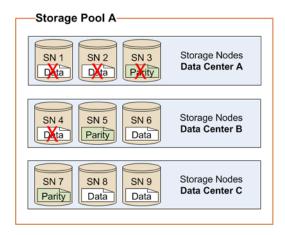
The following example illustrates the use of an erasure coding algorithm on an object's data. In this example, the ILM rule uses a 6+3 erasure coding scheme. Each object is sliced into six equal data fragments, and three parity fragments are computed from the object data. Each of the nine fragments is stored on a different node across multiple sites to provide data protection for node failures or site loss.



The 6+3 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 six of the nine fragments (data or parity) remain available. Up to three 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 three Storage Nodes are lost, the object is not retrievable.



#### Related concepts

What erasure coding schemes are on page 155

#### Related tasks

Configuring an Erasure Coding profile on page 177

## 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, confirm that the storage pool includes exactly one site or three or more sites. Do not use the default storage pool, All Storage Nodes, or a storage pool that includes the default site, All Sites.

Note: 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

The following table lists the erasure coding schemes currently supported by StorageGRID for storage pools that include three or more sites. It specifies the recommended number of sites and Storage Nodes for each scheme. The supported erasure coding schemes are designed to provide site loss protection. One site can be lost, and the object will still be accessible.

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?
4+2	3	3*	9	Yes
6+2	4	3*	12	Yes
8+2	5	3*	15	Yes
6+3	3	4	12	Yes

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?
9+3	4	4	16	Yes
2+1	3	3*	9	Yes
4+1	5	3*	15	Yes
6+1	7	3*	21	Yes

<sup>\*</sup> At minimum, each site requires three Storage Nodes.

Additional erasure coding schemes are available. Contact your account manager.

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.

**Note:** If you are unsure of which scheme to use, select 4+2 or 6+3, or contact support. In general, you should avoid using the m+1 schemes unless your application does not require a high degree of fault tolerance.

#### Erasure coding schemes for one-site storage pools

Erasure coding is well suited for single-site deployments that require efficient data protection with only a single erasure-coded copy rather than multiple replicated copies. A storage pool that includes only one site supports all of the erasure coding schemes listed in the previous table, assuming that the site includes an adequate number of Storage Nodes. For example, the 2+1 erasure coding scheme requires a storage pool with three or more Storage Nodes, while the 6+3 scheme requires a storage pool with at least nine Storage Nodes.

#### 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.

**Note:** 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:

- Increased number of Storage Nodes and sites 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 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, higher WAN network traffic usage for retrievals and repairs, especially for frequently retrieved objects or object repairs over WAN network connections.
- · Higher usage of compute resources.

## Requirements for erasure coding

Erasure coding is best suited for the following requirements:

• Objects larger than 1 MB in size.

Attention: Due to the overhead of managing the number of fragments associated with an erasure-coded copy, do not use erasure coding for objects smaller than 200 KB.

- 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.

# Configuring ILM rules

Before you can create and activate the ILM policy for your StorageGRID system, you must define object storage locations, determine the types of copies you want, optionally configure S3 regions, and create one or more rules.

## About this task

When configuring 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.

#### **Steps**

- 1. Configuring storage pools on page 158
- 2. Using Cloud Storage Pools on page 165
- 3. Configuring an Erasure Coding profile on page 177
- 4. Configuring regions (optional and S3 only) on page 179
- 5. Creating an ILM rule on page 181

## Configuring storage pools

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.

#### **Steps**

- 1. Creating and assigning storage grades on page 158
- 2. Guidelines for creating storage pools on page 160
- 3. Using multiple storage pools for cross-site replication on page 161
- 4. Creating a storage pool on page 163
- **5.** Editing a storage pool on page 164
- **6.** Removing a storage pool on page 165

## Creating and assigning storage grades

A storage grade identifies the type of storage used by a Storage Node to store object data, for example, flash or spinning disk. If you use more than one type of storage, you can optionally create a storage grade to identify each type. Creating storage grades makes it easier to select a specific type of Storage Node when configuring storage pools.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

If storage grade is not a concern (for example, your StorageGRID system only uses spinning disks), you can skip this procedure and use the All Storage Nodes storage grade when configuring storage pools.

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

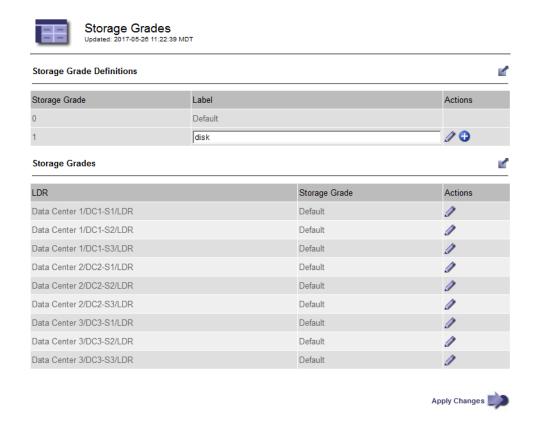
nodes. Storage grades assigned to only one node can cause ILM backlogs if that node becomes unavailable.

**Note:** You cannot configure storage grades for Archive Nodes.

#### **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 LDR services added during a StorageGRID system expansion.



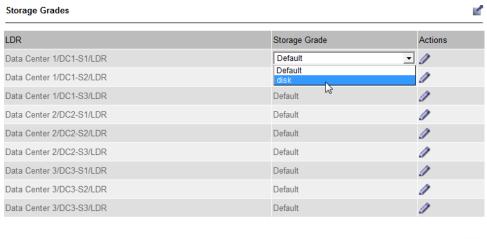
b. To edit an existing storage grade, click **Edit** and modify the label as required.

Note: You cannot delete storage grades.

c. Click Apply Changes.

These storage grades are now available for assignment to LDR services.

- 3. Assign a storage grade to an LDR service:
  - a. For each Storage Node's LDR service, click **Edit** and select a storage grade from the list.





**Attention:** 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 once the ILM policy is activated. If the assignment is changed, data is stored based on the new storage grade.

b. Click Apply Changes.

## 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 called All Storage Nodes that uses the default site,
  All Sites, and the All Storage Nodes storage grade. Because this storage pool is automatically
  updated whenever you add new data center sites, review the guidelines for replicated and erasurecoded copies before using this storage pool or the default site, All Sites.
- 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 that 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 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 Compliance setting is enabled and you are creating a compliant ILM rule, you
  cannot use a storage pool that includes Archive Nodes. See "Managing S3 buckets and objects for
  compliance."

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 "Configuring connection settings for S3 API" for information.

## Guidelines for storage pools used for replicated copies

If your StorageGRID system includes more than one data center site, considering creating a storage pool for each site instead of using the default storage pool, All Storage Nodes, or a storage pool that includes the default site, All Sites. You can then specify those storage pools in the rule's placement instructions to enable site-loss protection of replicated copies.

## 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. No erasure coding schemes are available for a storage pool that has two sites.
- If you plan to create erasure-coded copies, create a new storage pool and manually add each site you want to include. Do not use the default storage pool, All Storage Nodes, or a storage pool that includes the default site, All Sites. This is to ensure that the erasure coding scheme does not become invalid when new sites are added. For example, if you currently have one site and use the All Storage Nodes storage pool, your Erasure Coding profile will become invalid if you add a second site.
- 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.

#### Related concepts

What replication is on page 153 What erasure coding is on page 154 What erasure coding schemes are on page 155 Using multiple storage pools for cross-site replication on page 161

## Related tasks

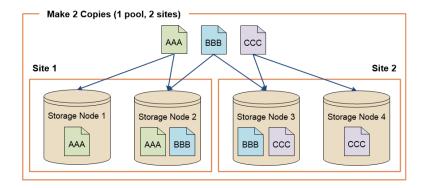
Managing S3 buckets and objects for compliance on page 220 Configuring connection settings for S3 API on page 260

#### Using multiple storage pools for cross-site replication

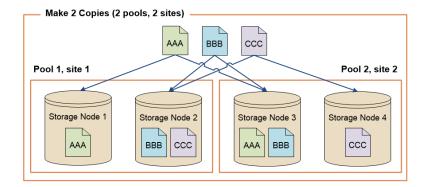
If your StorageGRID deployment includes more than one site, you can enable site-loss protection by creating a storage pool for each site and specifying both storage pools in the rule's placement instructions. For example, if you configure an ILM rule to make two replicated copies and specify storage 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.

## 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.

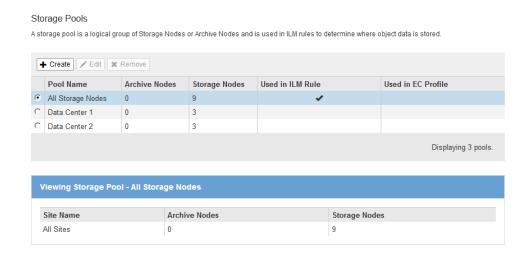
#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You have reviewed the guidelines for creating storage pools.

#### **Steps**

#### 1. Select ILM > Storage Pools.

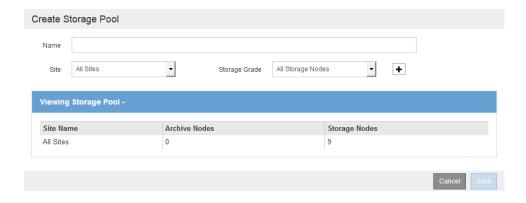
The Storage Pools page appears.



This page lists all defined storage pools, shows the number of Archive Nodes or Storage Nodes in each pool, and specifies whether the storage pool is currently being used in an ILM rule or EC profile. The table below the list shows details for the currently selected storage pool, including the sites in the storage pool and the number and type of nodes at each site.

#### 2. Click Create.

The Create Storage Pool page appears. By default, All Sites and the All Storage Nodes storage grade are selected.



**3.** Enter a name for the storage pool.

Use a name that will be easy to identify when you configure Erasure Coding profiles and ILM rules. When you enter the name, the table heading updates automatically.

**4.** From the **Site** drop-down list, select the site to which object data will be copied if an ILM rule or Erasure Coding profile uses this storage pool.

The default value, All Sites, includes all data center sites in your StorageGRID grid.

**Attention:** When creating a storage pool that will be used in an Erasure Coding profile, you must manually add each site you want to include. Do not include the default site, **All Sites**, in the 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 to which object data will be copied if an ILM rule uses this storage pool.

The **All Storage Nodes** default value includes all Storage Nodes at the selected site. The **Archive Nodes** default value includes all Archives 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.** Optionally, click + to add another entry to the storage pool.

For each entry, specify a unique combination of Site and Storage Grade.

**Note:** 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.

To remove an entry, click \*.

7. When you are satisfied with your selections, click **Save**.

The new storage pool is added to the list.

#### Related concepts

Guidelines for creating storage pools on page 160

## Editing a storage pool

You can edit a storage pool to change its name or to update sites and storage grades.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You have reviewed the guidelines for creating storage pools.
- If you plan to edit a storage pool that is used by a rule in the active ILM policy, you have considered how your changes will affect object data placement.

#### **Steps**

1. Select ILM > Storage Pools.

The Storage Pools page appears.

**2.** Select the radio button for the storage pool.

You cannot edit the All Storage Nodes storage pool.

- 3. Click Edit.
- **4.** As required, change the storage pool name or select other sites and storage grades.

Note: 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.

5. Click Save.

## Removing a storage pool

You can remove a storage pool that is not being used.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### **Steps**

1. Select ILM > Storage Pools.

The Storage Pools page appears.

2. Select the radio button for an unused storage pool.

You cannot remove the All Storage Nodes storage pool or any storage pool that is being used in a saved ILM rule or in an Erasure Coding profile.

- 3. Click Remove.
- 4. Click OK.

## **Using Cloud Storage Pools**

You can use a Cloud Storage Pool to move StorageGRID objects to an external storage location, such as Amazon Glacier. Moving objects outside of the grid lets you take advantage of a low-cost storage tier for long-term archive.

### **Steps**

- 1. What a Cloud Storage Pool is on page 165
- 2. Lifecycle of a Cloud Storage Pool object on page 167
- 3. Considerations for Cloud Storage Pools on page 168
- **4.** Creating a Cloud Storage Pool on page 170
- 5. Editing a Cloud Storage Pool on page 173
- 6. Removing a Cloud Storage Pool on page 174
- 7. Troubleshooting Cloud Storage Pools on page 175

## 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 low-cost Amazon Glacier

storage, or you might want to free up on-premise storage by storing older versions of objects externally.

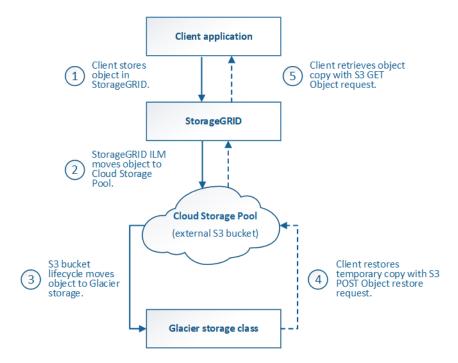
From an ILM perspective, a Cloud Storage Pool is similar to a storage pool. To store objects in either location, you select the pool when creating the placement instructions for an ILM rule. However, while storage pools consist of Storage Nodes or Archive Nodes within the StorageGRID system, a Cloud Storage Pool consists of an external S3 bucket.

The following table compares storage pools to Cloud Storage Pools and shows the high-level similarities and differences.

	Storage pool	Cloud Storage Pool
How is it created?	Using the <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 ILM > Storage Pools option in Grid Manager. You must set up the external S3 bucket before you can create the Cloud Storage Pool.
How many pools can you create?	Multiple	Only one
Where are objects stored?	On one or more Storage Nodes or Archive Nodes within StorageGRID.	In a single S3 bucket that is external to the StorageGRID system.  Note: Optionally, you can configure a bucket lifecycle for the external bucket if you want to transition objects to low-cost, long-term storage, such as Amazon Glacier. The external storage system must support the Glacier storage class and the S3 POST Object restore API.
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	Replication
How many copies are allowed?	Multiple	One The replicated copy is moved to the Cloud Storage Pool and stored externally to StorageGRID. Additional replicated or erasure coded copies of the object cannot exist in StorageGRID during the same time period.
What are the advantages?	You can quickly access the object at any time.  Redundant copies are available if a Storage Node or site fails.	Low-cost storage.  Note: Redundant copies are not available locally, and restore operations might take longer. For these reasons, a Cloud Storage Pool is an appropriate option for rarely accessed or older versions of data.

## Lifecycle of a Cloud Storage Pool object

The figure shows the lifecycle stages of an object that is stored in a Cloud Storage Pool.



#### 1. Object stored in StorageGRID

To start the lifecycle, a client application stores an object in StorageGRID.

## 2. Object moved to Cloud Storage Pool

- When the object is matched by an ILM rule that uses a 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 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 storage

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.

**Note:** 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.

Note: Objects ingested by Swift tenants do not support POST Object restore requests, and should not be transitioned to Glacier storage. The external S3 bucket should not have a lifecycle configured for Glacier storage.

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 temporary copy to the 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).

**Note:** When the expiration date of the temporary copy is reached, the copy is automatically removed from the Cloud Storage Pool. The original object remains in Glacier.

## 5. Object retrieved

Once an object has been restored to the Cloud Storage Pool, the client application can issue a GET Object request to retrieve the restored object.

## Moving objects back to StorageGRID

Cloud Storage Pool objects can be moved back to StorageGRID, as follows:

- If the objects have not been transitioned to Glacier storage, use ILM to move the objects back to a storage pool within the StorageGRID system.
- If the objects have been transitioned Glacier storage:
  - 1. Issue S3 POST Object restore requests to restore the objects back to the Cloud Storage Pool. Specify an extended length of time for **Days** to ensure that ILM has adequate time to process them—several weeks or more is recommended.

**Note:** If you need to restore multiple objects, it is more efficient to issue the S3 POST Object restore requests directly to the external S3 bucket.

- 2. Use ILM to move the restored objects back to a storage pool within the StorageGRID system.
- Once an object is moved back to StorageGRID, the copy in the Cloud Storage Pool is deleted.

## **Deleting Cloud Storage Pool objects**

You can delete Cloud Storage Pool objects, as follows:

- Use ILM to delete the objects at the end of a specified time period.
- Issue an S3 DELETE Object request.

#### **Related information**

Implementing S3 client applications

## **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.

## Information required to create a Cloud Storage Pool

Before you can create a Cloud Storage Pool, you must create the external S3 bucket that you will use for the Cloud Storage Pool. When you create the Cloud Storage Pool in StorageGRID, you must enter the following information about the bucket:

- The Uniform Resource Identifier (URI) used to access the S3 bucket
- The exact name of the S3 bucket
- If an access key is required to access the S3 bucket, the access key ID and the secret access key
- Optionally, a custom CA certificate to verify TLS connections to the S3 bucket

## 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 (\$3:\*); 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

## Considerations for the bucket lifecycle configuration

The movement of objects from StorageGRID to 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 Cloud Storage Pool to AWS Glacier (or to a storage solution that implements the Glacier storage class) is controlled by the external bucket's lifecycle configuration.

If you want to transition objects from the Cloud Storage Pool, you must create the appropriate 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.

For example, suppose you want all objects that are moved from StorageGRID to the Cloud Storage Pool to be transitioned to Glacier storage immediately. You would create a lifecycle configuration rule for 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 Glacier on the day they were created (that is, on the day they were moved from StorageGRID to the Cloud Storage Pool).

**Attention:** When configuring bucket lifecycle rules for the external bucket, 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.

## 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 proxy settings to allow messages to be sent to external endpoints, such as an endpoint on the internet.

## Considerations for object segmentation

The **Segmentation** storage option must be enabled in the Grid Manager. This option is enabled by default.

#### **Considerations for AWS S3 costs**

When StorageGRID connects to the external Cloud Storage Pool bucket, it issues various S3 requests to monitor connectivity and to ensure it can perform the required operations. These requests might include requests to put, post, copy, and list objects. While some additional AWS costs will be associated with these StorageGRID requests, the overall cost of using a Cloud Storage Pool will be only a small fraction of what you will pay to store the objects in S3.

#### Related tasks

Configuring proxy settings on page 288

## **Creating a Cloud Storage Pool**

When you create a Cloud Storage Pool, you specify the name and location of the external bucket that StorageGRID will use to store objects. A StorageGRID system can have only one Cloud Storage Pool.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You have reviewed the guidelines for configuring Cloud Storage Pools.
- The external bucket referenced by the Cloud Storage Pool has already been created.

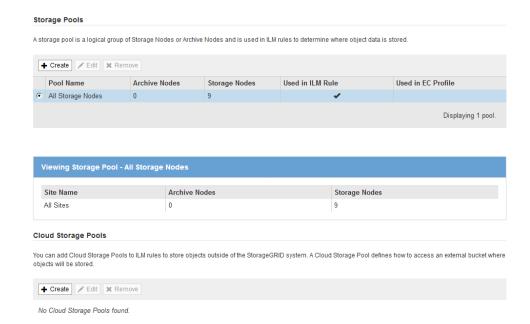
## About this task

A Cloud Storage Pool specifies a single external bucket. StorageGRID validates the Cloud Storage Pool as you create it, so you must ensure that the bucket specified in the Cloud Storage Pool exists and is reachable.

#### **Steps**

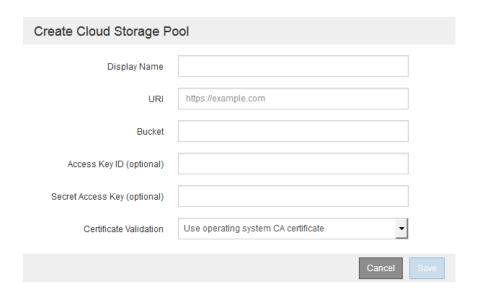
#### 1. Select ILM > Storage Pools.

The Storage Pools page appears. This page includes two sections: Storage Pools and Cloud Storage Pools.



2. In the Cloud Storage Pools section of the page, click Create.

The Create Cloud Storage Pool dialog appears.



**3.** 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.

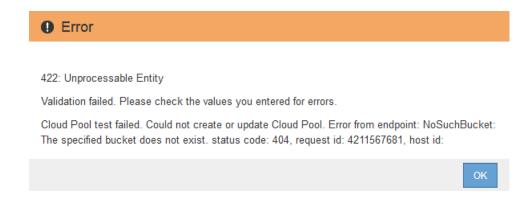
Field	Description
URI	The Uniform Resource Identifier (URI) used to access the S3 bucket used for the Cloud Storage Pool.
	Specify the URI in one of the following formats:
	• https://host:port
	• http://host:port
	If you do not specify a port, by default port 443 is used for HTTPS URIs and port 80 is used for HTTP URIs.
	For example, the URI for a Cloud Storage Pool bucket hosted on AWS might be https://s3-aws-region.amazonaws.com
Bucket	The name of the external S3 bucket that was created for the Cloud Storage Pool. The name you specify here must exactly match the S3 bucket's name or Cloud Storage Pool creation will fail. You cannot change this value after the Cloud Storage Pool is saved.
Access Key ID	Optionally, the Access Key ID for the account that owns the external bucket.
	For anonymous access to the bucket, omit both the Access Key ID and the Secret Access Key.
Secret Access	If you specified an Access Key ID, the associated Secret Access Key.
Key	A Secret Access Key is required when you specify an Access Key ID.
Certificate Validation	The method used to validate the certificate for TLS connections to the Cloud Storage Pool:
	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 security certificate.
	If you select this setting, copy and paste the custom security certificate in the <b>CA Certificate</b> text box.
	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:

- Validates that the bucket and the URI 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.



See the instructions for troubleshooting Cloud Storage Pools, resolve the issue, and then try saving the Cloud Storage Pool again.

## **Related concepts**

Considerations for Cloud Storage Pools on page 168 Troubleshooting Cloud Storage Pools on page 175

## **Editing a Cloud Storage Pool**

You can edit a Cloud Storage Pool to change its name, URI, or other details; however, you cannot change the bucket for a Cloud Storage Pool.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You have reviewed the guidelines for configuring Cloud Storage Pools.

#### **Steps**

1. Select ILM > Storage Pools.

The Storage Pools page appears.

- 2. In the Cloud Storage Pools table, select the radio button for the Cloud Storage Pool.
- 3. Click Edit.

Edit Cloud Storage Pool	- Example Cloud Storage Pool
Display Name	Example Cloud Storage Pool
URI	https://10.96.104.168:18082
Bucket	cloud-bucket
Access Key ID (optional)	*******QLQD
Secret Access Key (optional)	*****
Certificate Validation	Use operating system CA certificate
	Cancel

4. As required, change the display name, URI, credentials, or certificate validation method.

Attention: You cannot change the bucket for a Cloud Storage Pool.

#### 5. Click Save.

When you save a Cloud Storage Pool, StorageGRID validates that the bucket and the URI exist, and can be reached using the credentials that you specified.

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.

See the instructions for troubleshooting Cloud Storage Pools, resolve the issue, and then try saving the Cloud Storage Pool again.

#### Related concepts

Considerations for Cloud Storage Pools on page 168 Troubleshooting Cloud Storage Pools on page 175

#### 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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You have confirmed that the bucket 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."

**Note:** When you create a Cloud Storage Pool, StorageGRID writes a marker file to the bucket to identify the bucket 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.

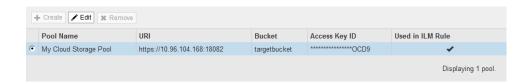
#### **Steps**

#### 1. Select ILM > Storage Pools.

The Storage Pools page appears.

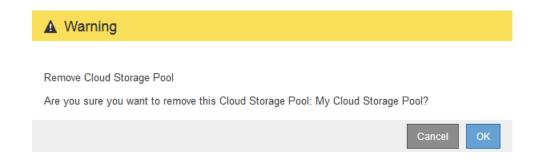
2. Select the radio button for a Cloud Storage Pool that is not currently used in an ILM rule.

You cannot remove a Cloud Storage Pool if it is used in an ILM rule. The Remove button is disabled.



#### 3. Click Remove.

A confirmation warning is displayed.



#### 4. Click OK.

The Cloud Storage Pool is removed.

## Related concepts

Troubleshooting Cloud Storage Pools on page 175

## **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.

## **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 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 bucket.

Try these steps to correct the issue:

- Check to make sure that no one in your organization is also using this bucket as a Cloud Storage Pool.
- Delete the x-ntap-sgws-cloud-pool-uuid file from the bucket 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 bucket.

To correct the issue, review the error message from the endpoint. Then, try one or more of the following:

- Create an external bucket with the same name you entered for the Cloud Storage Pool, and try to save the new Cloud Storage Pool again.
- Correct the 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-pooluuid 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.

## Related concepts

Lifecycle of a Cloud Storage Pool object on page 167

## **Configuring an Erasure Coding profile**

deleted object will not be found.

You create an Erasure Coding profile 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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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 two-site deployment.

#### About this task

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. 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.

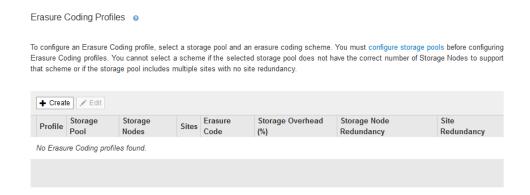
**Note:** You must select a storage pool that contains Storage Nodes. You cannot use Archive Nodes for erasure-coded data.

After you save the Erasure Coding profile, you can change its name, but you cannot select a different storage pool or erasure coding scheme. You also cannot delete the profile.

#### **Steps**

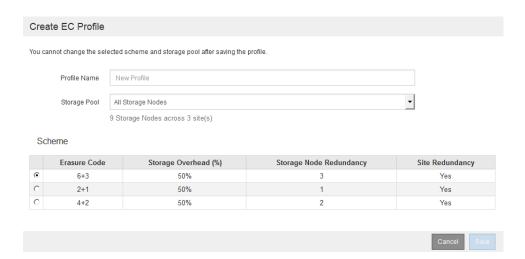
## 1. Select ILM > Erasure Coding.

The Erasure Coding Profiles page appears.



#### 2. Click Create.

The Create EC Profile dialog box appears. By default, the **Storage Pool** field shows the default storage pool, **All Storage Nodes**, and lists any available erasure coding schemes, based on the total number of Storage Nodes and sites available in your StorageGRID system.



**3.** Enter a name for the Erasure Coding profile.

This name is automatically appended to the storage pool name when you configure ILM rules for erasure coding. Enter a short, but meaningful name.

**4.** Select the storage pool you created for this Erasure Coding profile.

**Attention:** Do not use the default storage pool, **All Storage Nodes**, or a storage pool that uses the default site, **All Sites**.

When you select a storage pool, the list of available erasure coding schemes is updated to reflect the number of Storage Nodes and sites in that pool. The following information is listed for each available 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.



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 only 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.

Erasure Code	Storage Overhead (%)	Storage Node Redundancy	Site Redundancy
	3 ( )	,	•

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.

#### Related concepts

What erasure coding schemes are on page 155 Configuring storage pools on page 158 What erasure coding is on page 154

#### Related tasks

Creating an ILM rule on page 181

# 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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

When creating an S3 bucket, you can specify that the bucket be created in a specific region. Specifying a region allows the bucket to be geographically close to its users, which can help optimize latency, minimize costs, and address regulatory requirements.

When you create an ILM rule, you might want to use the region associated with an S3 bucket as an advanced filter. For example, you can design a rule that applies only to objects in S3 buckets created in the us-west-2 region. You can then specify that copies of those objects be placed on Storage Nodes at a data center site within that region to optimize latency.

When configuring regions, follow these guidelines:

- By default, all buckets are considered to belong to the us-east-1 region.
- You must create the regions using the Grid Manager before you can specify a non-default region when creating buckets using the Tenant Manager or Tenant Management API or with the LocationConstraint request element for S3 PUT Bucket API requests. An error occurs if a PUT Bucket request uses a region that has not been defined in StorageGRID.
- You must use the exact region name when you create the S3 bucket. Region names are case sensitive and must contain between 2 and 32 characters. Valid characters are numbers, letters, and hyphens.

**Note:** 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 us-east-1 (required)

#### 2. To add a region:

Region 2

a. Click the insert icon + to the right of the last entry.

us-west-1

- 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 \*.

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 ILM rule wizard.

## Related concepts

Using advanced filters in ILM rules on page 186

## 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 have specific access permissions.
- If you plan to use last access time metadata, Last Access Time updates must be enabled by bucket for S3 or by container for Swift.
- If you are creating erasure-coded copies, you have configured an Erasure Coding profile.

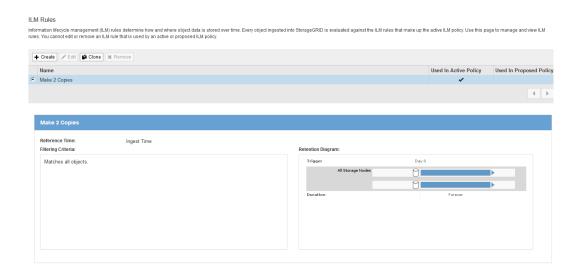
#### About this task

When StorageGRID evaluates objects against an ILM rule, it looks first at the rule's basic and advanced filtering criteria. If an object matches the filtering criteria, the object is copied and placed according to the rule's placement instructions. Placement instructions determine where, when, and how object data is stored. If the rule includes more than one placement instruction, when a set time expires, the content placement instructions for the next time period are applied to objects at the next ILM evaluation time.

## Steps

## 1. Select ILM > Rules.

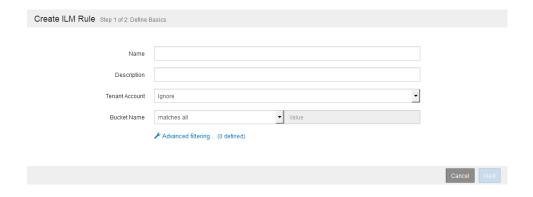
The ILM Rules page appears, with the stock rule, Make 2 Copies, selected.



**Note:** If the global Compliance setting has been enabled for the StorageGRID system, the ILM Rules page indicates which ILM rules are compliant. The summary table includes a **Compliant** column, and the details for the selected rule include a **Compliance Compatible** field. See "Managing S3 buckets and objects for compliance" for more information.

#### 2. Click Create.

Step 1 of the Create ILM Rule wizard appears.



- 3. Complete Step 1 of the Create ILM Rule wizard.
  - a. Enter a unique name for the rule in the **Name** field. You must enter between 1 and 64 characters.
  - b. 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	Make 3 Copies
Description	Save 1 copy at 3 sites for 1 year. Then, save EC copy forever

- c. From the **Tenant Account** drop-down list, optionally select the S3 or Swift tenant account to which this rule applies. If this rule applies to all tenants, select **Ignore** (default).
- d. Use the **Bucket Name** field to 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.

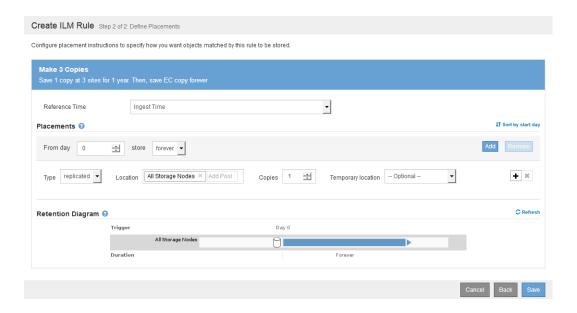
e. Optionally, click Advanced filtering, and specify additional filtering criteria.

If you do not configure advanced filtering, the rule applies to all objects that match the currently configured criteria (tenant account and bucket).

See "Using advanced filters in ILM rules" for information about the types of metadata, operators, and metadata values you can specify.

## 4. Click Next.

Step 2 of the wizard appears.



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.
Noncurrent Time	The time an object version became noncurrent because a new version was ingested and replaced it as the current version.
	<b>Note:</b> The Noncurrent Time applies only to S3 objects in versioning-enabled buckets.
	You can use this option to reduce the storage impact of versioned objects by filtering for noncurrent object versions updated with a new current version or delete marker.
User Defined Creation Time	A time specified in user-defined metadata.

Note: If you want to create a compliant rule, you must select Ingest Time. See "Managing S3 buckets and objects for compliance."

6. 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.

**Attention:** Do not specify an Archive Node or a Cloud Storage Pool as the location for the placement instruction starting on day 0. On day 0 (ingest), objects should only be stored on Storage Nodes. To store object copies in other locations, click **Add** to add a second time period starting on day 1 or later.

- 7. If you want to create replicated copies on one or more storage pools:
  - a. From the **Type** drop-down list, select **replicated**.
  - b. In the **Location** field, optionally click **Add Pool**. Then, select one or more storage pools. If you are specifying more than one storage pool, keep these rules in mind:
    - If you are specifying more than one storage pool and creating *n* copies, you must add *n* or more pools. For example, if you plan to specify 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, while ensuring that one copy goes
      only to one site.
    - 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 storage pool
      (All Storage Nodes) and another storage pool.



c. Select the number of copies you want to make.

**Attention:** In general, do not configure ILM rules to create only one replicated copy. If the only replicated copy is lost or corrupt, data will be lost. In addition, you might lose access to the object during maintenance or recovery operations. Use ILM to create a single replicated copy only if you are willing to risk the loss of object data when a failure occurs.

**Note:** 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, be aware that only three copies will be made—one copy for each Storage Node.

d. If you are using only a single storage pool, specify a temporary storage pool.



Specifying a temporary storage pool is optional, but recommended. If the preferred storage pool is unavailable, a copy is made in the temporary storage pool. As soon as the preferred

storage pool becomes available, a copy is made in the preferred storage pool, and the copy in the temporary storage pool is deleted.

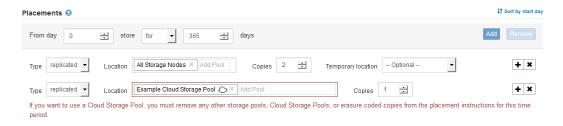
**Attention:** Failing to specify a temporary storage pool puts object data at risk if the preferred pool is unavailable.

- **8.** If you want to move objects to a Cloud Storage Pool:
  - a. From the **Type** drop-down list, select **replicated**.
  - b. In the Location field, remove All Storage Pools, and click Add Pool. Then, select the Cloud Storage Pool



If you want to use a Cloud Storage Pool, keep these rules in mind:

- You can create only one replicated copy.
- You cannot store the object in another storage pool or Cloud Storage Pool, or as an erasure coded copy, during the same time period. In the following example, the All Storage Nodes storage pool was selected from day 0 to day 365. An error message appears when a Cloud Storage Pool is selected for the same time period.



- **9.** If you want to create an erasure-coded copy:
  - a. From the **Type** drop-down list, select **erasure coded**.

The number of copies changes to 1.

b. Select the storage location.

The storage locations for an erasure-coded copy include the name of the storage pool, followed by the name of the Erasure Coding profile.



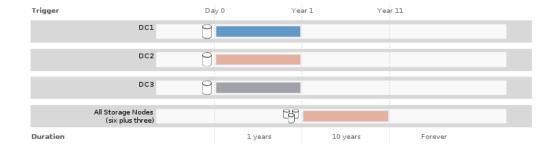
**Attention:** When adding a rule that makes an erasure-coded copy to the ILM policy, you must ensure that the policy has at least one rule that filters by Object Size. Due to the overhead of managing the number of fragments associated with an erasure-coded copy, do not erasure code objects smaller than 200 KB.

- 10. Optionally, add different time periods or create additional copies at different locations:
  - Click **Add** to add a different time period to the placement instructions.
  - Click the plus icon to create additional copies at a different location during the same time period.
- 11. 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, three replicated copies will be saved to three storage pools (DC1, DC2, and DC3) for one year. Then, an erasure-coded copy will be saved for 10 years, using a 6+3 erasure-coding scheme. After 10 years, the objects will be deleted from StorageGRID.



## 12. 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 concepts

Using advanced filters in ILM rules on page 186
Configuring storage pools on page 158
Creating, simulating, and activating an ILM policy on page 192

### Related tasks

Using Ingest Time or Last Access Time in ILM rules on page 190 Configuring an Erasure Coding profile on page 177 Managing S3 buckets and objects for compliance on page 220

## 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.

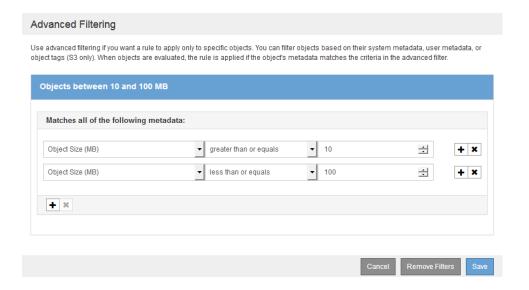
Metadata type	Supported operators	Metadata value	Obje	ct type
			S3	Swift
Ingest Time (microseconds)	<ul> <li>equals</li> <li>does not equal</li> <li>less than</li> <li>less than or equals</li> <li>greater than</li> <li>greater than or equals</li> </ul>	Time and date the object was ingested, in microseconds since Unix Epoch.  See "Using Ingest Time or Last Access Time in ILM rules" for more information on how to calculate this value.	Yes	Yes
Key	<ul> <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>	All or part of a unique S3 or Swift object key. For example, you might want to match objects that end with .txt or start with test-object/.	Yes	Yes
Last Access Time (microseconds)	<ul> <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>	Time and date the object was last retrieved (read or viewed) in microseconds since Unix Epoch.  See "Using Ingest Time or Last Access Time in ILM rules" for more information on how to calculate this value.  Note: 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.	Yes	Yes

Metadata type	Supported operators	Metadata value	Obje	ct type
			S3	Swift
Object Tag (S3 only)	<ul> <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>	Key-value pair For example, to filter on objects that have an object tag of Image=True, specify Image for Object Tag Name, equals for the operator, and True for Object Tag Value.  Note: Object tag names and object tag values are case sensitive. You must enter these items exactly as they were defined for the object.	Yes	No

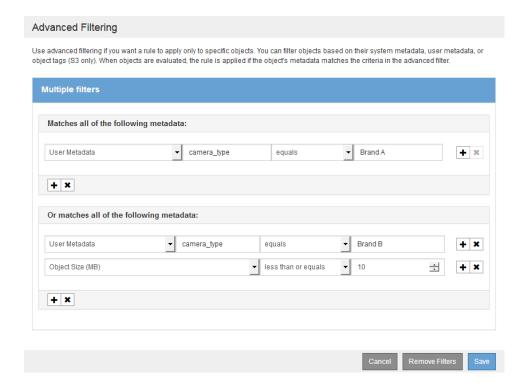
## Specifying multiple metadata types and values

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



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.



#### Related tasks

Using Ingest Time or Last Access Time in ILM rules on page 190 Configuring regions (optional and S3 only) on page 179

### Related information

Implementing S3 client applications
Implementing Swift client applications

## Using Ingest Time or Last Access Time in ILM rules

You can use advanced filtering if you want an ILM rule to apply only to objects that were ingested or last accessed on a specific date. You can also use Ingest Time or 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 past month on local Storage Nodes, while moving objects that have not been viewed as recently to an off-site location.

## About this task

When using Ingest Time or Last Access Time as an advanced filter, you must convert the desired time and date to microseconds since Unix Epoch.

When using Last Access Time as an advanced filter or as a reference time, you must enable last access time updates for S3 buckets.

**Note:** Last access time updates are always enabled for Swift containers, but are disabled by default for S3 buckets.

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 addisabled (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><li>No, for the source copy</li><li>Yes, for the destination copy</li></ul>	<ul><li>No, for the source copy</li><li>Yes, for the destination copy</li></ul>	<ul><li>Yes, for the source copy</li><li>Yes, for the destination copy</li></ul>	<ul><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. If you are using Ingest Time or Last Access Time as advanced filters, determine the UTC date and time you want to use in the filter.

You might need to convert from your local time zone to UTC.

2. Convert the UTC date and time to microseconds since Unix Epoch.

For example, use date from a Linux command prompt.

```
# date -d '2015-03-14 00:00:00 UTC' +%s000000
14262912000000
```

3. If you are using Last Access Time as an advanced filter or as a reference time, enable last access time updates on each S3 bucket specified in that rule.

You can use the Tenant Manager or the Tenant API to enable updates to last access time for S3 buckets. See the instructions for using tenant accounts.

Attention: Be aware that enabling last access time updates can reduce performance, especially in systems with small objects. The performance impact occurs 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

#### Related information

# Creating, simulating, and activating an ILM policy

When you create a 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.

When creating an ILM 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. If no rules match an object, the policy's default rule controls where that object is placed and for how long it is retained.
- Make sure that the rules in the policy are in the correct order. After an object has been matched by a rule, none of the following rules in the policy are applied to that object.
- 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.

**Caution:** 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.

Creating a policy consists of these main tasks:

## **Steps**

- 1. Creating a proposed ILM policy on page 192
- 2. Simulating an ILM policy on page 196
- **3.** Activating the ILM policy on page 203
- 4. Verifying an ILM policy with object metadata lookup on page 204

## **Related concepts**

What an information lifecycle management policy is on page 150

## 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.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You have created the rules you want to add to the proposed policy. Note that you can save a proposed policy, create additional rules, and then edit the policy to add the new rules.

#### About this task

Typical reasons for creating a proposed ILM policy include:

- You made changes to a storage pool.
- You started to use a Cloud Storage Pool.

New storage retention requirements were defined. For example, you need to comply with regulations that require object data tobe preserved for a specified amount of time.

The proposed ILM policy must include at least one ILM rule.

## **Steps**

#### 1. Select ILM > Policies.

The ILM Policies page appears. The Baseline 2 Copies Policy is the system's default policy. This policy includes the stock ILM rule (Make 2 Copies), and it applies if no other policy has been activated.

From this page, you can review the proposed, active, and historical policies; create, edit, or remove a proposed policy; clone the active policy; or view the details for any policy.



Note: If the global Compliance setting is enabled, the ILM Policies page indicates which ILM rules are compliant. See "Managing S3 buckets and objects for compliance."

## 2. Determine how you want to create the proposed ILM policy.

Option	St	eps
Create a new proposed policy that has no rules already selected	a.	If a proposed ILM policy currently exists, select that policy, and click <b>Remove</b> .  You cannot create a new proposed policy if a proposed policy already exists.
	b.	Click Create Proposed Policy.
Create a proposed policy based on the active policy	a.	If a proposed ILM policy currently exists, select that policy, and click <b>Remove</b> .  You cannot clone the active policy if a proposed policy already exists.
	b.	Select the active policy from the table.
	c.	Click Clone.
Edit the existing proposed policy	a. b.	Select the proposed policy from the table.  Click <b>Edit</b> .

The Configure ILM Policy dialog box appears.

If you are creating a new proposed policy, all fields are blank and no rules are selected.



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.



3. Enter a unique name for the proposed policy in the Name field.

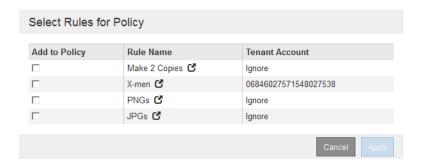
You must enter between 1 and 64 characters. If you are cloning the active policy, you can use the current name with the appended version number or you can enter a new name.

**4.** Enter the reason you are creating a new proposed policy in the **Reason for change** field.

You must enter between 1 and 128 characters.

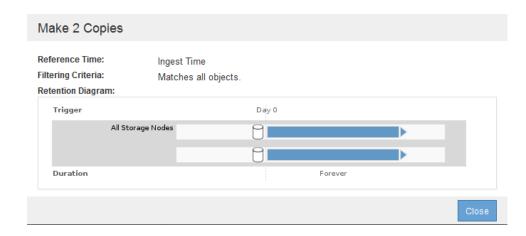
**5.** To add rules to the policy, click **Select Rules**.

The Select Rules for Policy dialog box appears, with all defined rules listed. If you are cloning the active policy, any rules that are currently in use are selected.



**6.** Select and unselect the check boxes to choose the rules you want to add to the proposed policy.

You can click the rule name or the more details icon to view the settings for each rule. Click Close when you are done viewing rule details.

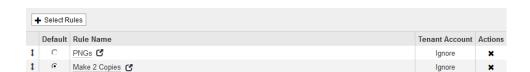


- 7. When you are done selecting rules for the proposed policy, click **Apply**.
- **8.** Drag the rows to reorder the rules in the proposed policy.

**Caution:** You must confirm that the ILM rules are in the correct order. When the policy is applied to an object, object metadata is evaluated by the rules in the order listed, starting at the top.

9. Select the radio button to specify which rule you want to be the default rule for this policy. Every ILM policy must contain one default ILM rule. The placement instructions for the default

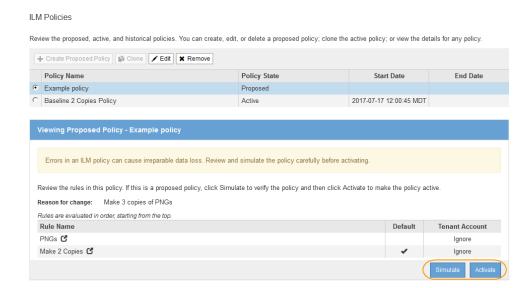
rule are applied to any objects that are not matched by the other rules in the policy.



**Note:** If the global Compliance setting is enabled, the default rule must be a compliant rule. See "Managing S3 buckets and objects for compliance."

- 10. As required, click the delete icon \* to delete any rules that you do not want in the policy, or click Select Rules to add more rules.
- 11. When you are done, click Save.

The ILM Policy page updates, and the policy you saved is shown as Proposed. Proposed policies do not have start and end dates. The Simulate and Activate buttons are now enabled.



## Related concepts

What an information lifecycle management policy is on page 150

#### Related tasks

Managing S3 buckets and objects for compliance on page 220

## 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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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.

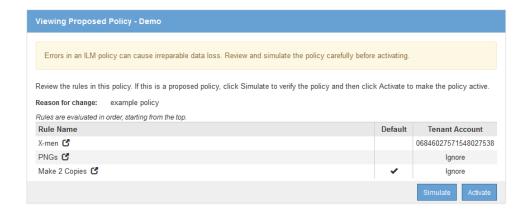
- The Simulation page retains the objects you tested until you close, navigate away from, or refresh the ILM Policies page.
- Simulation returns the name of the matched rule. To determine which storage pool or Erasure Coding profile is in effect, you can view the Retention Diagram by clicking the rule 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 Demo policy in this example has three rules:

- The first rule, X-men, applies only to objects in a specific tenant account and uses an advanced filter to match series=x-men user metadata.
- The second rule, PNGs, matches keys ending in .png.
- The last rule is the stock rule, Make 2 Copies, and it is selected as the default. This rule applies to any objects that do not match the other rules.



## 2. Click Simulate.

The Simulation ILM Policy dialog box appears.

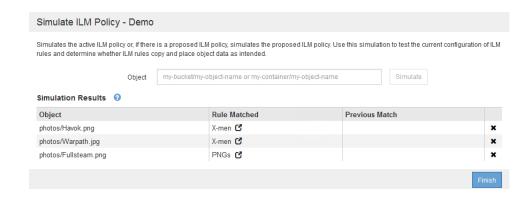
3. In the Object field, enter the S3 bucket/object-key or the Swift container/object-name for a test object, and click Simulate.

Note: A message appears if you specify an object that has not been ingested.



4. 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. None of the test objects was matched by the Make 2 Copies rule



## **Examples for simulating ILM policies**

These examples show how you can verify ILM rules by simulating the ILM policy before activating it.

#### Choices

- Example 1: Verifying rules when simulating a proposed ILM policy on page 198
- Example 2: Reordering rules when simulating a proposed ILM policy on page 199
- Example 3: Correcting a rule when simulating a proposed ILM policy on page 201

## Example 1: Verifying rules when simulating a proposed ILM policy

This example shows how to verify rules when simulating a proposed policy.

#### About this task

In this example, the "xmen photo files" policy is being simulated against the ingested objects in a bucket called "photos." The policy includes three rules, as follows:

- The first rule, X-men, applies only to objects in a specific tenant account and filters for series=x-men user metadata. This rule is also marked as the default rule for the policy.
- The second rule, PNGs, filters for keys ending with .png.
- The last rule, JPGs, filters for keys ending with . jpg.



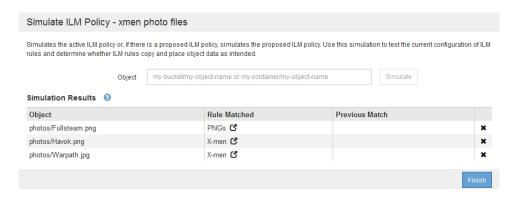
#### **Steps**

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.



3. Confirm that each object was matched by the correct rule.

In this example:

- a. The Fullsteam.png object did not match the X-men rule, but did match the PNGs rule.
- b. The Havok.png and the Warpath.jpg objects both matched the X-men rule, which was evaluated first.

Note that even if these two files had not matched the X-men rule, they would have matched one of the subsequent rules: either PNGs or JPGs.

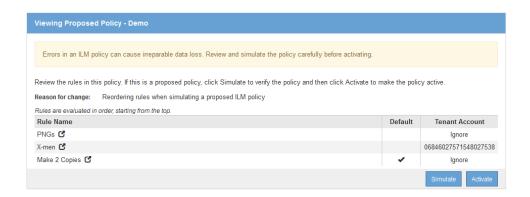
## Example 2: Reordering rules when simulating a proposed ILM policy

This example shows how you can reorder rules to change the results when simulating a policy.

#### About this task

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:

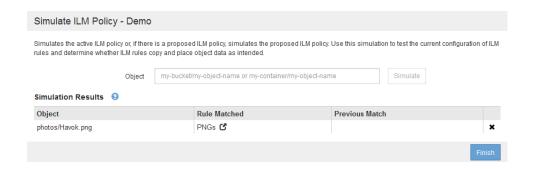
- The first rule, PNGs, filters for key names that end in .png.
- The second rule, X-men, applies only to objects in a specific tenant account and filters for series=x-men user metadata.
- The last rule is the default rule, Make 2 Copies, which will match any objects that do not match the first two rules.



## **Steps**

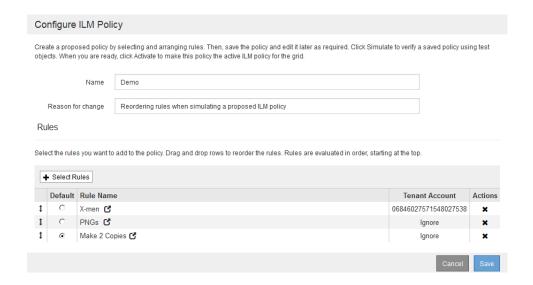
- 1. After adding the rules and saving the policy, click **Simulate**.
- 2. In the **Object** field, enter the S3 bucket/object-key or the Swift container/object-name for a test object, and click **Simulate**.

The Simulation Results appear, showing that the Havok.png object was matched by the PNGs rule.



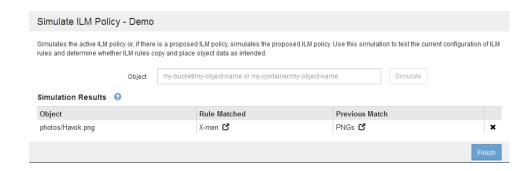
However, the rule that the Havok.png object was meant to test was the X-men rule.

- **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.



- 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 <code>Havok.png</code> 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.



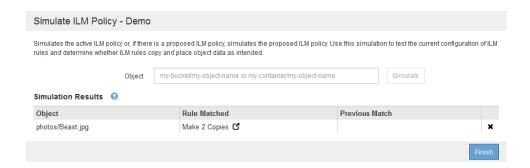
Note: If you stay on the Configure Policies page, you can re-simulate a policy after making changes without needing to re-enter the names of the test objects.

## 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.

#### About this task

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, Make 2 Copies, rule.



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."



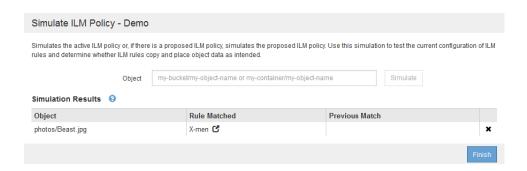
- **3.** To resolve the error, correct the rule, as follows:
  - If the rule is part of the proposed policy, you can either clone the rule or remove the rule from the policy and then edit it.
  - If the rule is part of the active policy, you must clone the rule. You cannot edit or remove a rule from the active policy.

Option	De	scription
Cloning the rule	a.	Select ILM > Rules.
	<b>b.</b>	Select the incorrect rule, and click Clone.
	c.	Change the incorrect information, and click Save.
	d.	Select ILM > Policies.
	e.	Select the proposed policy, and click <b>Edit</b> .
	f.	Click Select Rules.
	g.	Select the check box for the new rule, uncheck the check box for the original rule, and click <b>Apply</b> .
	h.	Click Save.
Editing the rule	a.	Select the proposed policy, and click <b>Edit</b> .
	<b>b.</b>	Click the delete icon * to remove the incorrect rule, and click Save.
	c.	Select <b>ILM &gt; Rules</b> .
	d.	Select the incorrect rule, and click Edit.
	е.	Change the incorrect information, and click Save.
	f.	Select ILM > Policies.
	g.	Select the proposed policy, and click <b>Edit</b> .
	h.	Select the corrected rule, click <b>Apply</b> , and click <b>Save</b> .

**4.** Perform the simulation again.

**Note:** 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.



## 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.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You have saved and simulated the proposed ILM policy.

Caution: 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 a new ILM policy goes into effect, the StorageGRID system immediately uses it to manage all objects in the grid, including existing objects and newly ingested objects.

#### 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 uses a Make 3 Copies rule instead of a Make 2 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 Make 2 Copies rule instead of a Make 3 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.



#### 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.



• 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.

## 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.

## Before you begin

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

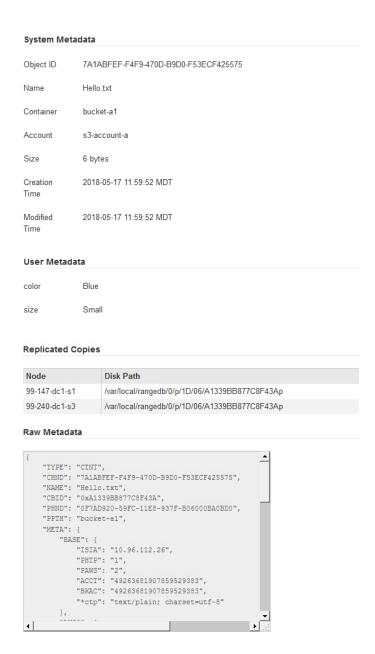
Object Me	Object Metadata Lookup		
Enter the ide	Enter the identifier for any object stored in the grid to view its metadata.		
Identifier	bucket-a1/Hello.txt	Look Up	

## 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, including the name of the grid node and the full path to the disk location of the object.
- For erasure-coded object copies, the current storage location of each fragment, including the name of the grid node and the type of fragment (data or parity).
- 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, 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.

The following example shows the object metadata lookup results for an S3 test object that is stored as two replicated copies.



**5.** Confirm that the object is stored in the correct location or locations and that it is the correct type of copy.

**Note:** 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 concepts

Configuring audit client access on page 290

## **Related information**

Understanding audit messages
Implementing S3 client applications

Implementing Swift client applications

# 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.

#### **Choices**

- Deleting an ILM rule on page 207
- Editing an ILM rule on page 207
- Cloning an ILM rule on page 208
- Viewing the ILM policy activity queue on page 209

## Deleting an ILM rule

To keep the list of current ILM rules manageable, delete any ILM rules that you are not likely to use. You cannot delete the stock ILM rule (Make 2 Copies), ILM rules listed in the active policy, or ILM rules currently listed in the proposed policy.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### Steps

- 1. Select ILM > Rules.
- 2. Select the ILM rule you want to delete, and click **Remove**.
- 3. Click **OK** to confirm that you want to delete the ILM rule.

The ILM rule is deleted.

## Related concepts

Creating, simulating, and activating an ILM policy on page 192

## Editing an ILM rule

You might need to edit an ILM rule to change a filter or placement instruction. You cannot edit the stock ILM rule (Make 2 Copies), ILM rules listed in the active policy, ILM rules listed in the proposed policy, or ILM rules created before StorageGRID version 10.3.

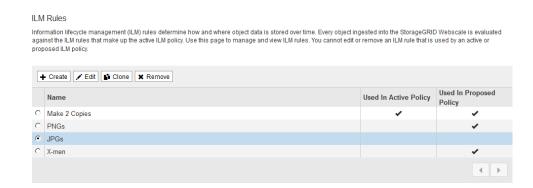
## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### Steps

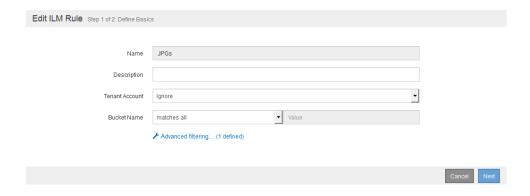
1. Select ILM > Rules.

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.



2. Select a rule that is not being used, and click Edit.

The Edit ILM Rule wizard opens.



**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.

## **Related concepts**

Using advanced filters in ILM rules on page 186

## Related tasks

Creating an ILM rule on page 181

## 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.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- · You have specific access permissions.

## 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.



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 concepts

Using advanced filters in ILM rules on page 186

#### Related tasks

Editing an ILM rule on page 207

## Viewing the ILM policy activity queue

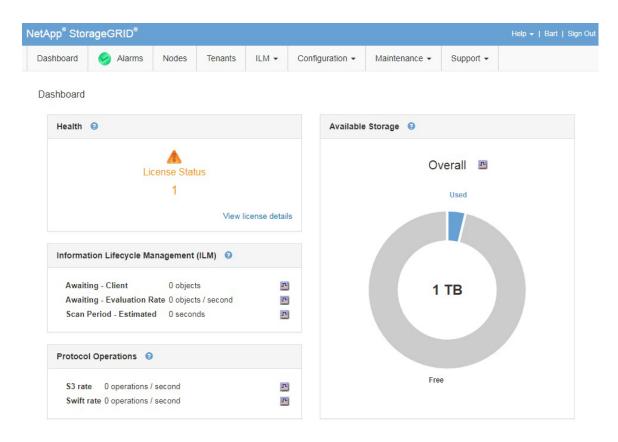
You can view the number of objects that are in the queue to be evaluated against the ILM policy at any time. You might want to monitor the ILM processing queue to determine system performance. A large queue might indicate that the system is not able to keep up with the ingest rate, the load from the client applications is too great, or that some abnormal condition exists.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

## **Steps**

1. Select Dashboard.



2. Monitor the Information Lifecycle Management (ILM) section.

You can click the question mark ② to see a description of the items in this section.

# **Example ILM rules and policies**

You can use the following examples as starting points for defining 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 on page 214
- Example 3: ILM rules and policy for better protection for image files on page 216

## Example 1: ILM rules and policy for object storage

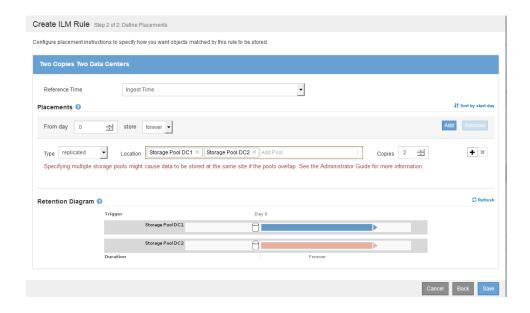
You can use the following example rules and policy as a starting point when defining an ILM policy to meet your object protection and retention requirements.

**Caution:** The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Carefully analyze your ILM rules before adding them to an ILM policy to confirm that they 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.

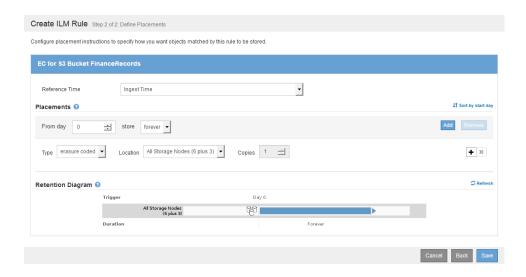
Rule definition	Example value
Storage Pools	Two storage pools, each at different data centers, named Storage Pool DC1 and Storage Pool DC2.
Rule Name	Two Copies Two Data Centers
Reference Time	Ingest Time
Content Placement	On Day 0, keep two replicated copies forever—one in Storage Pool DC1 and one in Storage Pool DC2.



## 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.

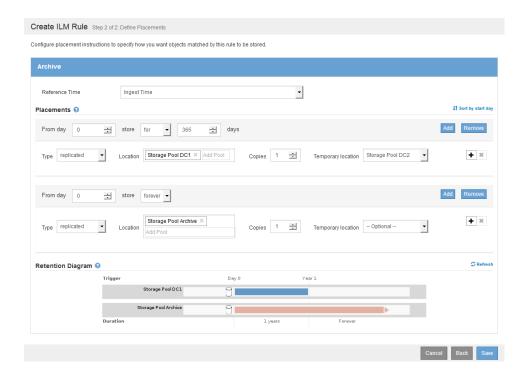
Rule definition	Example value
Erasure Coding Profile	<ul> <li>One storage pool across three data centers</li> <li>Use 6+3 Erasure Coding scheme</li> </ul>
Rule Name	EC for S3 Bucket FinanceRecords
Reference Time	Ingest Time
Content Placement	For objects in the S3 Bucket FinanceRecords, create one erasure-coded copy in the pool specified by the Erasure Coding profile. Keep this copy forever.



## ILM rule 3 for example 1: Store object to DC1 and Archive

This example ILM rule creates two copies. One copy is stored in Data Center 1 for one year, and the second copy is stored in an Archive Node forever.

Rule definition	Example value
Storage Pools	A disk storage pool and an archive storage pool.
Rule Name	Archive
Reference Time	Ingest Time
Content Placement	<ul> <li>On Day 0, keep a replicated copy in Storage Pool DC1 for 365 days; temporary copies in Storage Pool DC2</li> <li>On Day 0, keep a replicated copy in Storage Pool Archive forever; temporary</li> </ul>
	copies in All Storage Nodes

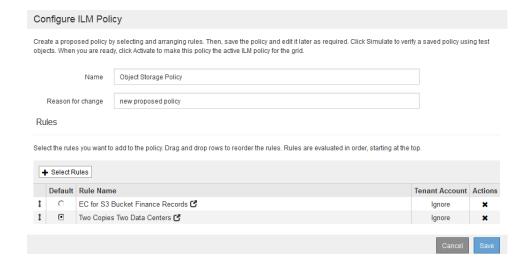


## ILM policy for example 1

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 4+2 Erasure Coding to store all objects belonging to the S3 Bucket FinanceReports 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.



## Example 2: ILM rules and policy for EC object size filtering

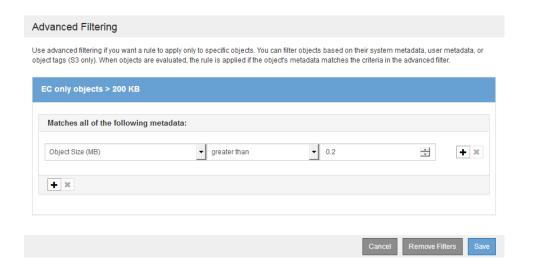
You can use the following example rules and policy as starting points to define an ILM policy that filters by object size to meet recommended EC requirements.

**Caution:** The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Carefully analyze your ILM rules before adding them to an ILM policy to confirm that they will work as intended to protect content from loss.

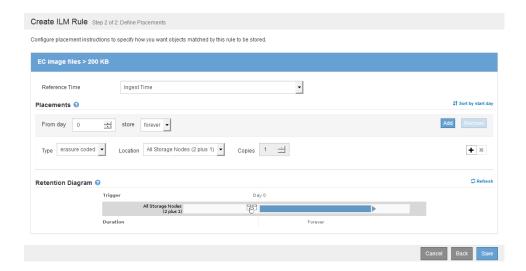
## ILM rule 1 for example 2: Use EC for all objects larger than 200 KB

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 1 erasure-coded copy in all Storage Nodes



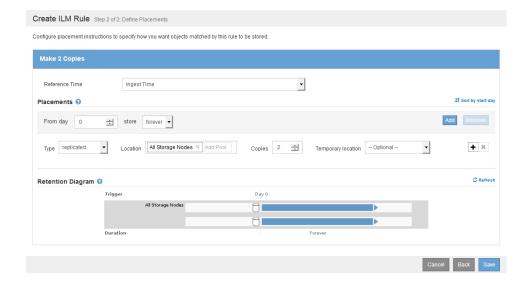
The placement instructions specify that one erasure-coded copy be created in all Storage Nodes.



## ILM rule 2 for example 2: Make 2 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	Make 2 Copies
Reference Time	Ingest Time
Advanced Filtering for Object Size	None
Content Placement	Create 2 replicated copies in all Storage Nodes

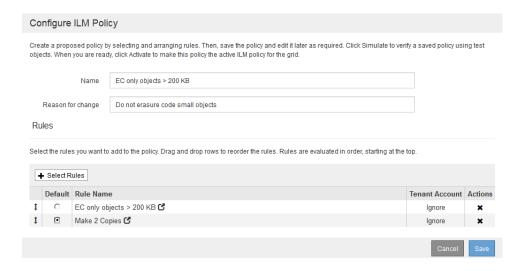


## ILM policy for example 2: Use EC for objects larger than 200 KB

In this example policy, objects larger than 200 KB are erasure coded, and any other objects that are smaller than 200 KB are replicated using the default catch-all Make 2 Copies rule.

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.



## Example 3: ILM rules and policy for better protection for image files

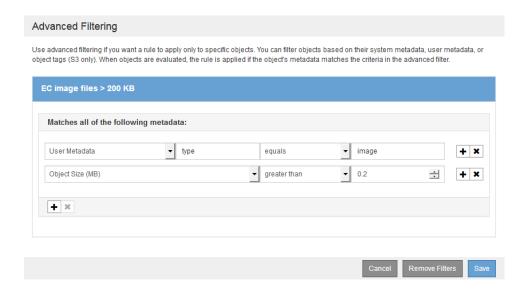
You can use the following example rules and policy to ensure that images larger than 200 KB are erasure coded and that three copies are made of smaller images.

**Caution:** The following ILM rules and policy are only examples. There are many ways to configure ILM rules. Carefully analyze your ILM rules before adding them to an ILM policy to confirm that they will work as intended to protect content from loss.

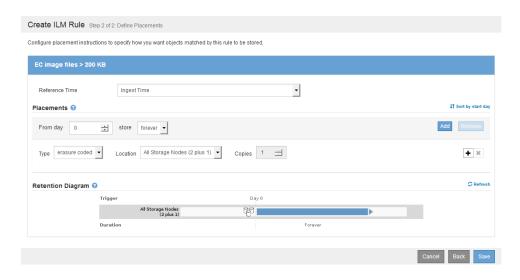
## ILM rule 1 for example 3: Use EC for image files larger than 200 KB

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



Because this rule is configured as the first rule in the policy, the erasure coding placement instructions will only apply to images that are larger than 200 KB.

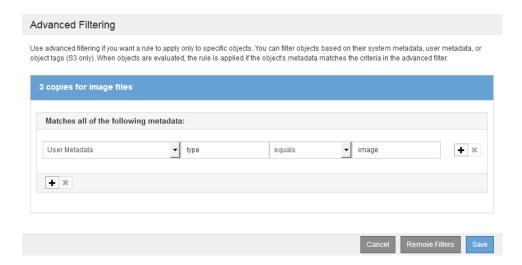


# ILM rule 2 for example 3: Replicate 3 copies for all remaining image files

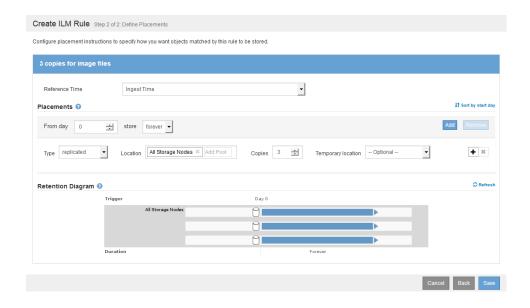
This example ILM rule uses advanced filtering to specify that image files be replicated.

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



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.

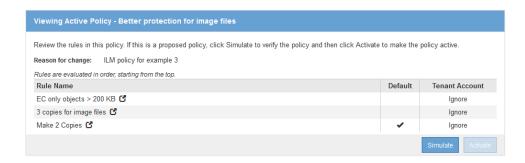


# ILM policy for example 3: Better protection for image files

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 stock rule, Make 2 Copies, as the default to any remaining objects (that is, all nonimage files).



# Managing S3 buckets and objects for compliance

In some cases, an S3 tenant account might need to comply with regulations that require object data to be preserved for a specified amount of time. If your StorageGRID system includes such tenants, you can enable the global Compliance setting and create compliant ILM rules to manage the objects in compliant S3 buckets.

# How StorageGRID protects compliant data

When StorageGRID is properly configured and when compliant S3 buckets, ILM rules, and ILM policies have been correctly applied, StorageGRID provides functionality that prevents objects in S3 buckets from being overwritten, deleted, or altered until the specified retention period has expired.

StorageGRID meets the relevant storage requirements of these regulations:

- US SEC (Securities and Exchange Commission) regulation 17 CFR § 240.17a-4(f), which regulates exchange members, brokers or dealers.
- US CFTC (Commodity Futures Trading Commission) regulation 17 CFR § 1.31(b)-(c)1.31(b)-(c), which regulates commodity futures trading.

# Compliance and retention

When the global Compliance setting is enabled for the StorageGRID system, S3 tenant users can create compliant buckets for object data, such as legal and financial records, that needs to be preserved for a certain amount of time. When creating a compliant bucket, users can specify the retention period for bucket objects and select whether object data will be automatically deleted when the retention period expires.

Each object's retention period starts when the object is ingested into the bucket. During the retention period, the object can be retrieved, but it cannot be modified or deleted. As required, tenant users can increase a bucket's retention period, place the bucket under a legal hold (meaning that objects cannot be deleted when their retention period expires), remove a legal hold, or change the auto-delete setting.

### Compliance and the storage of duplicate data

StorageGRID ensures that duplicate copies of each compliant object are stored on the grid. When the global Compliance setting is enabled, the active and any proposed ILM policies must use a compliant ILM rule as their default rule. Compliant rules create at least two replicated object copies or one erasure-coded copy on Storage Nodes. These copies must exist from day 0 until the retention period expires and the objects are deleted.

#### Compliance and security features

StorageGRID protects compliant objects with the following platform security features:

- Internal public key infrastructure and node certificates are used to authenticate and encrypt internode communication. Internode communication is secured by TLS.
- Rules for firewalls and iptables are automatically configured to control incoming and outgoing network traffic, as well as closing unused ports.
- The base operating system of StorageGRID appliances and virtual nodes is hardened; unrelated software packages are removed.

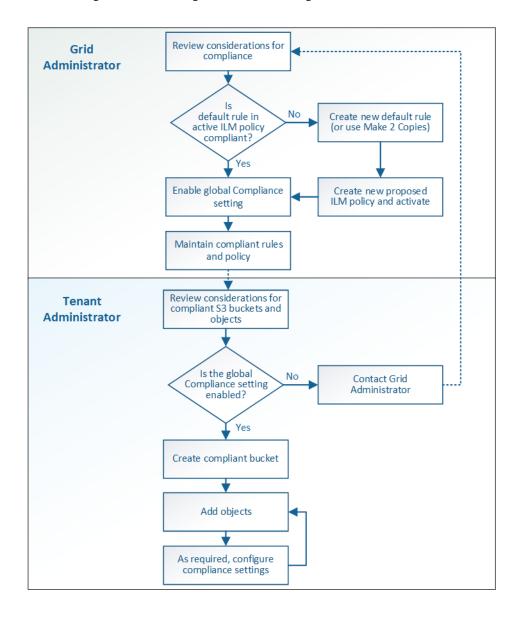
- Root login over SSH is disabled on all grid nodes. SSH access between nodes uses certificate authentication.
- Separate networks are available for Client, Admin, and internal Grid traffic

# **Compliance workflow**

The workflow diagram shows the high-level steps for enabling the global Compliance setting, for creating and managing compliant ILM rules and ILM policies, and for creating and managing S3 buckets that are compliant. As a grid administrator, you must coordinate closely with the tenant administrator to ensure that the objects in compliant buckets are protected in a manner that satisfies regulatory requirements.

As the workflow diagram shows, a grid administrator must enable the global Compliance setting for the entire StorageGRID system before an S3 tenant can create compliant buckets. The grid administrator must also ensure that the default rule in the grid's active ILM policy satisfies the dataprotection requirements of objects in compliant buckets.

Then, once the global Compliance setting has been enabled, tenants can create and manage compliant buckets using the Tenant Manager, the Tenant Management API, or the S3 REST API.



### **Related information**

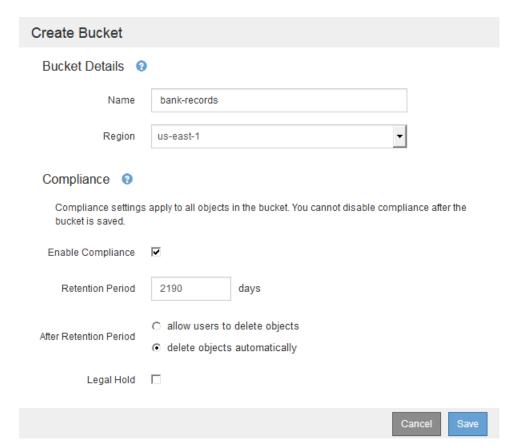
Using tenant accounts
Implementing S3 client applications

# **Considerations for compliance**

Make sure you review the considerations for using the global Compliance setting as well as the restrictions StorageGRID places on compliant buckets, compliant objects, and compliant ILM rules and policies.

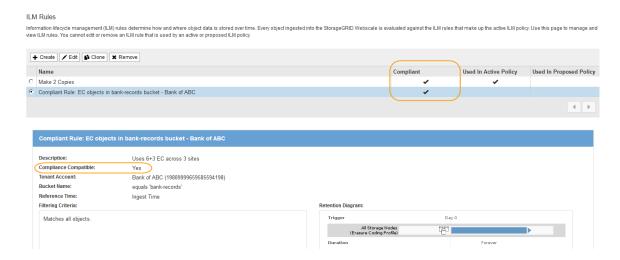
# Considerations for using the global Compliance setting

- You must enable the global Compliance setting before any S3 tenant can create a compliant bucket.
- After you enable the global Compliance setting, you cannot disable this setting.
- Enabling the global Compliance setting allows all S3 tenant accounts to use the Tenant Manager, the Tenant Management API, or the S3 REST API to create and manage compliant buckets. Users with the appropriate permissions can create compliant buckets, set and increase the retention period for objects in the bucket, specify how objects can be deleted at the end of their retention period, and optionally place all objects in the bucket under a legal hold or lift a legal hold. For example, this tenant user is creating a compliant bucket named bank-records in the default us-east-1 region. Objects in this bucket will be retained for 6 years and then deleted automatically. This bucket is not currently under a legal hold.



When the global Compliance 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 satisfies the requirements of S3 compliant buckets. The ILM Rules and ILM Policies pages indicate which ILM rules are compliant.

In the following example, the ILM Rules page lists two rules that are compatible with compliant buckets.



## Restrictions for using compliant buckets

- If S3 tenants need to create compliant buckets, they must enable compliance and specify compliance settings when they create the bucket. After a bucket has been saved, compliance cannot be disabled for the bucket.
- The retention period for the bucket specifies the minimum amount of time each object in that bucket must be preserved (stored) within StorageGRID.
- Tenant users can edit bucket settings to increase the retention period, but they can never decrease this value.
- If a tenant account is notified of a pending legal action or regulatory investigation, users can preserve relevant information by placing a legal hold on the bucket. When a bucket is under a legal hold, no object in that bucket can be deleted even if its retention period has ended. As soon as the legal hold is lifted, objects in the bucket can be deleted when their retention periods end.
- Objects can be added to a compliant bucket at any time, regardless of the bucket's compliance settings.
- Objects can be retrieved from a compliant bucket at any time, regardless of the bucket's compliance settings.
- Versioning is not supported for compliant buckets.

# Restrictions for objects in compliant buckets

Each object that is saved in a compliant bucket goes through three stages:

## 1. Object ingest

When an object is ingested, the system generates metadata for the object that includes a unique object identifier (UUID) and the ingest date and time. The object inherits the compliance settings from the bucket.

- After an object is ingested into a compliant bucket, its data, S3 user-defined metadata, or S3 object tags cannot be modified, even after the retention period expires.
- StorageGRID maintains three copies of all object metadata at each site to provide redundancy and protect object metadata from loss. Metadata is stored independently of object data.

## 2. Retention period

- The retention period for an object starts when the object is ingested into the bucket.
- Each time the object is accessed or looked up, the compliance settings for the bucket are also looked up. The system uses the object's ingest time and date and the bucket's retention period setting to calculate when the object's retention period will expire.
- During an object's retention period, 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.

**Note:** As required, you might need to add new ILM rules to manage the objects in a particular bucket.

 During an object's retention period, or when legal hold is enabled for the bucket, the object cannot be deleted.

# 3. Object deletion

- When an object's retention period ends, all copies of the object can be deleted, unless legal hold is enabled for the bucket.
- When an object's retention period ends, a bucket-level compliance setting allows tenant users to control how objects are deleted: by users when required or automatically by the system.
- If the bucket setting is to delete objects automatically, all copies of the object are removed by the scanning ILM process in StorageGRID. When an object's retention period ends, the object is scheduled for deletion. You can look for the IDEL (ILM Initiated Delete) message in the audit log to determine when ILM has started 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).

**Note:** The actual amount of time needed to delete all object copies can vary, depending on the number of objects in the grid and how busy the grid processes are.

# **Restrictions for compliant ILM rules**

If you want to enable the global Compliance setting, you must ensure that the default rule in your active ILM policy is compliant. A compliant rule satisfies the requirements of compliant S3 buckets:

- 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." The actual meaning of "forever" is determined by the compliance settings for each bucket.

For example, this rule satisfies the requirements of compliant S3 buckets. It stores three replicated object copies from Ingest Time (day 0) to "forever." The objects will be stored on Storage Nodes at three data centers.



**Note:** The Make 2 Copies stock rule is compliant. You can use it as the default rule in a compliant policy.

When you configure the placement instructions for a compliant rule, you must consider where the object copies will be stored. For example, if your deployment includes more than one site, you can enable site-loss protection for compliant objects by creating a storage pool for each site and specifying both storage pools in the rule's placement instructions. See "Using multiple storage pools for cross-site replication."

## Restrictions for active and proposed ILM policies

When the global Compliance 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 non-compliant buckets.
- Compliant rules can apply to objects in any compliant or non-compliant bucket.

As illustrated in "Example: Using compliant ILM rules in an ILM policy," a compliant ILM policy might include these three rules:

- 1. A compliant rule that creates erasure-coded copies of the objects in a specific compliant S3 bucket. 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 non-compliant buckets because it stores only one object copy forever and it uses Archive Nodes.
- 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 compliant or non-compliant bucket that was not filtered out by the first two rules.

#### Related concepts

Using multiple storage pools for cross-site replication on page 161 Example: Using compliant ILM rules in an ILM policy on page 227

#### Related information

Using tenant accounts Implementing S3 client applications Understanding audit messages

# **Enabling compliance**

Enabling the global Compliance setting allows all S3 tenant accounts to create and manage compliant buckets. If S3 tenant accounts need to comply with regulatory requirements when saving object data, you can enable compliance for your entire StorageGRID system.

## Before you begin

- You must have specific access permissions. For details, see information about controlling system
  access with administration user accounts and groups.
- You must be signed in to the Grid Manager using a supported browser.
- You must have reviewed the compliance workflow, and you must understand the considerations for compliance.

#### About this task

Users with the appropriate permissions can create compliant buckets, set and increase the retention period for bucket objects, specify how objects can be deleted at the end of their retention period, and optionally place all objects in the bucket under a legal hold or lift a legal hold.

Attention: If you enable this setting, you will not be able to disable it in the future.

### **Steps**

1. Select Configuration > Compliance.

The Global Compliance Settings page appears.

Global Compliance Settings

Enable compliance for your StorageGRID Webscale system if S3 tenant accounts need to comply with regulatory requirements when saving object data. Be aware that you cannot disable this feature after it has been enabled.

# Compliance

Before enabling compliance for the entire StorageGRID Webscale grid, you must ensure that the default rule in the active ILM policy is compliant. A compliant rule satisfies the requirements of compliant S3 buckets.

- 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". The actual meaning of "forever" is determined by the compliance setting for each bucket.

The Make 2 Copies stock rule is compliant and can be used as the default rule

☐ Enable Compliance



- 2. Select Enable Compliance.
- 3. Click Apply.

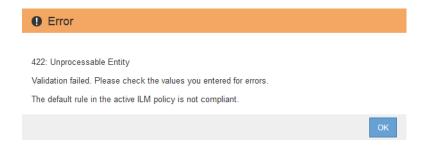
A confirmation dialog box appears.



4. If you are sure you want to enable compliance for the grid, click OK.

When you click OK:

- If the default rule in the active ILM policy is compliant, compliance 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. Click **OK**, and create a new proposed policy, simulate it, and activate it.



## Related concepts

Creating, simulating, and activating an ILM policy on page 192

### Related tasks

Creating an ILM rule on page 181

# Example: Using compliant ILM rules in an ILM policy

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 compliant S3 buckets.

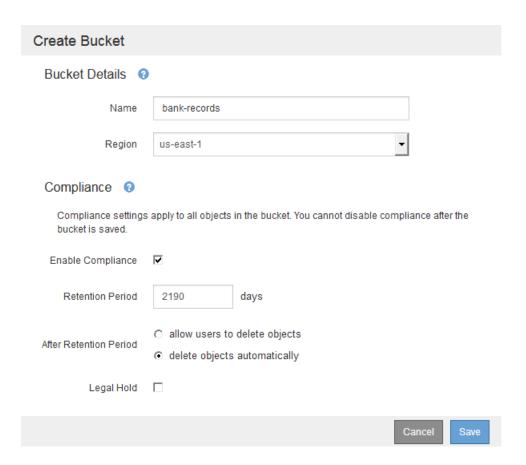
**Attention:** The following S3 bucket, ILM rules, and ILM policy are examples. There are many ways to configure compliant ILM rules. Carefully analyze your ILM rules before adding them to an ILM policy to confirm that they will work as intended to protect content from loss.

# S3 bucket for compliance example: bank-records

In this example, an S3 tenant account named Bank of ABC has created a compliant bucket, bank-records, to store critical bank records. To comply with regulatory requirements, objects in this bucket must be retained for six years. At the end of that time, objects must be deleted.

<b>Bucket definition</b>	Example value
Tenant Account Name	Bank of ABC

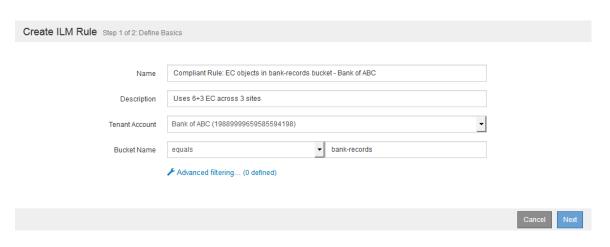
<b>Bucket definition</b>	Example value
Bucket Name	bank-records
Bucket Region	us-east-1 (default)
Retention Period	6 years (2,190 days)
After Retention Period	Objects will be deleted automatically
Legal Hold	Not in effect



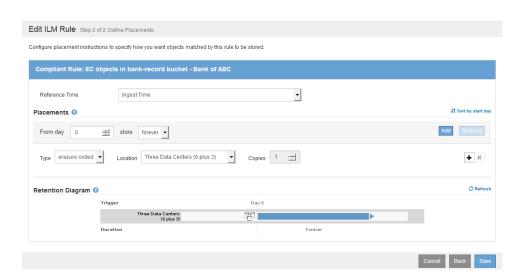
# ILM rule 1 for compliance example: Erasure Coding profile with bucket matching

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 places an erasure-coded copy of the object on Storage Nodes at three data center sites using a 6+3 Erasure Coding profile. This rule satisfies the requirements of compliant S3 buckets: erasure-coded copies are 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
Reference Time	Ingest Time
Placements	From day 0 store forever
Erasure Coding Profile	<ul> <li>Create one erasure-coded copy on Storage Nodes at three data center sites</li> <li>Uses 6+3 erasure coding scheme</li> </ul>

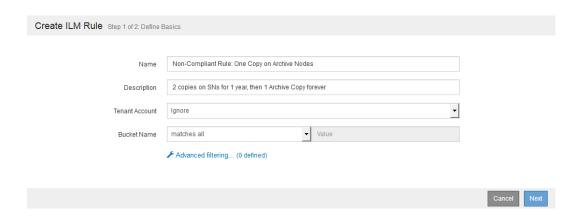


# ILM rule 2 for compliance example: Non-compliant rule

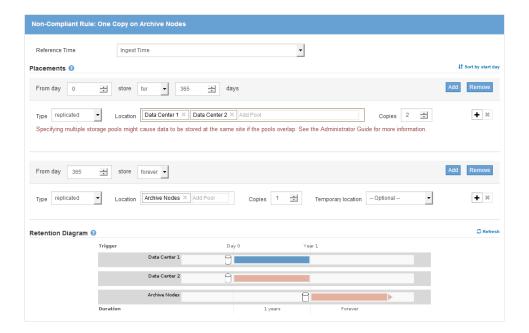
This example ILM rule stores two replicated object copies on Storage Nodes for one year. After that, it moves one object copy to Archive Nodes and stores this copy forever. Because this rule uses

Archive Nodes and because it only saves only one object copy from day 365 to forever, it is not compliant and will not apply to the objects in compliant S3 buckets.

Rule definition	Example value
Rule Name	Non-Compliant Rule: One Copy on Archive Nodes
Tenant Accounts	Not specified
Bucket Name	Not specified, but will only apply to non-compliant buckets



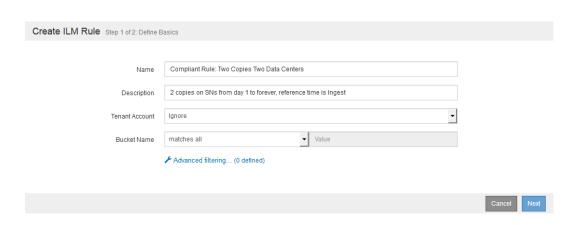
Rule definition	Example value
Reference Time	Ingest Time
Placements	<ul> <li>On Day 0, keep two replicated copies on Storage Nodes in Data Center 1 and Data Center 2 for 365 days</li> <li>On Day 365, keep one replicated copy on Archive Nodes forever; temporary copies in Data Center 1</li> </ul>



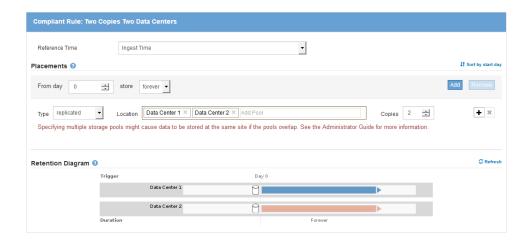
# ILM rule 3 for compliance example: Default rule

This example ILM rule copies object data to storage pools in two data centers. This rule is designed to be the default rule in the ILM policy. It satisfies the requirements of compliant S3 buckets: 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	Compliant Rule: Two Copies Two Data Centers
Tenant Account	Not specified
Bucket Name	Not specified
Advanced filtering	Not specified



Rule definition	Example value
Reference Time	Ingest Time
Placements	From Day 0 to forever, keep two replicated copies—one on Storage Nodes in Data Center 1 and one on Storage Nodes in Data Center 2.



# **ILM** policy for compliance example

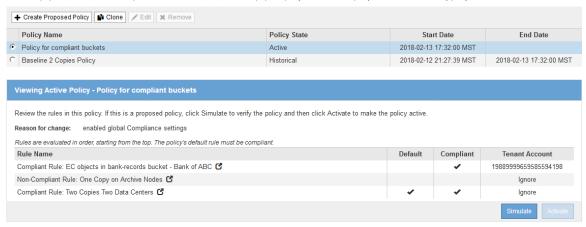
To create an ILM policy that will effectively protect objects in compliant S3 buckets as well as objects in non-compliant buckets, you must select ILM rules that satisfy the storage requirements for both types of objects. Then, you must simulate and activate the proposed policy.

# Adding rules to the policy

In this example, the active ILM policy includes three ILM rules, in the following order:

- 1. A compliant rule that creates erasure-coded copies of the objects in a specific compliant S3 bucket. 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 non-compliant buckets because it stores only one object copy forever and it uses Archive Nodes.
- A compliant rule that creates two replicated object copies on Storage Nodes from day 0 to forever.

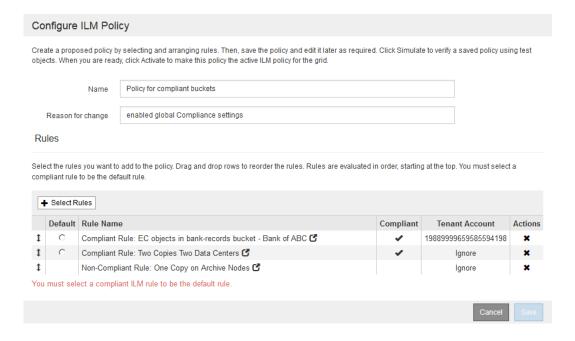
Review the proposed, active, and historical policies. You can create, edit, or delete a proposed policy; clone the active policy; or view the details for any policy.



# Selecting a default rule

If the global Compliance setting is enabled, the default rule in the active or any proposed ILM policy must be compliant. In the example, the default rule is the second compliant rule. This rule applies to any object in any compliant or non-compliant bucket that was not matched by the first two rules.

**Note:** When the global Compliance setting is enabled, you might see an error message when you initially select the rules for a proposed ILM policy. The message indicates you must select a compliant ILM rule to be the default rule. Select the **Default** radio button for the compliant rule that you want to be the default, and drag that rule to the appropriate position in the list.



# Simulating the proposed policy

After you have added rules in your proposed policy, arranged them, and chosen a default compliant rule, you should simulate the policy by testing objects from both compliant and non-compliant buckets. For example, if you simulated the example policy, you would expect test objects to be evaluated as follows:

- A test object in the bucket bank-records for the Bank of ABC tenant would be matched by the EC objects compliant rule.
- A test object in any non-compliant bucket for any tenant account would be matched by the noncompliant rule.
- A test object in a compliant bucket named customer-records for Bank of ABC or any other tenant would be matched by the default rule. This is because the bucket name does not match bank-records and the non-compliant rule does not apply to objects in compliant buckets.

## Activating the policy

When you are completely satisfied that the new policy protects object data as expected, you can activate it.

#### Related concepts

Creating, simulating, and activating an ILM policy on page 192

# Resolving consistency errors when updating a bucket's 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 a bucket's compliance configuration.

#### About this task

Tenant users who have created compliant buckets can configure the bucket's compliance settings from the Tenant Manager, Tenant Management API, or the S3 REST API. For example, a tenant user might need to increase the retention period or put a bucket under a legal hold.

When a tenant user updates the compliance settings for an S3 bucket, StorageGRID attempts to immediately update the bucket's metadata across the grid. If the system is unable to update the bucket's 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.

# 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 compliance changes are consistently applied across the grid.
- 3. Once the underlying issue has been resolved, remind the tenant user to retry their compliance configuration changes.

# **Related information**

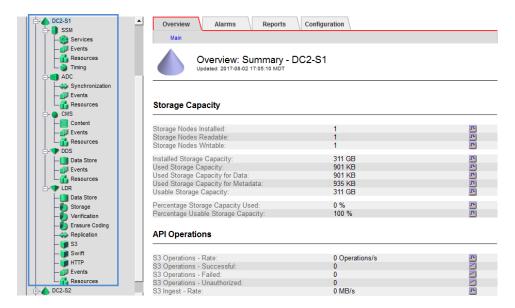
Using tenant accounts Implementing S3 client applications Recovery and maintenance

# Managing disk storage

Storage Nodes provide disk storage capacity and services.

# What a Storage Node is

A Storage Node includes the services and processes required to store, move, verify, and retrieve object data and metadata on disk.

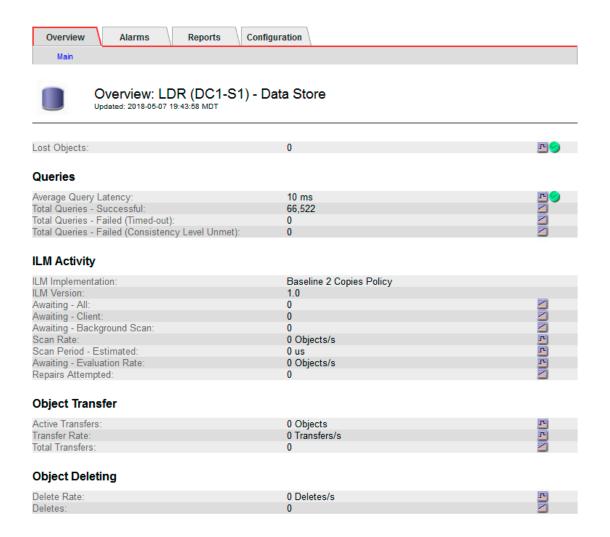


# What the LDR service is

Hosted by a 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 deleting
- Object data storage
- Object data transfers from another LDR service (Storage Node)
- Data storage management
- Protocol interfaces (S3 and Swift)



## 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.

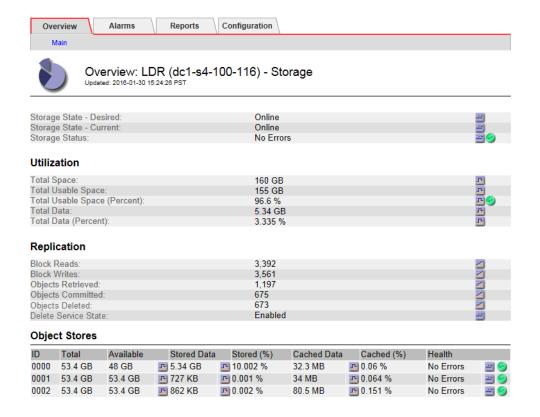
You can also view the total number of queries that failed because of consistency failures. Consistency level failures result from an insufficient number of available metadata stores at the time a query is performed through the specific LDR service.

### **ILM Activity**

Information Lifecycle Management (ILM) metrics allow you to monitor the rate at which objects are evaluated for ILM implementation. You can view some of these metrics on the Dashboard.

# **Object stores**

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



The object stores in a Storage Node are identified by a hexadecimal number from 0000 to 000F, which is known as the volume ID. By default, 3 TB of 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.

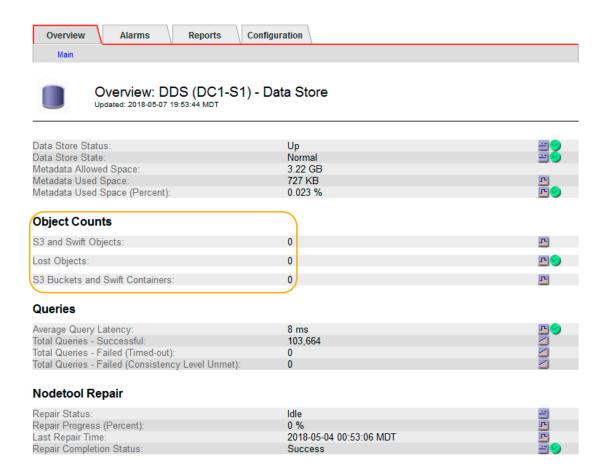
### What the DDS service is

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

The DDS 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.

# **Object counts**

The DDS service lists 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).



Because object metadata synchronization occurs over time, object count attributes (see **DDS** > **Data Store** > **Overview** > **Main**) can differ between DDS services. Eventually, all metadata stores will synchronize and counts should become the same.

#### 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.

### 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.

# **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 DDS service.

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

# What the nodetool repair operation is

Periodically, the StorageGRID system runs the nodetool repair operation on Storage Nodes checking for and repairing metadata replication inconsistencies that may occur over time.

Nodetool repair is run every 12 to 14 days at random times on different Storage Nodes, so that it does not run on every Storage Node at the same time. The nodetool repair operation is a seamless activity that occurs in the background of normal system operations.

## 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.

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.

# **Managing Storage Nodes**

Managing Storage Nodes entails monitoring the amount of usable space on each node, using watermark settings, and applying Storage Node configuration settings.

# What watermarks are

You use watermark settings to globally manage a Storage Node's usable storage space. Watermark settings trigger alarms that assist you in monitoring available storage and determining when to add Storage Nodes.

A Storage Node becomes read-only when all of a Storage Node's object stores reach the Storage Volume Hard Read-Only Watermark. If available space falls below this configured watermark

amount, a Notice alarm is triggered for the Storage Status (SSTS) attribute. The alarm notifies you to manage storage proactively so that you add capacity only when necessary.

You can view the StorageGRID system's current watermark values at any time. Go to Configuration > Storage Options > Overview.



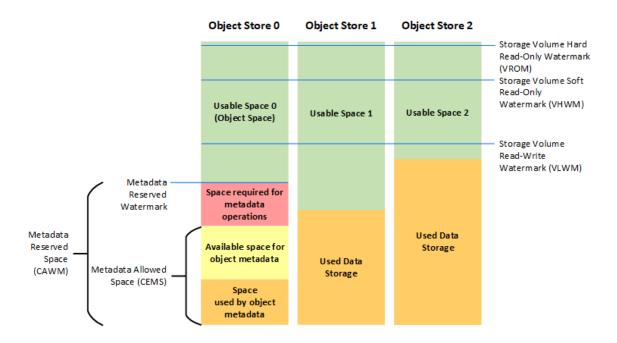
### **Object Segmentation**

Description	Settings
Segmentation	Enabled
Maximum Segment Size	1 GB

#### **Storage Watermarks**

Storage Volume Soft Read-Only Watermark 10 GB Storage Volume Hard Read-Only Watermark 5 GB	Description	Settings
Storage Volume Hard Read-Only Watermark 5 GB	Storage Volume Read-Write Watermark	30 GB
,	Storage Volume Soft Read-Only Watermark	10 GB
Notadata Reserved Space 3 000 GB	Storage Volume Hard Read-Only Watermark	5 GB
netauata reserveu opace 5,000 OD	Metadata Reserved Space	3,000 GB

#### Watermark related attributes



Watermar k Name	Default Setting	Code	Description
Storage Volume Hard Read- Only Watermark	5 GB	VROM	Indicates when a Storage Node transitions to hard read- only mode. Hard read-only mode means that the Storage Node is read-only and no longer accepts write requests.  The Storage Volume Hard Read-Only Watermark value is calculated against the Total Space value for the Storage Node, but measured against the Total Usable Space value for the Storage Node. When the value of Total Usable Space falls below the value of Storage Volume Hard Read-Only Watermark, the Storage Node transitions to hard read-only mode.  The Storage Volume Hard Read-Only Watermark value must be less than value for The Storage Volume Soft Read-Only Watermark.
Storage Volume Soft Read- Only Watermark	10 GB	VHWM	Indicates when a Storage Node transitions to soft readonly 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.  The Storage Volume Soft Read-Only Watermark value is calculated against the Total Space value for the Storage Node, but measured against the Total Usable Space value for the Storage Node. When the value of Total Usable Space falls below the value of Storage Volume Soft Read-Only Watermark, the Storage Node transitions to soft read-only mode:  The Storage State – Current (SSCR) changes to Read-Only. If Storage State – Desired is set to Online, Storage Status (SSTS) changes to Insufficient Free Space and a Notice alarm is triggered.  An alarm for Total Usable Space (Percent) (SAVP) can be triggered, depending on the relationship between the watermark setting (in bytes) and the alarm settings (in percent).  The Storage Node is writable again if Total Usable Space (STAS) becomes greater than Storage Volume Soft Read-Only Watermark.
Storage Volume Read-Write Watermark	30 GB	VLWM	Indicates when a Storage Node that has transitioned to read-only mode is allowed to become read-write again, specifically when one of the constituent volumes determines that its available space exceeds this setting.  The Storage Volume Read-Write Watermark value must be greater than the value for Storage Volume Soft Read-Only Watermark.  The value for a volume's available space is located on the Support > Grid Topology > storage node > LDR > Storage > Overview page.

Watermar k Name	Default Setting	Code	Description
Metadata Reserved Space	3 TB	CAWM	The amount of space reserved on storage volume 0 for metadata storage. Metadata Reserved Space is subdivided into space for object metadata (Metadata Allowed Space, or CEMS) and space for essential database operations, such as compaction and repair.  Note: If the storage capacity of volume 0 is less than 500 GB (non-production use only), 10% of the storage volume's capacity is reserved for metadata.

# **Related concepts**

Managing full Storage Nodes on page 246

# **Storage Node configuration settings**

Depending on your requirements, there are several configuration settings for Storage Nodes that you can apply.

This table summarizes Storage Node configuration settings.

Service/ Component	Attribute Name	Code	Description
LDR	HTTP State	HSTE	HSTE is related to the HTTP protocol for all LDR traffic, including S3, Swift, and other internal StorageGRID traffic. Set the LDR service to one of the following choices:  • Offline: No operations are allowed, and any client application that attempts to open an HTTP session to the LDR service receives an error message. Active sessions are gracefully closed.  • Online: Operation continues normally
	Auto-Start HTTP	HTAS	HTAS is related to the HTTP protocol for all LDR traffic, including S3, Swift, and other internal StorageGRID traffic.  Enable the HTTP component when the LDR service is restarted. If not selected, the HTTP interface remains Offline until explicitly enabled.  If Auto-Start HTTP is selected, the state of the system on restart depends on the state of the LDR > Storage component. If the LDR > Storage component is Readonly on restart, the HTTP interface is also Read-only. If the LDR > Storage component is Online, then HTTP is also Online. Otherwise, the HTTP interface remains in the Offline state.
LDR > Data Store	Reset Lost Objects Count	RCOR	Reset to zero the counter for the number of lost objects on this service.

Service/ Component	Attribute Name	Code	Description
LDR > Storage	Storage State  – Desired	SSDS	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.  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.
			The attribute can take one of the following values:
			Offline: When the desired state is Offline, the LDR service takes the LDR > Storage component offline.
			Read-only: When the desired state is Read-only, the LDR service moves the storage state to read-only and stops accepting new content. Note that content might continue to be saved to the Storage Node for a short time until open sessions are closed.
			Online: Leave the value at Online during normal system operations. The Storage State – Current of the storage component will be dynamically set by the service based on the condition of the LDR service, such as the amount of available object storage space. If space is low, the component becomes Read-only.
	Health Check Timeout	SHCT	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	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 priority rate at which background verification takes place. See information on configuring the background verification rate.
	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 information about troubleshooting.

Service/ Component	Attribute Name	Code	Description
LDR > Erasure Coding	Reset Writes Failure Count	RSWF	Reset to zero the counter for write failures of erasure-coded object data to the Storage Node.
	Reset Reads Failure Count	RSRF	Reset to zero the counter for read failures of erasure-coded object data from the Storage Node.
	Reset Deletes Failure Count	RSDF	Reset to zero the counter for delete failures of erasure-coded object data from the Storage Node.
	Reset Corrupt Copies Detected Count	RSCC	Reset to zero the counter for the number of corrupt copies of erasure-coded object data on the Storage Node.
	Reset Corrupt Fragments Detected Count	RSCD	Reset to zero the counter for corrupt fragments of erasure-coded object data on the Storage Node.
	Reset Missing Fragments Detected Count	RSMD	Reset to zero the counter for missing fragments of erasure-coded object data on the Storage Node. Use only after foreground verification completes.
LDR > Replication	Reset Inbound Replication Failure Count	RICR	Reset to zero 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 to zero the counter for outbound replication failures. This can be used to clear the RORF (Outbound Replications – Failed) alarm.
	Disable DSIR Inbound Replication		Select to disable inbound replication as part of a maintenance or testing procedure. Leave unchecked during normal operation.
			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.
	Disable Outbound Replication	DSOR	Select 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, 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.
LDR > HTTP	Reset HTTP Counts	LHAC	Reset to zero the counter for all HTTP transactions.

#### Related tasks

Configuring the background verification rate on page 253

#### Related information

Troubleshooting StorageGRID

# **Managing full Storage Nodes**

As Storage Nodes reach capacity, you must be expand the StorageGRID system through the addition of new storage. There are two options available when considering how to increase storage capacity: adding Storage Nodes and adding storage volumes.

# **Adding Storage Nodes**

You can increase storage capacity by adding Storage Nodes. Careful consideration of currently active ILM rules and capacity requirement must be taken when adding storage. See the instructions for expanding a StorageGRID grid.

## Adding storage volumes

Each Storage Node supports a maximum of 16 storage volumes. If a Storage Node includes fewer than 16 storage volumes, you can increase its capacity by adding storage volumes up to the maximum of 16.

#### **Related information**

Expanding a StorageGRID system

# **Setting Grid Options**

Grid Options allow you to manage the settings for stored objects, including stored object encryption, stored object hashing, and stored object compression. You can also set the Prevent Client Modify and the HTTP options.

#### Choices

- Configuring stored object encryption on page 246
- Configuring stored object hashing on page 247
- Configuring stored object compression on page 248
- Enabling Prevent Client Modify on page 249
- Enabling HTTP for client communications on page 250

# Configuring stored object encryption

Stored object encryption enables the encryption of stored object data so that if an object store is compromised data cannot be retrieved in a readable form. By default, objects are not encrypted.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

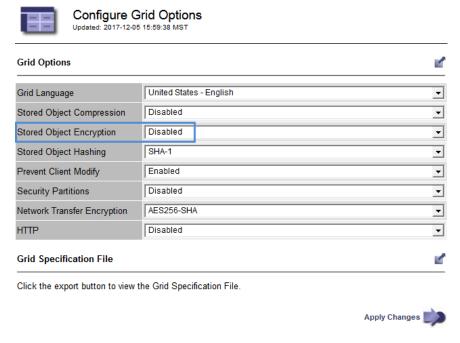
#### About this task

Objects can also be encrypted using the AES-128 or AES-256 encryption algorithm. Stored object encryption enables the encryption of all object data ingested through S3 or Swift. If disabled, currently encrypted objects remain encrypted. For S3 objects, the Stored Object Encryption setting can be overridden by the x-amz-server-side-encryption header. If you use the x-amzserver-side-encryption header, you must specify the AES-256 encryption algorithm in the request.

**Note:** 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 > Grid Options.
- **2.** From the Grid Options menu, select **Configuration**.
- 3. Change Stored Object Encryption to Disabled, AES-256, or AES-128.



4. Click Apply Changes.

# Configuring stored object hashing

The Stored Object Hashing option specifies the hashing algorithm used by the LDR service to hash data when new content is stored. These hashes are verified during retrieval and verification to protect the integrity of data.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

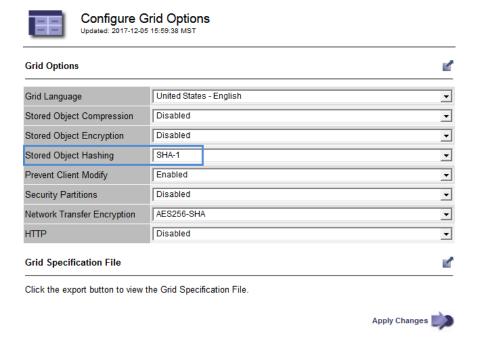
#### About this task

By default, object data is hashed using the SHA-1 algorithm. Object data can also be hashed using the SHA-256 algorithm.

**Note:** 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 > Grid Options.
- 2. From the Grid Options menu, select Configuration.
- 3. Change Stored Object Hashing to SHA-256 or SHA-1.



4. Click Apply Changes.

# Configuring stored object compression

You can use the Stored Object Compression grid option to reduce the size of objects stored in StorageGRID, so that objects consume less storage.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

The Stored Object Compression grid option is disabled by default. If you enable this option, StorageGRID attempts to compress each object when saving it, using lossless compression.

**Note:** 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 setting, be aware of the following:

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
Stored Object Compression will not further reduce an object's size.

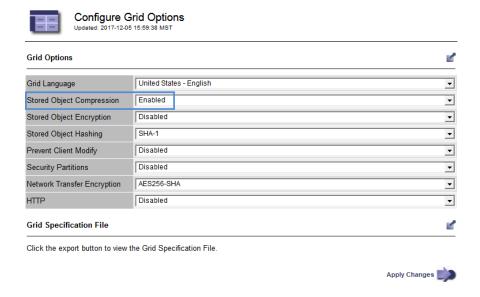
If the Stored Object Compression grid option is enabled, S3 and Swift client applications should avoid performing GET Object operations that specify a range of bytes be returned. These "range read" operations are inefficient because StorageGRID must effectively uncompress the objects to access the requested bytes. GET Object operations that request a small range of bytes from a very large object are especially inefficient; for example, it is inefficient to read a 10 MB range from a 50 GB compressed object.

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

Note: 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 > Grid Options.
- From the Grid Options menu, select **Configuration**.
- Change Stored Object Compression to **Enabled**.



4. Click Apply Changes.

# **Enabling Prevent Client Modify**

You can set the Prevent Client Modify option to **Enabled** to override the permissions defined for HTTP profiles and to deny specific HTTP client operations.

# Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

# About this task

When the Prevent Client Modify option is enabled, 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
 Note: This setting does not apply to buckets with versioning enabled. Versioning already prevents modifications to object data, user-defined metadata, and object tagging.

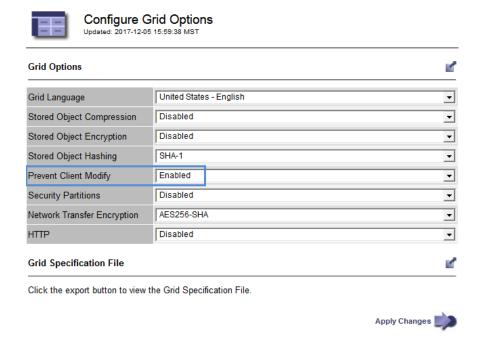
### Swift REST API

- Delete Container requests
- Requests to modify any existing object. For example, the following operations are denied: Put Overwrite, Delete, Metadata Update, and so on.

Prevent Client Modify is a system wide setting.

#### **Steps**

- 1. Select Configuration > Grid Options.
- **2.** From the Grid Options menu, select **Configuration**.
- Change Prevent Client Modify to Enabled.



4. Click Apply Changes.

# **Enabling HTTP for client communications**

By default, client applications use HTTPS for all connections with the Storage Nodes and API Gateway Nodes. However, you can set the HTTP option to **Enabled** if you want to use HTTP communications between S3 and Swift clients and StorageGRID in addition to HTTPS communications. For example, you might use HTTP when testing a non-production grid.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

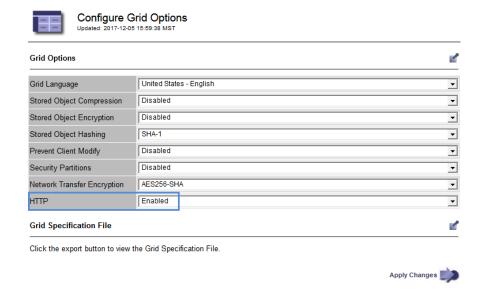
#### About this task

When the HTTP option is enabled, either HTTP or HTTPS can be used for communications between S3 and Swift clients and StorageGRID. HTTP communications require the use of different ports between the S3/Swift clients and the API Gateway Nodes and Storage Nodes. See the instructions for implementing S3 and Swift client applications for more information.

Attention: HTTP mode is intended for use in testing and debugging environments. Be careful when enabling HTTP for a production grid since requests will be sent unencrypted.

## **Steps**

- 1. Select Configuration > Grid Options.
- From the Grid Options menu, select Configuration.
- Change HTTP to **Enabled**.



4. Click Apply Changes.

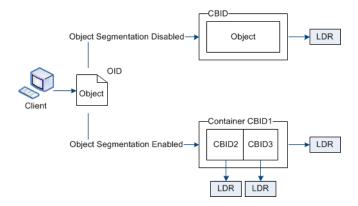
## **Related information**

Implementing S3 client applications Implementing Swift client applications

# 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.

# 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.

## Related concepts

What background verification is on page 253 What foreground verification is on page 255

# 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:

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.

**Note:** 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.

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 be quarantined.

If background verification cannot replace a corrupted object because it cannot locate another copy, the LOST (Lost Objects) alarm is triggered.

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.

# Configuring the background verification rate

You can change the rate at which background verification checks replicated object data on a Storage Node if you have concerns about data integrity.

### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### About this task

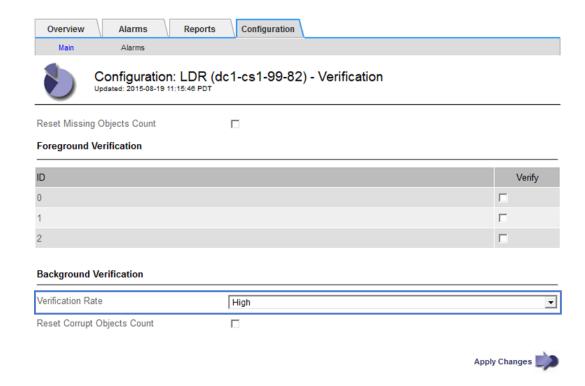
You can change the Verification Rate for background verification on a Storage Node:

- Adaptive: Default setting. The task is designed to verify at a maximum of 4 MB/s or 10 objects/s (whichever is exceeded first).
- High: Storage verification proceeds quickly, at a rate that can slow ordinary system activities.

Use the High verification rate only when you suspect that a hardware or software fault might have corrupted object data. After the High priority background verification completes, the Verification Rate automatically resets to Adaptive.

## **Steps**

- 1. Select Support > Grid Topology.
- 2. Select Storage Node > LDR > Verification.
- 3. Click Configuration > Main.
- 4. Go to LDR > Verification > Configuration > Main.
- Under Background Verification, select Verification Rate > High or Verification Rate > Adaptive.



**Note:** Setting the Verification Rate to High triggers a Notice level alarm for VPRI (Verification Rate).

- 6. Click Apply Changes.
- 7. Monitor the results of background verification.
  - a. Go to **LDR** > **Verification** > **Overview** > **Main** and monitor the attribute Corrupt Objects Detected (OCOR).

If background verification finds corrupt replicated object data, the attribute Corrupt Objects Detected is incremented. The LDR service recovers by quarantining the corrupt object data and sending a message to the DDS service to create a new copy of the object data. The new copy can be made anywhere in the StorageGRID system that satisfies the active ILM policy.

b. Go to **LDR > Erasure Coding > Overview > Main** and monitor the attribute Corrupt Fragments Detected (ECCD).

If background verification finds corrupt fragments of erasure-coded object data, the attribute Corrupt Fragments Detected is incremented. The LDR service recovers by rebuilding the corrupt fragment in place on the same Storage Node.

8. If corrupt replicated object data is found, contact technical support to clear the quarantined copies from the StorageGRID system and determine the root cause of the corruption.

# 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. If foreground verification finds issues when verifying erasure-coded objects, you must perform a recovery procedure for any affected storage volumes.

Foreground verification checks both replicated object data and erasure-coded object data. If foreground verification finds a missing copy of object data, it automatically attempts to replace it.

If a copy of replicated object data is found to be missing, the StorageGRID system automatically attempts to replace it 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 object 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.

If a fragment of a copy of an erasure-coded object is found to be missing, the StorageGRID system 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, 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 no other copies of a missing replicated object or a corrupted erasure-coded object can be found in the grid, the LOST (Lost Objects) alarm is triggered.

If foreground verification identifies an issue with erasure coding on a storage volume, the foreground verification grid task pauses with an error message that identifies the volume that is affected.

# 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.

#### Before you begin

- You have ensured that the following grid tasks are not running:
  - o Grid Expansion: Add Server (GEXP), when adding a Storage Node
  - Storage Node Decommissioning (LDCM) on the same Storage Node

If these grid tasks are running, wait for them to complete or release their lock.

- You have ensured that the storage is online. (Select **Support > 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 drive

Foreground 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.

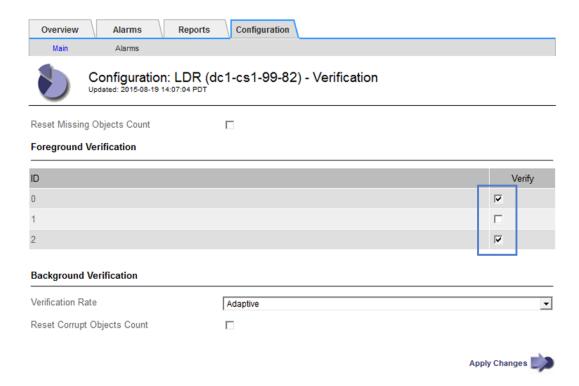
You can configure foreground verification to check all of a Storage Node's object stores or only specific object stores.

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.

- 1. Select Storage Node > LDR > Verification.
- 2. Click Configuration > Main.
- **3.** Under **Foreground Verification**, select the check box for each storage volume ID you want to verify.



## 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.

If the count for the attribute 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 count for the attribute 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 > 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

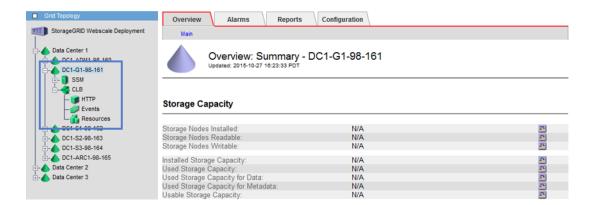
If you still have concerns about data integrity, go to **LDR > Verification > Configuration > Main** and increase the background Verification Rate. Background verification checks the correctness of all stored object data and repairs any issues that it finds. Finding and repairing potential issues as quickly as possible reduces the risk of data loss.

#### Related information

Recovery and maintenance

# How load balancing works

To balance ingest and retrieval workloads, optionally deploy the StorageGRID system with API Gateway Nodes, or integrate a third-party load balancer.



#### **API Gateway Node**

The API Gateway Node provides layers 3 and 4 load balancing to the StorageGRID system and distributes incoming network connections from client applications to Storage Nodes.

The Connection Load Balancer (CLB) service distributes incoming TCP network connections 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.

## **Related concepts**

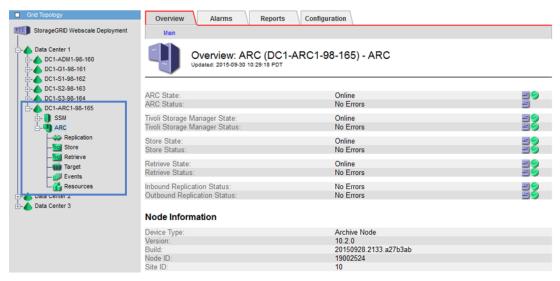
What link costs are on page 279

# Managing archival storage

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.

# 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.



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.

**Note:** 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 the cloud through the S3 API or 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.

# About supported archive targets

When you configure the 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 TSM middleware.

## Archiving to the cloud through the S3 API

You can configure an Archive Node to target any external archival storage system that is capable of interfacing with the StorageGRID system through the S3 API.

The Archive Node's ARC service can be configured to connect directly to Amazon Web Services (AWS) or to any other system that can interface to the StorageGRID system through the S3 API; for example, another instance of the StorageGRID system.

#### Archiving to tape through TSM middleware

You can configure an Archive Node to target a Tivoli Storage Manager (TSM) server which provides a logical interface for storing and retrieving object data to random or sequential access storage devices, including tape libraries.

The Archive Nodes's ARC service acts as a client to the TSM server, using Tivoli Storage Manager as middleware for communicating with the archival storage system.

#### **Tivoli Storage Manager 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).

# Managing connections to archival storage

You can configure an Archive Node to connect to an external archival storage system through either the S3 API or TSM middleware.

Once the type of archival target is configured for an Archive Node, the target type cannot be changed.

# Configuring connection settings for S3 API

You must configure a number of settings before the Archive Node can communicate with an external archival storage system that connects to the StorageGRID system through the S3 API.

**Attention:** The use of the **Cloud Tiering - Simple Storage Service** (**S3**) option, which moves objects from an Archive Node to an external archival storage system, 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 "Migrating objects from Cloud Tiering - S3 to a Cloud Storage Pool."

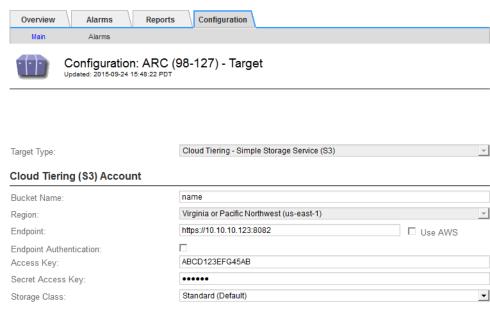
#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You need to create 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 Simple Storage Service (S3) is used as the external archival storage system.

#### 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 external archival storage system.

- 1. Select Support > Grid Topology.
- 2. Select Archive Node > ARC > Target.
- 3. Click Configuration > Main.





4. Select Cloud Tiering - Simple Storage Service (S3) from the Target Type drop-down list.

**Note:** Configuration settings are unavailable until you select a Target Type.

**5.** 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. Clear to disable endpoint SSL certificate and
  host name verification for the targeted external archival storage system. Only clear the
  checkbox if the network to the external archival storage system is trusted. If another instance
  of a StorageGRID system is the target archival storage device and the system is configured
  with publicly signed certificates, you do not need to clear the checkbox.
- Storage Class: Select Standard, the default value, for regular storage, or Reduced Redundancy, which provides lower cost storage with less reliability for objects that can be easily recreated.
   If the targeted archival storage system is another instance of the StorageGRID system, Storage Class controls the target system's dual-commit behavior.

## 6. Click Apply Changes.

The specified configuration settings are validated and applied to your StorageGRID system. Once configured, the target cannot be changed.

#### Related tasks

Migrating objects from Cloud Tiering - S3 to a Cloud Storage Pool on page 262

# 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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You have already stored objects in the S3 bucket configured for Cloud Tiering.

**Note:** 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.

Note: 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 tasks

Configuring connection settings for S3 API on page 260 Creating a Cloud Storage Pool on page 170

Cloning an ILM rule on page 208

Creating a proposed ILM policy on page 192

Simulating an ILM policy on page 196

Activating the ILM policy on page 203

# 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.

## Before you begin

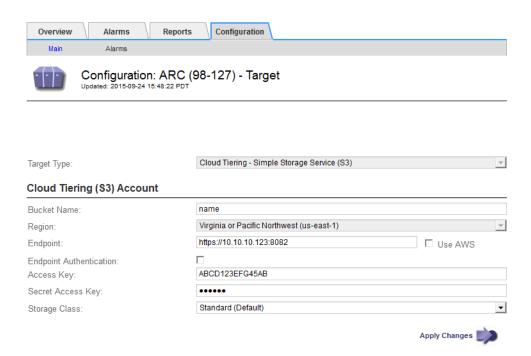
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### 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 > Grid Topology.
- 2. Select Archive Node > ARC > Target.
- 3. Click Configuration > Main.



**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.

Note: 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.

## Before you begin

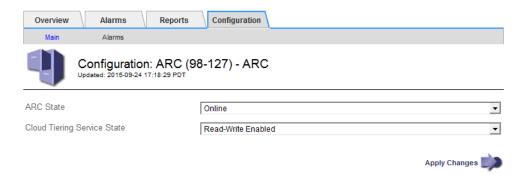
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- The Archive Node must be configured.

#### 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 > Grid Topology.
- Select Archive Node > ARC.
- 3. Click Configuration > Main.



- 4. Select a Cloud Tiering Service State.
- 5. Click Apply Changes.

# Configuring connections to Tivoli Storage Manager middleware

Before the Archive Node can communicate with Tivoli Storage Manager (TSM) middleware, you must configure a number of settings.

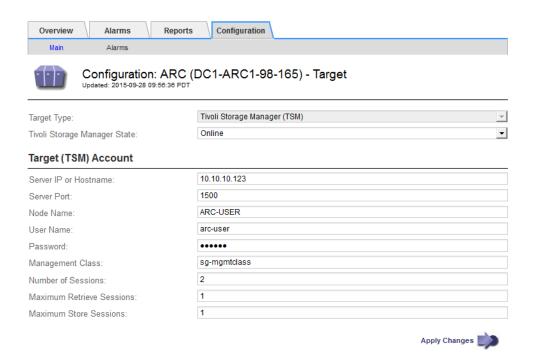
## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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.

- 1. Select Support > Grid Topology.
- 2. Select Archive Node > ARC > Target.
- 3. Click Configuration > Main.



- **4.** From the **Target Type** drop-down list, select **Tivoli Storage Manager** (**TSM**).
- 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.

# Managing Archive Nodes

You can configure an Archive Node to optimize Tivoli Storage Manager 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.

#### Choices

- Optimizing Archive Node for TSM middleware sessions on page 267
- Managing an Archive Node when TSM server reaches capacity on page 268
- Configuring Archive Node replication on page 269
- Configuring retrieve settings on page 270
- Configuring the archive store on page 271
- Setting Custom alarms for the Archive Node on page 273

# Optimizing Archive Node for TSM middleware sessions

You can optimize the performance of an Archive Node that connects to an external archival storage system through the S3 API by configuring the Archive Node's sessions.

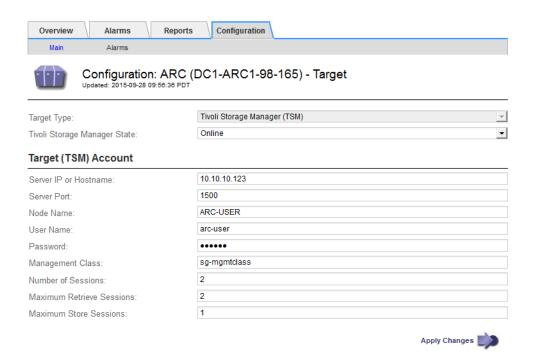
### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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.

- 1. Select Support > Grid Topology.
- 2. Select Archive Node > ARC > Target.
- 3. Click Configuration > Main.
- 4. Change Maximum Retrieve Sessions to be the same as Number of Sessions.



5. Click Apply Changes.

# 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.

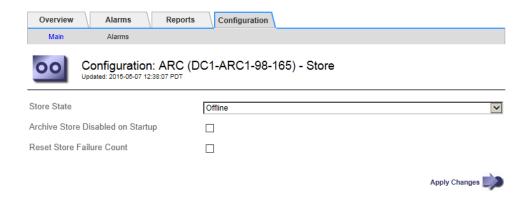
#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You 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.

- 1. Select Support > Grid Topology.
- 2. Select Archive Node > ARC > Store.
- 3. Click Configuration > Main.



- 4. Change Archive 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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### Steps

- 1. Select Support > Grid Topology.
- 2. Select Archive Node > ARC > Target.
- 3. Click Configuration > Main.
- 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.

Note: Changing Maximum Store Sessions to 0 is not necessary if the Archive Node is Readonly. Store sessions will not be created.

6. Click Apply Changes.

## Configuring Archive Node replication

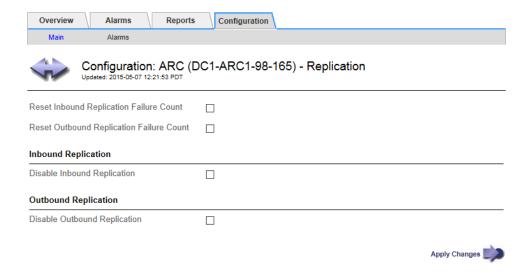
You can configure the replication settings for an Archive Node and disable inbound and outbound replication, or reset the failure counts being tracked for the associated alarms.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### **Steps**

- 1. Select Support > Grid Topology.
- 2. Select Archive Node > ARC > Replication.
- 3. Click Configuration > Main.



- **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 checkbox 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.

# Configuring 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.

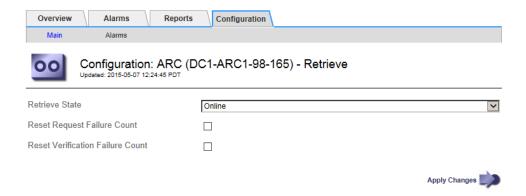
#### Before you begin

You must be signed in to the Grid Manager using a supported browser.

You have specific access permissions.

#### **Steps**

- 1. Select Support > Grid Topology.
- 2. Select Archive Node > ARC > Retrieve.
- 3. Click Configuration > Main.



- **4.** Modify the following settings, as necessary:
  - Archive 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 checkbox to reset the counter for request failures. This can be used to clear the ARRF (Request Failures) alarm.
  - Reset Verification Failure Count: Select the checkbox 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 the archive store

You can configure store setting for an Archive Node.

#### About this task

Store settings differ based on the configured target type for the Archive Node.

#### Related tasks

Managing an Archive Node when TSM server reaches capacity on page 268

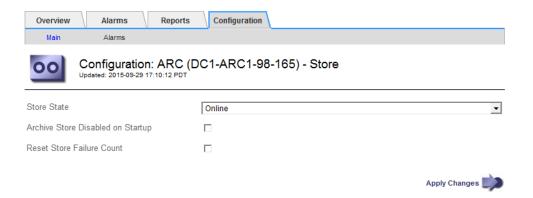
# Configuring the archive store for TSM middleware connection

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.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

- 1. Select Support > Grid Topology.
- 2. Select Archive Node > ARC > Store.
- 3. Click Configuration > Main.



- **4.** Modify the following settings, as necessary:
  - Archive 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.

## Configuring store settings for S3 API connection

If your Archive Node connects to an archival storage system through the S3 API, you can reset the Store Failures count, which can be used to clear the ARVF (Store Failures) alarm.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

#### **Steps**

- 1. Select Support > Grid Topology.
- 2. Select Archive Node > ARC > Store.
- 3. Click Configuration > Main.



- 4. Select Reset Store Failure Count.
- 5. Click Apply Changes.

The Store Failures attribute resets to zero.

# 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.

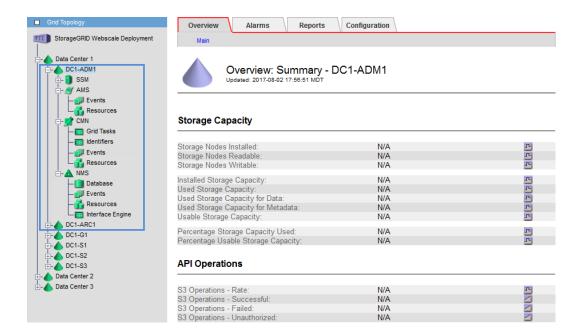
#### Related tasks

Creating custom service or component alarms on page 127

# What an Admin Node is

You perform most day-to-day activities using the Grid Manager, which resides on Admin Nodes. Admin Nodes provide services for the web interface, system configuration, and audit logs. Each site in a StorageGRID deployment can have one or more Admin Nodes.

Admin Nodes use the AMS service, the CMN service, and the NMS service.



#### 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

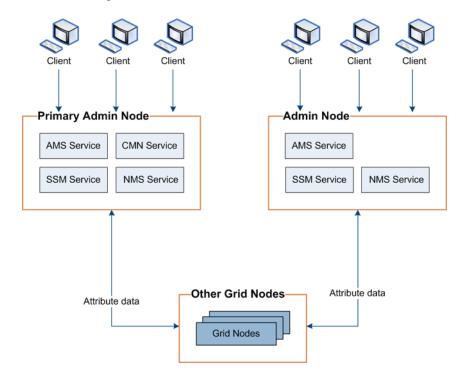
The Network Management System (NMS) service powers the monitoring, reporting, and configuration options that are displayed through the Grid Manager, the StorageGRID system's browser-based interface.

#### Related tasks

Verifying an ILM policy with object metadata lookup on page 204

A StorageGRID system can include multiple Admin Nodes. This provides you with the redundancy of multiple IP addresses from which you can to sign in to StorageGRID system and perform various monitoring and configuration procedures.

Having multiple Admin Nodes provides you with the capability to continuously monitor and configure your StorageGRID system in the event that an Admin Node fails. If an Admin Node becomes unavailable, web clients can reconnect to any other available Admin Node and continue to view and configure the system. Meanwhile, attribute processing continues, alarms are still triggered, and related notifications sent. However, multiple Admin Nodes does not provide failover protection except for notifications and AutoSupport messages. Alarm acknowledgments made from one Admin Node are not copied to other Admin Nodes.



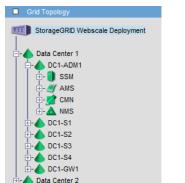
#### Related concepts

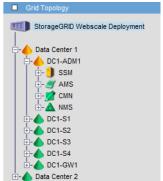
Alarm acknowledgments on page 275

# Alarm acknowledgments

Alarm acknowledgments made from one Admin Node are not copied to any other Admin Node. Because acknowledgments are not copied to other Admin Nodes, it is possible that the Grid Topology tree will 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.

# **Email notifications and AutoSupport messages**

In a multi-site StorageGRID system, one Admin Node is configured as the preferred sender of notifications and AutoSupport messages. This preferred sender can be any Admin Node. All other Admin Nodes become "standby" senders.

Under normal system operations, only the preferred sender sends notifications and AutoSupport messages. The standby sender monitors the preferred sender and if it detects a problem, the standby sender switches to online status and assumes the task of sending notifications and AutoSupport messages.

#### Preferred and standby senders

There are two scenarios in which both the preferred sender and the standby sender can send notifications and AutoSupport messages:

- It is possible that while the StorageGRID system is running in this "switch-over" scenario, where the standby sender assumes the task of sending notifications and AutoSupport messages, the preferred sender will maintain the ability to send notifications and AutoSupport messages. If this occurs, duplicate notifications and AutoSupport messages are sent: one from the preferred sender and one from the standby sender. When the Admin Node configured as the standby sender no longer detects errors on the preferred sender, it switches to "standby" status and stops sending notifications and AutoSupport messages. Notifications and AutoSupport messages are once again sent only by the preferred sender.
- If the standby sender cannot detect the preferred sender, the standby sender switches to online and sends notifications and AutoSupport messages. In this scenario, the preferred sender and standby senders are "islanded" from each other. Each sender (Admin Node) can be operating and monitoring the system normally, but because the standby sender cannot detect the other Admin Node of the preferred sender, both the preferred sender and the standby sender send notifications and AutoSupport messages.

When sending a test email, all NMS services send a test email.

#### Related concepts

About alarms and email notifications on page 112 What AutoSupport is on page 135

#### Related tasks

Selecting a preferred sender on page 120

# Managing networking

Because the topology of your StorageGRID system is that of a group of interconnected servers, over time as your system changes and grows you may be required to perform various updates to the system's networking.

You can change the configuration of the Grid, Client, or Admin Networks, or you can add new Client and Admin Networks. You can also update external NTP source IP addresses and DNS IP addresses at any time.

**Note:** To use the Grid Network editor to modify or add a network for a grid node, see the recovery and maintenance instructions. For more information about network topology, see *Grid primer*.

#### **Grid Network**

Required. The Grid Network is the communication link between grid nodes. All hosts on the Grid Network must be able to talk to all other hosts. This network is used for all internal StorageGRID system communications.

#### **Admin Network**

Optional. The Admin Network allows for restricted access to the StorageGRID system for maintenance and administration.

#### Client Network

Optional. The Client Network can communicate with any subnet reachable through the local gateway.

#### Guidelines

- A StorageGRID grid node requires a dedicated network interface, IP address, subnet mask, and gateway for each network it is assigned to.
- A grid node is not permitted to have more than one interface on a network.
- A single gateway, per network, per grid node is supported, and it has to 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
  - Grid Network or eth0: hic2 and hic4 (10-GbE network ports)
  - Admin Network or eth1: mtc1 (the leftmost 1-GbE port)
  - Client Network or eth2: hic1 and hic3 (10-GbE network ports)
- 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 VM deployment to allow access to the installation UI before the grid is fully installed.

#### Related information

Recovery and maintenance
Grid primer
SG6000 appliance installation and maintenance
SG5700 appliance installation and maintenance
SG5600 appliance installation and maintenance

# **Viewing IP addresses**

You can view the IP address for each grid node that makes up 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.

#### Before you begin

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.

ode Information	0		
Name	SGA-X33-005-024		
Туре	Storage Node		
Software Version	11.1.1 (build 20180814.2117.6c160fc)		
IP Addresses	172.16.5.24, 10.224.5.24, 47.47.5.24, 169.254.0.1 Show less •		
	Interface	IP Address	
	eth0	172.16.5.24	
	eth0	2001:aaaa:aaaa:0:526b:4bff.fe42:d711	
	eth0	fe80::526b:4bff:fe42:d711	
	eth1	10.224.5.24	
	eth1	fd20:8b1e:b255:8154:dac4:97ff:fe2a:e49e	
	eth1	fe80::dac4:97ff:fe2a:e49e	
	eth2	2001:ffff:ffff:0:526b:4bff:fe42:d711	
	eth2	47.47.5.24	
	eth2	fe80::526b:4bff:fe42:d711	
	hic1	172.16.5.24	
	hic1	2001:aaaa:aaaa:0:526b:4bff;fe42:d711	

## Related concepts

Managing networking on page 277

### **Related information**

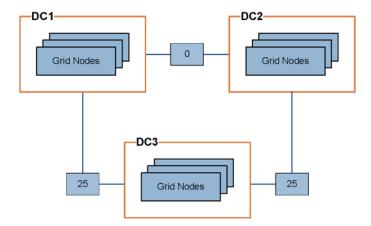
Recovery and maintenance

# 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 configure link costs to control how API Gateway Nodes direct traffic, and which replicated copy of an object is used to fulfill a retrieval.

- API 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, an API Gateway Node at data center site 1 (DC1) equally distributes client connections to Storage Nodes at DC1 and to Storage Nodes at DC2. An API 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.

# **Updating link costs**

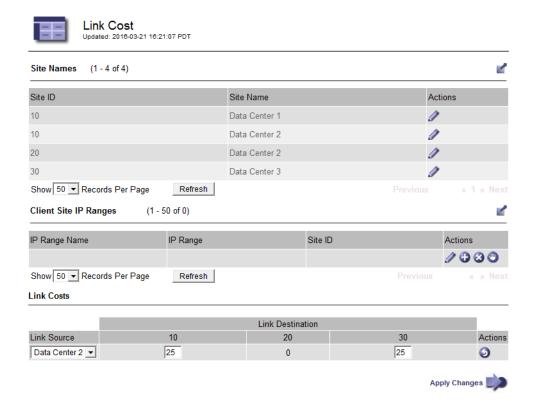
You can update the link costs between data center sites.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

## Steps

1. Select Configuration > Link Cost.



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.

# 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.

#### Before you begin

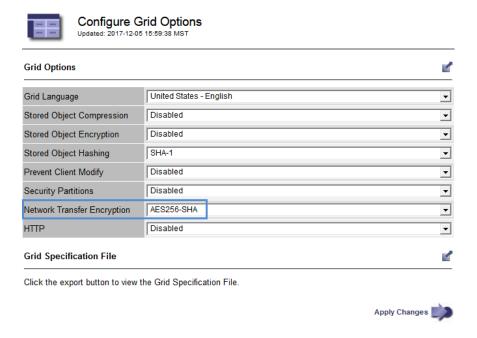
- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.

## About this task

By default, network transfer encryption uses the AES256-SHA algorithm. Control traffic can also be encrypted using the AES128-SHA algorithm.

- 1. Select Configuration > Grid Options.
- **2.** From the Grid Options menu, select **Configuration**.

3. Change Network Transfer Encryption to AES256-SHA or AES128-SHA.



4. Click Apply Changes.

# **Configuring certificates**

You can customize the certificates used by the StorageGRID system.

The StorageGRID system uses security certificates for two 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 API 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.

# Configuring custom server certificates for the Grid Manager and the Tenant Manager

You can replace the default StorageGRID server certificates with custom certificates that allows users to access the Grid Manager and the Tenant Manager without encountering security warnings.

#### About this task

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.

- 1. Select Configuration > 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).
  - **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.

**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 > 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 custom server certificates for storage API endpoints

You can replace the default object storage API service endpoint server certificates with a single custom server certificate that is specific to your organization.

#### About this task

API service endpoints on Storage Nodes are secured and identified by X.509 server certificates. By default, every Storage 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.

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 in the API client they will use to access the system.

- 1. Select Configuration > 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).

- **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 API client connections.

**5.** Refresh the page to ensure the web browser is updated.

# Restoring the default server certificates for storage API endpoints

You can revert to using the default server certificates for the storage API endpoints.

#### **Steps**

- 1. Select Configuration > 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 object storage API service endpoints 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 API client connections.

**4.** Refresh the page to ensure the web browser is updated.

# Copying the StorageGRID system's CA certificate

You can copy the StorageGRID system's certificate authority (CA) certificate from the StorageGRID system for client applications that require server verification. If a custom server certificate has been configured, then client applications should verify the server using the root CA certificate that issues the custom server certificate, rather than copy the CA certificate from the StorageGRID system.

## Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions. For details, see information about controlling system
  access with administration user accounts and groups.

#### **Steps**

- 1. Select Configuration > Grid Options.
- **2.** From the Grid Options menu, click **Overview**.
- 3. Under CA Certificates, expand CA Certificate.
- 4. Select the CA certificate.

Include the "----BEGIN CERTIFICATE----" and the "----END CERTIFICATE----" in your selection.

#### **Grid Information**

 Grid ID:
 674114

 Configured:
 2016-08-04 13:26:28 MDT

 Vendor:
 NetApp Inc.

 Software Suite Interoperability Version:
 10.3.0

#### **CA Certificates**

CA Certificate: ☐ CN=GPT, OU=NetApp StorageGRID, O=NetApp Inc., L=Sunnyvale, ST=California, C=US -BEGIN CERTIFICATE-MIIETjCCAzagAwIBAgIJAOeOVW5ik86sMAOGCSqGSIb3DQEBCwUAMHcxCzAJBgNV BAYTA1VTMRMwEQYDVQQIEwpDYWxpZm9ybmlhMRIwEAYDVQQHEw1TdW5ueXZhbGUx FDASBgNVBAoTC051dEFwcCBJbmMuMRswGQYDVQQLExJOZXRBcHAgU3RvcmFnZUdS  ${\tt SUQxDDAKBgNVBAMTA0dQVDAeFw0xNjA4MDQxOTIzNTRaFw0z0DAxMTcxOTIZNTRAFw0z0DAxMTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTRAFw0z0DAxWTcxOTIZNTFAFw0z0DAxWTcxOTIZNTFAFw0z0DAxWTcxOTIZNTFAFw0z0DAxWTcxOTIZNTFAFw0z0DAxWTcxOTIZNTFAFw0z0DAxWTcxOTIZNTFAFw0z0DAxWTcxOTIZNTFAFw0z0DAxWTcxOTIZN$  $\verb|MHCxCzAJBgNVBAYTA1VTMRMwEQYDVQQIEwpDYWxp2m9ybm1hMRIwEAYDVQQHEw1T| \\$ dW5ueXZhbGUxFDASBgNVBAoTC051dEFwcCBJbmMuMRswGQYDVQQLExJ0ZXRBcHAg U3RvcmFnZUdSSUQxDDAKBgNVBAMTA0dQVDCCASIwDQYJKoZIhvcNAQEBBQADggEP  $\verb|ADCCAQoCggEBAMCLfwGW5EnQutnWVGXDCvJ+Fpwao+8cD2rh09VHfSzfgDfBv7rA||$ aIHh01kBgPdSPzVWiNu1Kf7v9iyoq5Mu2mnc3FbbHIKUKywngk4ObuyijnZ8ww1X EHCssRjxvYc7vify+5VkaExi0FvuaKiji/9209sWStkAJbQsUi3WNitJkUP3jzYi DtJzCa6AuJ199RcqSNgrdtvVJPY1yao4mNaWV+06uHNMOTEnjaeNzFxH7ZxPQ+c7 dcM6qQIHF478Yo05uX1cXq+HQvswALh4tUAbNiaVKajDJk9wrVrpW7vtUa36IYW/ 6gNvWD5PY1XCQy9gWih9I06TRv7D99K6dZkCAwEAAaOB3DCB2TAdBgNVHQ4EFgQU TfxScVUbD1wvZvmZTVA4KecPzggwgakGA1UdIwSBoTCBnoAUTfxScVUbD1wvZvmZ TVA4KecPzgihe6R5MHcxCzAJBgNVBAYTA1VTMRMwEQYDVQQIEwpDYWxpZm9ybm1h  ${\tt MRIwEAYDVQQHEwlTdW5ueXZhbGUxFDASBgNVBAoTC051dEFwcCBJbmMuMRswGQYD}$ VQQLExJOZXRBcHAgU3RvcmFnZUdSSUQxDDAKBgNVBAMTA0dQVIIJA0e0VW5ik86s  ${\tt MAwGA1UdEwQFMAMBAf8wDQYJKoZIhvcNAQELBQADggEBAGBJUp9k4wvAGNXT/Pd5}$ LGBr6rsZzTLoKIm6cx2LcGzz+eVow48bjemhEcXYJhh4stv0yBlZqN9hGoObNBLQ 90vJ4vLMJ4BCvjBxSbOCtRS8oYdioCvBsmmfoAmwD6G8m+gwDLvVpvVeJZrmPtdz tRE6snEyP8Ee0RCizIJYMr0Gl7IHiDCBCot+PChpiRV2MIWkoHx5YRZ7CWnpFB0u WokFRF+3L3BL4JtCAe/kaR2/W5YAJUY/tlx3IbOit01HySFK7UKoJ+LruqAS8mvR ucNd5pnJJlNRnJxcLmyFiDcHSnrMzX+22/xUY5B/Xvm1rIjY01F01UYMDNpL+DhB

**5.** Right-click the selected certificate, and then select **Copy**.

# Configuring StorageGRID certificates for ONTAP clients using FabricPool

You can use a script to generate a self-signed server certificate for S3 clients that perform strict hostname validation and do not support disabling strict hostname validation, such as ONTAP clients using FabricPool. 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.

## Before you begin

- You have specific access permissions.
- You must have the Passwords.txt file.

- 1. Obtain the fully qualified domain name (FQDN) of each API Gateway 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.
- **3.** Configure StorageGRID with a new self-signed certificate.
  - \$ sudo make-certificate --domains wildcard-gateway-node-fqdn --type storage

- For --domains, use wildcards to represent the fully qualified domain names of all API Gateway Nodes. For example, \*.sgws.foo.com uses the \* wildcard to represent gnl.sgws.foo.com and gnl.sgws.foo.com.
- Set --type to storage to configure the certificate used by S3 and Swift storage clients.
- 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.

**Note:** A certificate's validity period begins when make-certificate is run. You must ensure the S3 client is synchronized to the same time source as StorageGRID; otherwise, the client might reject the certificate.

#### **Example**

```
$ sudo make-certificate --domains *.s3.example.com --type storage --
days 730
```

The resulting output contains the public certificate needed by your S3 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 > Object Storage API Service Endpoints Server Certificate.
- 7. Configure your S3 client to use the public certificate you copied. Include the BEGIN and END tags.

# Configuring S3 API endpoint domain names

To support S3 virtual hosted-style requests, you must configure the list of endpoint domain names that S3 clients will be connecting to.

# Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You must have specific access permissions. For details, see information about controlling system
  access with administration user accounts and groups.
- You must have confirmed that a grid upgrade is not in progress.

**Warning:** Do not make any changes to the domain name configuration when a grid upgrade is in progress.

#### About this task

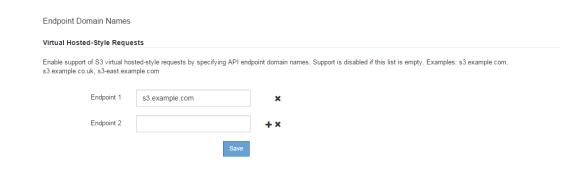
You configure API endpoint domain names after you create the fully qualified domain names on the DNS server, depending on the grid nodes that S3 clients will be connecting to:

- If S3 clients are connecting to one or more API Gateway Nodes, you must include the domain name of each API Gateway Node.
- If S3 clients are connecting to one or more Storage Nodes, you must include the domain name of each Storage Node.
- If S3 clients are connecting through an external load balancer, you must include the domain name of the load balancer.

## **Steps**

1. Select Configuration > Domain Names.

The Endpoint Domain Names page appears.



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. Obtain a custom server certificate with the wildcard Subject Alternative Name (SAN) for the endpoint domain name, the endpoint domain name, and any other domain names that must be supported.

This step is required to validate the SSL certificate and to verify the hostname when API client applications connect to the endpoint.

# **Example**

If the endpoint is s3.company.com, obtain a custom server certificate that includes the s3.company.com endpoint and the endpoint's wildcard SAN: \*.s3.company.com.

- 5. Select Configuration > Server Certificates. Then, install the custom certificate in the Object Storage API Service Endpoints Server Certificate section.
- **6.** Confirm that the DNS server also supports the endpoint and the wildcard SAN.

Now, when the endpoint bucket.s3.company.com is used, the DNS server resolves to the correct endpoint and the certificate authenticates the endpoint as expected.

#### Related tasks

#### Related information

Implementing S3 client applications

# **Configuring proxy settings**

If you are using platform services or Cloud Storage Pools, you can configure a non-transparent proxy to connect to external 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.

#### Before you begin

- You must have specific access permissions. For details, see information about controlling system
  access with administration user accounts and groups.
- You must be signed in to the Grid Manager using a supported browser.

#### About this task

You can configure the settings for a single proxy.

#### **Steps**

1. Select Configuration > Proxy Settings.

The Proxy Settings page appears.



- 2. Select the protocol for the non-transparent proxy.
- 3. Enter the hostname or IP address of the proxy server.
- **4.** 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.

5. Click Save.

After the proxy is saved, new endpoints for platform services or Cloud Storage Pools can be configured and tested.

**Note:** Proxy changes can take up to 10 minutes to take effect.

6. If you need to reset the proxy, click Clear.

## Related concepts

Considerations for Cloud Storage Pools on page 168

## 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.

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

#### Related concepts

What an Admin Node is on page 274

#### Related information

Understanding audit messages Upgrading StorageGRID

## 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.

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

#### Related information

Upgrading StorageGRID

#### **Configuring audit clients for Workgroup**

Perform this procedure for each Admin Node in a StorageGRID deployment from which you want to retrieve audit messages.

#### Before you begin

- 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. From the service laptop, 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 enable-disable-share add-user-to-share remove-user-from-share modify-group	set-authentication set-netbios-name join-domain add-password-server remove-password-server add-wins-server remove-wins-server	validate-config help exit

**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 step 6.

- 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:

workgroup\_name

Press Enter to use the Admin Node's host name 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

**Note:** 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.

**Note:** 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 step 7 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.** If the StorageGRID deployment is a single site, go to step *11*.

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 steps 4 through 9 to configure the audit share for each additional Admin Node.
- c. Close the remote secure shell login to the remote Admin Node: exit
- 11. Log out of the command shell: exit

#### Related information

Upgrading StorageGRID

#### Configuring audit clients for Active Directory

#### Before you begin

- 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).

#### About this task

Perform this procedure for each Admin Node in a StorageGRID deployment from which you want to retrieve audit messages.

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

#### **Steps**

- 1. From the service laptop, 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 enable-disable-share add-user-to-share remove-user-from-share	set-authentication set-netbios-name join-domain add-password-server	validate-config help exit

modify-group	remove-password-server	
	add-wins-server	
	remove-wins-server	
·		

#### 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 step 6.

- a. When prompted for Workgroup or Active Directory installation: ad
- b. When prompted, enter the name of the AD domain (short domain name).
- c. When prompted, enter the domain controller's IP address or DNS host name.
- d. When prompted, enter the full domain realm name.

Use uppercase letters.

- e. When prompted to enable winbind support, type y.
   Winbind is used to resolve user and group information from AD servers.
- f. When prompted, enter the NetBIOS name.
- g. When prompted, press Enter.The CIFS configuration utility is displayed.
- **6.** Join the domain:
  - a. If not already started, start the CIFS configuration utility: config\_cifs.rb
  - b. Join the domain: join-domain
  - **c.** 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**
  - **d.** When prompted, provide the Administrator's username: <code>administrator\_username</code> where <code>administrator\_username</code> is the CIFS Active Directory username, not the StorageGRID username.
  - e. When prompted, provide the Administrator's password: administrator\_password were administrator\_password is the CIFS Active Directory username, not the StorageGRID 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: adduser-to-share

A numbered list of enabled shares is displayed.

- **a.** Enter the number of the audit-export share.
- 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 step 9 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)

**Attention:** 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.
- 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 step *13*.

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 steps 4 through 11 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

Upgrading StorageGRID

#### 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.

#### Before you begin

- 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.

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

#### **Steps**

- 1. From the service laptop, 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  ${\tt Passwords.txt}$  file.

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

- Confirm that all services have a state of Running or Verified. Enter: storagegrid-statusIf 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 enable-disable-share add-user-to-share remove-user-from-share modify-group	set-authentication   set-netbios-name   join-domain   add-password-server   remove-password-server   add-wins-server   remove-wins-server	validate-config help exit

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 steps 5 to 8 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 step 14.
  - 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 steps 4 through 12 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.

#### Before you begin

- 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. From the service laptop, 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 enable-disable-share add-user-to-share remove-user-from-share modify-group	set-authentication set-netbios-name join-domain add-password-server remove-password-server add-wins-server remove-wins-server	validate-config     help   exit

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

Upgrading StorageGRID

#### 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 tasks

Adding a user or group to a CIFS audit share on page 296 Removing a user or group from a CIFS audit share on page 298

#### Related information

Upgrading StorageGRID

#### 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.

#### Before you begin

- 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. From the service laptop, 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 #.
- Confirm that all services have a state of Running or Verified. Enter: storagegrid-statusIf 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 enable-disable-share	add-ip-to-share remove-ip-from-share	validate-config   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 step 6 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 step 9.
  - 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 steps 4 through 7.c 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.

#### Before you begin

- 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. From the service laptop, 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-ip-to-share remove-ip-from-share	validate-config   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 from step 3 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 step *11*.

— 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 steps 2 through 9 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: 1s 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.

#### Before you begin

- 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. From the service laptop, 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     enable-disable-share	add-ip-to-share remove-ip-from-share	validate-config     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 steps 2 through 7 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

#### Steps

- 1. Add a new IP address to an existing NFS audit share.
- 2. Remove the original IP address.

#### Related tasks

Adding an NFS audit client to an audit share on page 301
Removing an NFS audit client from the audit share on page 303

## What data migration is

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 its day-to-day operations, and that the migrated data is handled appropriately by the StorageGRID system.

## 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 storage capacity and managing disk storage.

#### Related concepts

Managing disk storage on page 236

#### Related tasks

Monitoring storage capacity for the entire grid on page 100

## 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.

**Caution:** 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 concepts

What an information lifecycle management policy is on page 150

#### Related tasks

Configuring ILM rules on page 157

## 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. Controlling the rate of migration of data into the system is outside of the scope of StorageGRID functionality.

## 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 concepts

Monitoring data migration on page 307

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.

Monitor	Description
Number of objects waiting for ILM	1. Select Support > Grid Topology.
evaluation	2. Select deployment > Overview > Main.
	<b>3.</b> In the ILM Activity section, monitor the number of objects shown for the following attributes:
	Awaiting - All (XQUZ): The total number of objects awaiting ILM evaluation.
	<ul> <li>Awaiting - Client (XCQZ): The total number of objects awaiting ILM evaluation from client operations (for example, ingest).</li> </ul>
	<b>4.</b> 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.
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.
Archive Node > ARC > Store > Store Failures (ARVF)	If an alarm for this 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 to configure the StorageGRID system to send a notification email to the system administrator responsible for monitoring migration if the attribute values exceed their recommended maximum values.

#### Before you begin

- You must be signed in to the Grid Manager using a supported browser.
- You have specific access permissions.
- You must have configured email settings.
- You must have a mailing list.

#### **Steps**

1. Create an email list that includes all administrators responsible for monitoring the data migration.

Optionally, you can create a template to customize the subject line, header, and footer of data migration notification emails.

- 2. Create a Global Custom alarm for each attribute you need to monitor during data migration.
  - a. Select Configuration > Global Alarms.
  - b. Under Default Alarms, search for the Default alarms for the first attribute. Under Filter by, select Attribute Code, then type the four letter code for the attribute. For example, ARVF.
  - c. Click **Submit**
  - d. In the results list, click **Copy** next to the alarm you want to modify.

The alarm moves to the Global Custom Alarms table.

- e. Under Global Custom Alarms, in the Mailing List column for the copied attribute, add the mailing list.
- f. Repeat for each remaining attribute.
- g. When finished creating Global Custom alarms, click Apply Changes.

#### After you finish

Administrators responsible for monitoring data migration now receive an email notification if the values of key attributes exceed their maximum acceptable levels during migration.

Remember to disable these notifications after data migration is complete. Note that Global Custom alarms override Default alarms. If there are any, enable Custom alarms at the grid node level as Global Custom alarms cannot be triggered.

#### Related tasks

Configuring email server settings on page 114
Creating mailing lists on page 116

## 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**

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 than 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.

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 http://www.ibm.com/support/docview.wss?rs=577&uid=ssg1S7002972
- IBM Tape Device Drivers Programming Reference http://www.ibm.com/support/docview.wss?rs=577&uid=ssg1S7003032

#### Related concepts

Managing archival storage on page 259

### 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.

Note: 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

#### 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 you site may be different depending on the requirements of your deployment. For configuration details and for instructions, see the TSM documentation.

**Note:** You must log onto the server with administrative privileges and use the dsmadmc 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.

 ${\tt define\ path\ } server name\ tapelibrary\ {\tt srctype=server\ } desttype=library\ device=lib-device name$ 

- 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=1to library=tapelibrary format=tapetype

- DeviceClassName is the name of the device class
- 1to 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 maxscratch=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.

**Note:** 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 tsmdomain 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).

**Note:** You must log in to the TSM server with administrative privileges and use the dsmadmc 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:

```
{\tt 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

**Note:** 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.

**Note:** By default, registering a node creates an administrative user ID with client owner authority, with the password defined for the node.

## Copyright

Copyright © 2019 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, 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

## **Trademark**

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.

http://www.netapp.com/us/legal/netapptmlist.aspx

# How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

#### doccomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

• NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.

• Telephone: +1 (408) 822-6000

• Fax: +1 (408) 822-4501

• Support telephone: +1 (888) 463-8277